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

Prepare for your AZ-204 exam with additional products

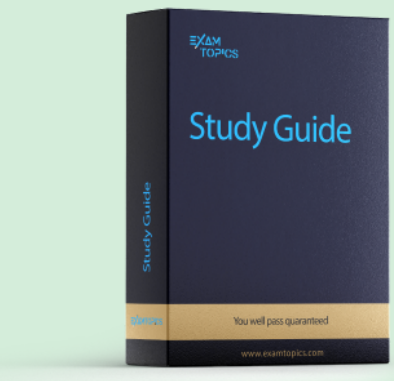

Study Guide

289 PDF Pages

Download Now

Video Course

162 Lectures

$19.99

Buy Now

Custom View Settings

Topic 1 - Question Set 1

Question #1 *Topic 1*

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

A. From the Update management blade, click Enable.

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

C. From the Redeploy blade, click Redeploy.

D. From the Profile blade, modify the usage location.

Correct Answer: *C*
When you redeploy a VM, it moves the VM to a new node within the Azure infrastructure and then powers it back on, retaining all your configuration options and associated resources.
Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/redeploy-to-new-node

Community vote distribution
C (100%)

adilkhan Highly Voted 6 months, 1 week ago
Just passed the exam 26/04/2023. All the questions were from Exam topics. Got 970 marks!
upvoted 14 times

collocollo Most Recent 1 month, 2 weeks ago
Selected Answer: C
C - ok
upvoted 1 times

ndoan 2 months ago
agree! C is the right answer
upvoted 2 times

notforme 2 months ago
i gave exam today, only 4 questions out of 15 pages of questions i went on this site
upvoted 2 times

Dianahu 2 months, 2 weeks ago
Just passed the exam with 947, all questions were from examtopics az-204 except this:
https://www.examtopics.com/discussions/microsoft/view/46665-exam-az-204-topic-5-question-4-discussion/
I had 1 case study - City Power & Light . GL!
upvoted 1 times

Utkarsh_ 1 week ago
https://www.examtopics.com/discussions/microsoft/view/46665-exam-az-204-topic-5-question-4-discussion/
upvoted 1 times

Dianahu 2 months, 2 weeks ago
Sry I meant I did not see this before exam :D
upvoted 1 times

ITRogers 3 months, 2 weeks ago
C is the right answer
upvoted 1 times

cmil30 4 months, 2 weeks ago
Go through all 34 pages and you should be good. Passed the exam on 18/06/2023.
upvoted 2 times

NavinD 4 months ago

Does it require contributor access to be bought for going through 34 pages? Do they have repeat questions in the paid portion or there are some undisclosed questions also as I am able to access only 50% of pages for free?

upvoted 2 times

vikasm361 2 months, 2 weeks ago

yup. you need to get contributor access

upvoted 2 times

torsanta 5 months ago

Selected Answer: C

Chosen Answer: C

upvoted 1 times

Kallem_Niharika 5 months ago

C is the correct answer

upvoted 1 times

nethala_swathi 5 months ago

Option:C

upvoted 1 times

D3D1997 5 months, 4 weeks ago

Selected Answer: C

redeploy

upvoted 2 times

SmartTechieGrv 6 months, 3 weeks ago

C is the correct ans

upvoted 1 times

Judy7 7 months, 1 week ago

Is this question not more for an admin rather than developer for AZ 204 exam?

upvoted 2 times

sbsd 8 months, 1 week ago

Option C.

If you have been facing difficulties troubleshooting Remote Desktop (RDP) connection or application access to Windows-based Azure virtual machine (VM), redeploying the VM may help. When you redeploy a VM, Azure will shut down the VM, move the VM to a new node within the Azure infrastructure, and then power it back on, retaining all your configuration options and associated resources.

https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/redeploy-to-new-node-windows

upvoted 1 times

[Removed] 8 months, 1 week ago

Option C.

If you have been facing difficulties troubleshooting Remote Desktop (RDP) connection or application access to Windows-based Azure virtual machine (VM), redeploying the VM may help. When you redeploy a VM, Azure will shut down the VM, move the VM to a new node within the Azure infrastructure, and then power it back on, retaining all your configuration options and associated resources.
https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/redeploy-to-new-node-windows

upvoted 1 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: C

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

upvoted 1 times

uffuchsi 9 months, 2 weeks ago

Selected Answer: C

Redeploy blade

upvoted 2 times

Question #2 *Topic 1*

DRAG DROP -

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

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

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

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

Select and Place:

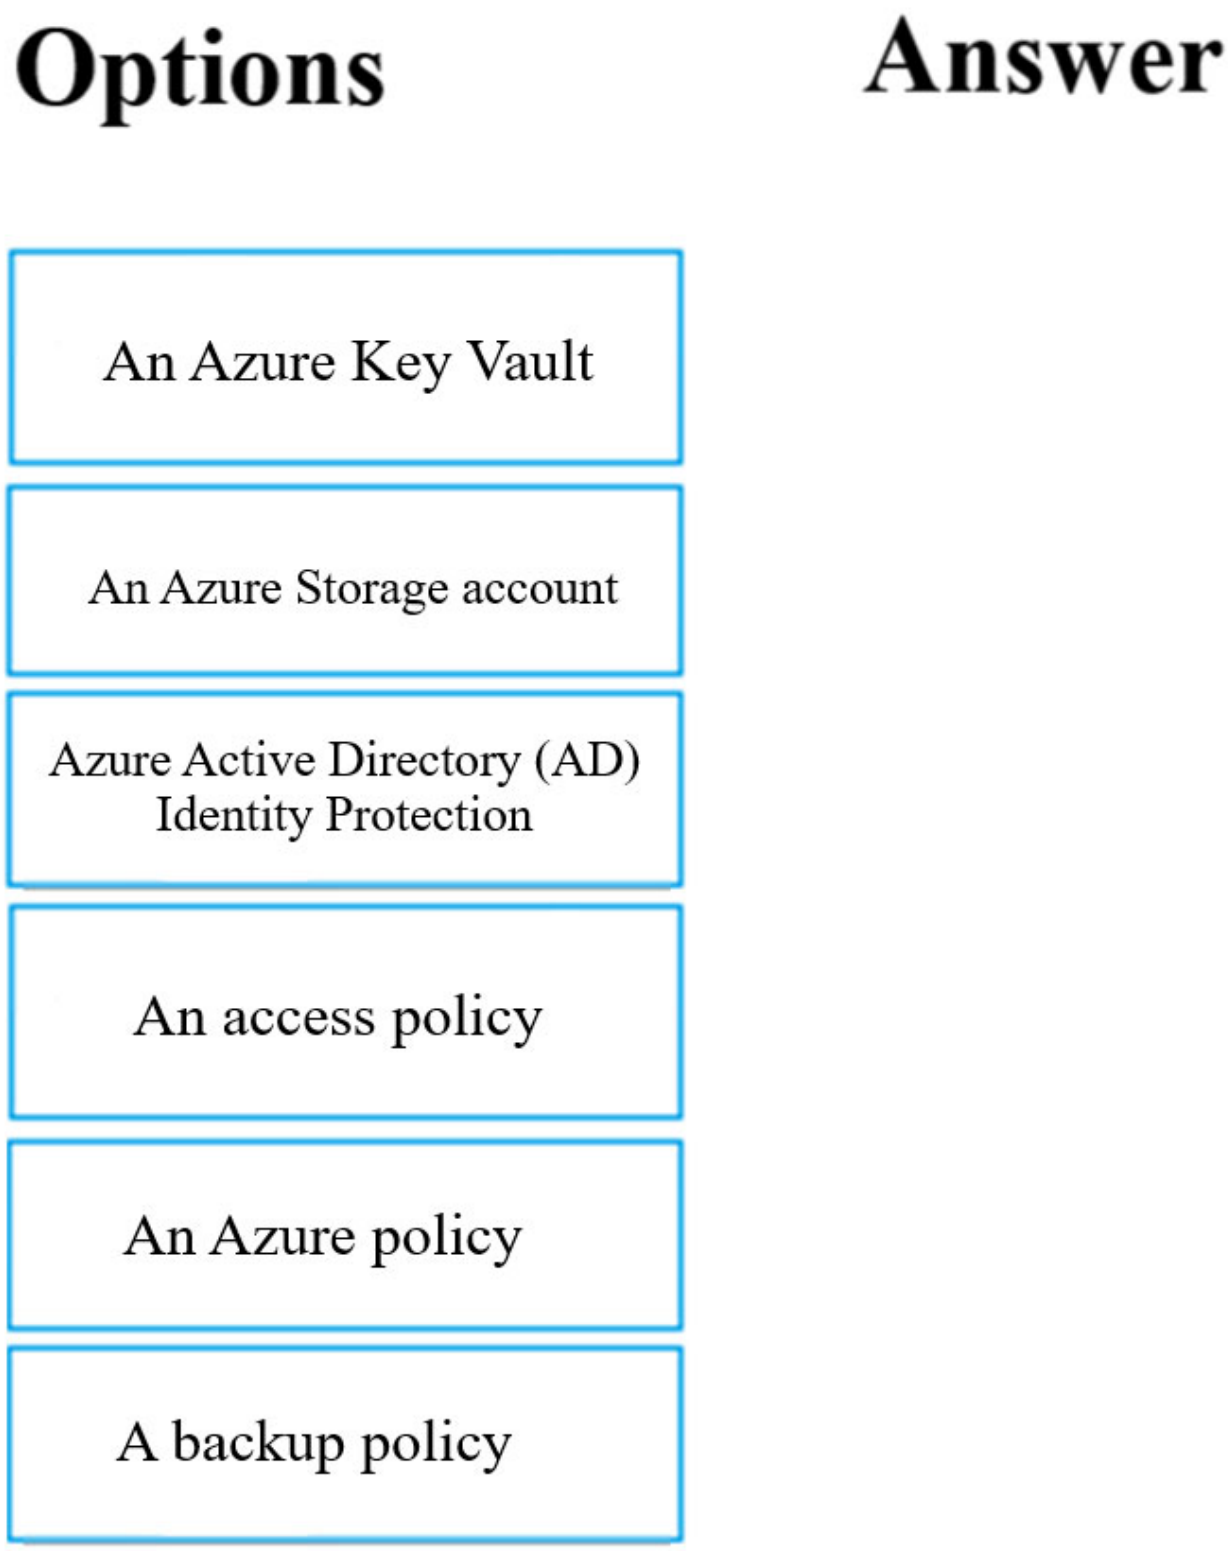

Correct Answer:

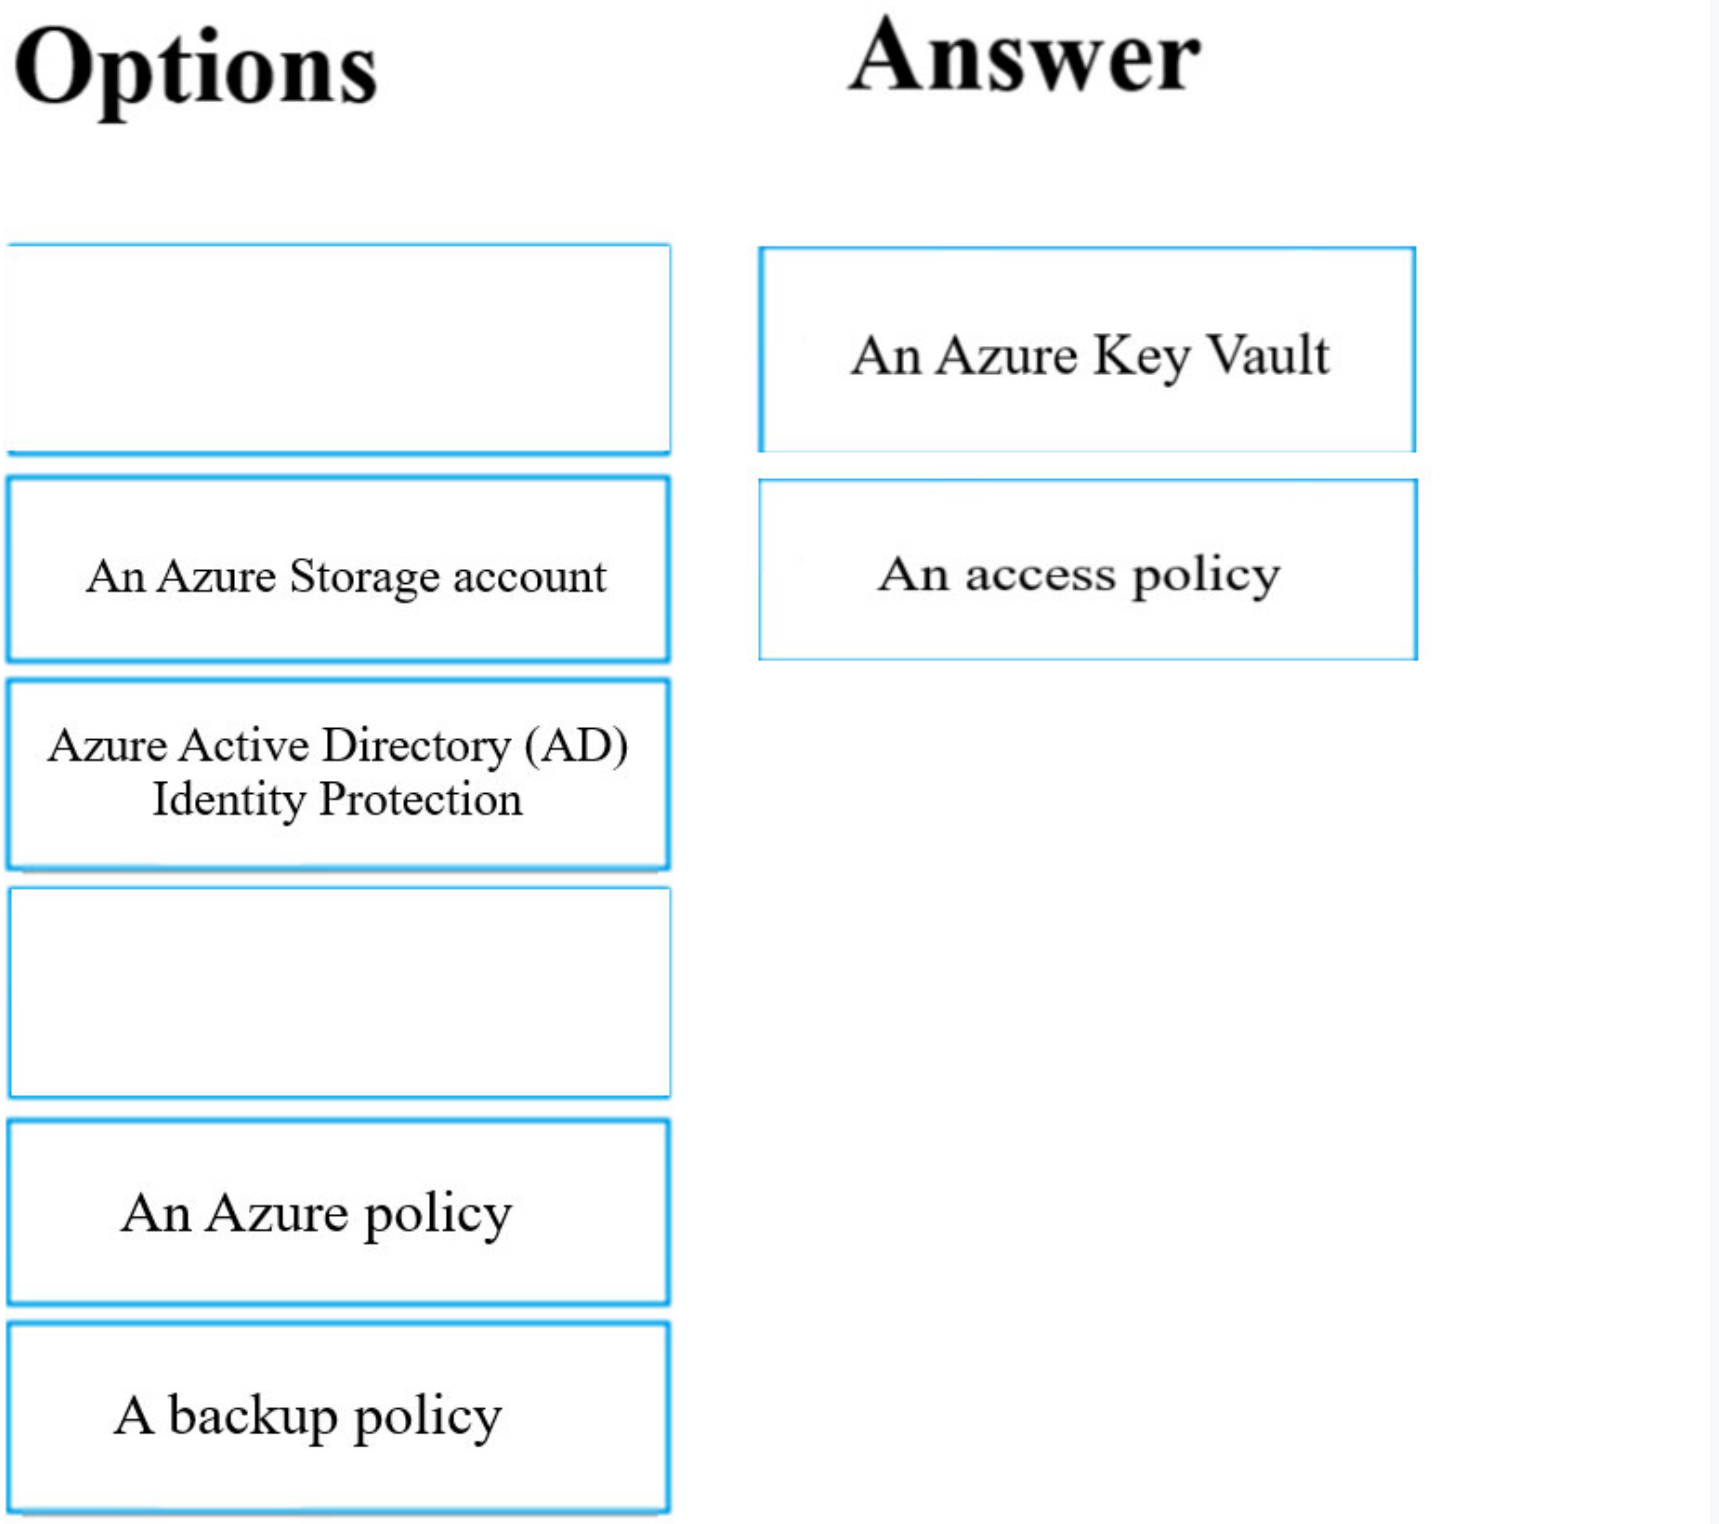

vincyN90 Highly Voted 1 year, 1 month ago

Key Vault + Access Policy.

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

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

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

upvoted 26 times

Obioha Highly Voted 1 year, 6 months ago

https://www.examtopics.com/exams/microsoft/az-104/view/2/

Answer is correct, see question in AZ104 #20

upvoted 5 times

dimipap3 Most Recent 6 days, 13 hours ago

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

upvoted 1 times

SmartTechieGrv 6 months, 3 weeks ago

Key Vault and Access Policy both

upvoted 2 times

moned 6 months, 4 weeks ago

Is it me, or is the question misleading since it says "option" - singuar?

upvoted 3 times

Rene_Clevers 6 months ago

I had the same problem
upvoted 1 times

Judy7 7 months, 1 week ago
This topic says AZ-204 but the questions seem to be from Az-104 which is the admin. This is not correct?
upvoted 1 times

p2006 8 months, 1 week ago
https://learn.microsoft.com/en-us/azure/key-vault/general/assign-access-policy
upvoted 1 times

forgetfulalligator 8 months, 1 week ago
Azure Key Vault is the recommended approach for securely storing and managing secrets such as the administrative password used in an Azure Resource Manager template. However, an access policy can also be used to control access to a secret stored in a different Azure resource, such as an Azure Storage account.
upvoted 1 times

uffuchsi 8 months, 2 weeks ago
Key Vault & Access Policy
upvoted 1 times

Santmax 9 months, 2 weeks ago
Key Vault + Access Policy.
upvoted 1 times

Question #3 *Topic 1*

Your company has an Azure Kubernetes Service (AKS) cluster that you manage from an Azure AD-joined device. The cluster is located in a resource group.
Developers have created an application named MyApp. MyApp was packaged into a container image.
You need to deploy the YAML manifest file for the application.
Solution: You install the Azure CLI on the device and run the kubectl apply `"f myapp.yaml command.
Does this meet the goal?

A. Yes

B. No

Correct Answer: *A*
kubectl apply -f myapp.yaml applies a configuration change to a resource from a file or stdin.
Reference:
https://kubernetes.io/docs/reference/kubectl/overview/
https://docs.microsoft.com/en-us/cli/azure/aks

Community vote distribution
A (81%) B (19%)

Nokaido Highly Voted 1 year, 1 month ago

Selected Answer: A
- apears as `" at examtopics
upvoted 9 times

jerrychan Most Recent 2 weeks, 2 days ago

reference:
https://www.devspace.sh/docs/configuration/deployments/kubectl/inline_manifests
kubectl apply -f ...
upvoted 1 times

RikinPatel 3 weeks, 2 days ago

Ans.: B
As per my understanding. step as follow as per microsoft..
- Create a resource group
- Create an AKS cluster
- Connect to the cluster
- Deploy the application

https://learn.microsoft.com/en-us/azure/aks/learn/quick-kubernetes-deploy-cli
Missting step: Connect to the cluster
upvoted 1 times

MikeAWS 1 month ago

A should be correct!
kubectl apply -f myapp.yaml
upvoted 1 times

collocollo 1 month, 2 weeks ago

Selected Answer: A
A - ok
upvoted 1 times

viking2 4 months ago

Selected Answer: A
As the command is written, the answer would be NO, as that is invalid syntax for kubectl.

However, assuming that this is just another transcription issue on ExamTopics, and that `" should be a - instead, the answer is A.
upvoted 3 times

LuchoGonzalez 4 months, 2 weeks ago

Chat gpt said Yes, and then give me the instruccion to achieve it:
Yes, your solution is correct. To deploy the YAML manifest file for the application on your Azure Kubernetes Service (AKS) cluster, you can follow these steps:

Install the Azure CLI on your Azure AD-joined device. You can download and install the Azure CLI from the official Microsoft Azure CLI

documentation.

Once the Azure CLI is installed, open a terminal or command prompt.

Log in to your Azure account by running the following command and following the prompts:...
upvoted 3 times

felipementel 6 months, 1 week ago
Install Azure CLI cannot enable you to use kubectl commands. For it, you need to install Kubernetes command-line tool (kubectl) - my opnion, B!
upvoted 1 times

LuchoGonzalez 4 months, 2 weeks ago
Chat gpt said Yes, and then give me the instruccion to achieve it:
Yes, your solution is correct. To deploy the YAML manifest file for the application on your Azure Kubernetes Service (AKS) cluster, you can follow these steps:

Install the Azure CLI on your Azure AD-joined device. You can download and install the Azure CLI from the official Microsoft Azure CLI documentation.

Once the Azure CLI is installed, open a terminal or command prompt.

Log in to your Azure account by running the following command and following the prompts:...
upvoted 1 times

SmartTechieGrv 6 months, 3 weeks ago
A is coorect
upvoted 1 times

vmakharashvili 7 months, 3 weeks ago
Correct answer is A
upvoted 1 times

Ayman99 8 months ago
For anyone wondering, most of the questions (in my case all the questions) are from exam topics. I just took the exam 27.02.2023 and passed with 870 score.
upvoted 2 times

VVVVVVSSSSS 5 months ago
Is contributor access necessarily required?
upvoted 1 times

necrodead93 7 months, 2 weeks ago
Did you study all the questions?
upvoted 2 times

forgetfulalligator 8 months, 1 week ago
Selected Answer: A
The "kubectl apply -f" command is used to deploy Kubernetes manifests, including YAML files, to a Kubernetes cluster
upvoted 4 times

uffuchsi 8 months, 2 weeks ago
A - Yes
upvoted 1 times

guschess 9 months ago
Connect to the cluster
To manage a Kubernetes cluster, use the Kubernetes command-line client, kubectl. kubectl is already installed if you use Azure Cloud Shell.

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

Azure CLI
- az aks install-cli

Why answer A is correct?
url: https://learn.microsoft.com/en-us/azure/aks/learn/quick-kubernetes-deploy-cli

thanks
upvoted 2 times

examinator007 9 months ago
AKS is not part of the AZ-204 - see: https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B (But ACI is)
upvoted 2 times

Net_IT 8 months, 2 weeks ago
Probably not but welcome to MS exams where everything and anythin is supposed to be known in order to pass. I'm starting to believe taht MS assumes you read all of the documentation regarding any subject they mention in the lesson material, e.g. Kubernetes in this case.
upvoted 1 times

Yolaw 9 months, 1 week ago

answer is A
https://learn.microsoft.com/en-us/azure/aks/learn/quick-kubernetes-deploy-cli

upvoted 1 times

Santmax 9 months, 2 weeks ago

Yes -answer is A

upvoted 1 times

Question #4 *Topic 1*

Your company has an Azure Kubernetes Service (AKS) cluster that you manage from an Azure AD-joined device. The cluster is located in a resource group.
Developers have created an application named MyApp. MyApp was packaged into a container image.
You need to deploy the YAML manifest file for the application.
Solution: You install the docker client on the device and run the docker run -it microsoft/azure-cli:0.10.17 command.
Does this meet the goal?

A. Yes

B. No

Correct Answer: *B*

Community vote distribution

B (100%)

jay158 Highly Voted 2 years, 3 months ago

Correct Answer NO

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

upvoted 13 times

lugospod Highly Voted 1 year, 10 months ago

K8 Out of scope for 2021!

upvoted 9 times

collocollo Most Recent 1 month, 2 weeks ago

Selected Answer: B

B - No

upvoted 1 times

viking2 4 months ago

Selected Answer: B

NO, use kubectl

upvoted 1 times

SmartTechieGrv 6 months, 3 weeks ago

Correct Answer No

upvoted 1 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: B

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

upvoted 2 times

uffuchsi 8 months, 2 weeks ago

B - No

upvoted 1 times

guschess 9 months ago

Selected Answer: B

The answer 3 helps with this topic.

upvoted 1 times

ExamTopicsAnkur 10 months, 3 weeks ago

Selected Answer: B

Answer is No

upvoted 2 times

Chouffie 1 year ago

The Answer is NO
To deploy a YAML file to Kubernetes you need to use the “kubectl” command. An example could be: kubectl apply -f myapp.yaml

upvoted 5 times

HafizSalmanMalik 1 year ago

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

upvoted 2 times

Mev4953 1 year, 9 months ago

Any reference?

upvoted 2 times

Question #5 *Topic 1*

Your company has a web app named WebApp1.
You use the WebJobs SDK to design a triggered App Service background task that automatically invokes a function in the code every time new data is received in a queue.
You are preparing to configure the service processes a queue data item.
Which of the following is the service you should use?

A. Logic Apps

B. WebJobs

C. Flow

D. Functions

Correct Answer: *B*
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-compare-logic-apps-ms-flow-webjobs

Community vote distribution

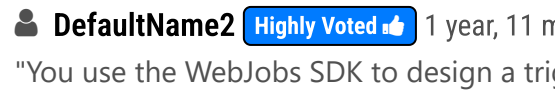

DefaultName2 Highly Voted 1 year, 11 months ago

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

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

upvoted 19 times

coffecold 1 year ago

Azure WebJobs is not a service, you'll host the WebJobs SDK functions in Azure WebApps.
With Webjobs you program code in your webapp and you'll be able to execute 'kind of Azure functions' inside your Web App. It is automatically deployed with the deployment of the web app. WebJobs is the predecessor of Azure functions.

upvoted 5 times

tracy161 9 months, 4 weeks ago

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

upvoted 1 times

coffecold 1 year ago

But still the best answer is B.

upvoted 1 times

eazybrave 1 year, 6 months ago

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

upvoted 3 times

Evo_Morales 1 year, 6 months ago

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

upvoted 1 times

noip Highly Voted 2 years, 2 months ago

the Correct Answer, D: Function

upvoted 9 times

edengoforit 1 year, 9 months ago

Without any rationale?

upvoted 5 times

person12345 Most Recent 4 days, 22 hours ago

There an issue with the question. There is no mention of the OS where the app of the question is running. In the documentation microsoft specifies that: "WebJobs aren't supported for App Service on Linux yet"
https://learn.microsoft.com/en-us/azure/app-service/webjobs-create

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

upvoted 1 times

ConanBarb 1 month ago

Selected Answer: D

D. The question is ambiguously phrased, but clearly (to me) they ask for an Azure Service that will replace the WebJobs SDK on Web Apps.

If nothing else WebJobs SDK is not a service and hence cannot be B.

upvoted 1 times

Aneeqa 1 month, 1 week ago

Azure Functions is another Azure service that allows you to run code in the cloud without having to worry about infrastructure. It is designed to run small pieces of code called functions that can be triggered by various events such as changes to data in Azure Storage, a new message in a Service Bus queue, or a new file in an Azure Blob container. Although Azure Functions can be used to process queue data items, it is not specifically designed for this purpose.

On the other hand, WebJobs is an Azure feature that allows you to run a program or script in the same context as a web app, API app, or mobile app. It is designed to run background tasks that can be invoked on a schedule, on demand, or continuously3. In your case, you are preparing to configure the service to process a queue data item. To achieve this goal, you should use WebJobs.

upvoted 1 times

MA17 3 months, 1 week ago

B- Webjobs
The WebJobs SDK makes it easy to write, run, and monitor such background tasks in Azure.
Azure function can do the same task but they are not part of the Azure app service and they're not considered as a background task.

upvoted 2 times

UkadikeToch 4 months, 2 weeks ago

The closest answer is B

upvoted 1 times

darsak10 5 months, 3 weeks ago

Selected Answer: B

Option B seems to be the correct answer. Webjob SDK can be used to automatically invoke a function every time new data is received in the queue.

upvoted 1 times

tony123 8 months, 1 week ago

Selected Answer: B

Webjobs run under Azure webapps

upvoted 1 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: B

My guts are saying B. Honestly it could either be WebJobs and Functions. The scenario mentions that the web app "WebApp1" is already being used and that the webjobs SDK is being used to design a triggered app service background task. Since WebJobs are designed to work specifically with Azure Web Apps, they provide a convenient way to run background tasks within the context of a web app

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

B. WebJobs or D, Functions

upvoted 1 times

fkaracan 8 months, 2 weeks ago

Selected Answer: B

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

upvoted 1 times

AzSimuz 8 months, 2 weeks ago

Functions and WebJobs both suits the need.
To me the answer is Function because when it comes to the most convenient one. Usually Function are more coast effective and according to the documentation "Azure Functions offers more developer productivity than Azure App Service WebJobs does. It also offers more options for programming languages, development environments, Azure service integration, and pricing. For most scenarios, it's the best choice."

upvoted 1 times

guschess 9 months ago

Selected Answer: B

The Azure WebJobs SDK is a framework that simplifies the task of writing background processing code that runs in Azure. The Azure WebJobs SDK includes a declarative binding and trigger system that works with Azure Storage Blobs, Queues and Tables as well as Service Bus.

URL: https://github.com/Azure/azure-webjobs-sdk

upvoted 1 times

obeesubz5 9 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/azure/app-service/webjobs-sdk-get-started

Just read first paragraph in that link "Get started with the Azure WebJobs SDK for Azure App Service to enable your web apps to run background tasks, scheduled tasks, and respond to events.

Use Visual Studio 2022 to create a .NET Core console app that uses the WebJobs SDK to respond to Azure Storage Queue messages, run the project locally, and finally deploy it to Azure."

upvoted 3 times

N9 9 months, 3 weeks ago

Selected Answer: B

You can use the Azure WebJobs SDK with WebJobs to simplify many programming tasks.

upvoted 1 times

armpro 9 months, 3 weeks ago

Selected Answer: B

B - WebJob
WebJob SDK is used to create web jobs. Below is link for exact queue processing scenario using WebJobs SDK
https://learn.microsoft.com/en-us/azure/app-service/webjobs-sdk-get-started

upvoted 2 times

Question #6 *Topic 1*

Your company has an Azure subscription.
You need to deploy a number of Azure virtual machines to the subscription by using Azure Resource Manager (ARM) templates. The virtual machines will be included in a single availability set.
You need to ensure that the ARM template allows for as many virtual machines as possible to remain accessible in the event of fabric failure or maintenance.
Which of the following is the value that you should configure for the platformFaultDomainCount property?

A. 10

B. 30

C. Min Value

D. Max Value

Correct Answer: *D*
The number of fault domains for managed availability sets varies by region - either two or three per region.
Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/manage-availability

Community vote distribution
D (100%)

jay158 Highly Voted 2 years, 3 months ago
Answer D Max Value. [values could be 1,2,3]

https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-manage-fault-domains
upvoted 20 times

argoth 2 years, 3 months ago
Correct. Values can be 1,2 or 3
upvoted 5 times

Bere Highly Voted 1 year, 11 months ago
As described here:
https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-manage-fault-domains
You can set the parameter --platform-fault-domain-count to 1, 2, or 3 (default of 3 if not specified).

And as described here:
https://docs.microsoft.com/en-us/azure/virtual-machines/availability-set-overview
Each virtual machine in your availability set is assigned an update domain and a fault domain by the underlying Azure platform. Each availability set can be configured with up to three fault domains and twenty update domains.

So answer is D Max Value
upvoted 7 times

nanotower Most Recent 4 weeks, 1 day ago
Vm are not in the exam anymore as of summer 2023.
upvoted 1 times

TheMCT 1 month, 3 weeks ago
Selected Answer: D
Azure allows you to specify a platform fault domain count of up to 3, which means that there can be up to three fault domains in an availability set. This setting helps distribute the virtual machines across multiple fault domains, reducing the risk of simultaneous failure due to hardware issues. Therefore, you should set it to the maximum value of 3 for optimal fault tolerance.
upvoted 1 times

SmartTechieGrv 6 months, 3 weeks ago
Answer is D
upvoted 1 times

p2006 8 months ago
https://learn.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-manage-fault-domains#azure-cli
upvoted 1 times

forgetfulalligator 8 months, 1 week ago
Selected Answer: D

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

upvoted 3 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 1 times

Santmax 9 months, 2 weeks ago

D Max value

upvoted 1 times

kkstays 10 months ago

Selected Answer: D

Max value is Correct

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: D

Answer D Max Value

upvoted 1 times

braj1234 1 year, 1 month ago

Answer is D

upvoted 1 times

mimi21212152 1 year, 11 months ago

az104 question, max value is correct.

upvoted 1 times

ranjitklive 2 years, 2 months ago

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

upvoted 5 times

Kalaisuran 2 years, 4 months ago

https://www.thatlazyadmin.com/2017/10/09/azure-fault-update-domains/

upvoted 2 times

akh13 9 months, 2 weeks ago

"A fault domain is a logical group of underlying hardware that share a common power source and network switch, similar to a rack within an on-premises datacenter."
So for maximum reliability of your availability set you want the maximum fault domains (which is 3 for a region with availability zones and 5 for a small region with just one availability zone). -> Answer D

upvoted 3 times

Question #7 *Topic 1*

Your company has an Azure subscription.
You need to deploy a number of Azure virtual machines to the subscription by using Azure Resource Manager (ARM) templates. The virtual machines will be included in a single availability set.
You need to ensure that the ARM template allows for as many virtual machines as possible to remain accessible in the event of fabric failure or maintenance.
Which of the following is the value that you should configure for the platformUpdateDomainCount property?

A. 10

B. 20

C. 30

D. 40

Correct Answer: *D*
Each virtual machine in your availability set is assigned an update domain and a fault domain by the underlying Azure platform. For a given availability set, five non-user-configurable update domains are assigned by default (Resource Manager deployments can then be increased to provide up to 20 update domains) to indicate groups of virtual machines and underlying physical hardware that can be rebooted at the same time.
Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/manage-availability

Community vote distribution
B (100%)

jay158 Highly Voted 2 years, 3 months ago
Answer is B 20
Each availability set can be configured with up to three fault domains and twenty update domains.
https://docs.microsoft.com/en-us/azure/virtual-machines/availability-set-overview
upvoted 45 times

ndh103 Highly Voted 2 years, 4 months ago
Should be 20
upvoted 10 times

regux Most Recent 2 weeks ago
Selected Answer: B
B 20 is the correct one.
upvoted 1 times

Kopparthi 2 months, 1 week ago
A higher value would generally be preferable to improve availability. Among the options, D. 40 would provide the highest degree of isolation and availability during maintenance or updates. Therefore, you should configure the platformUpdateDomainCount property with a value of 40
upvoted 1 times

Kopparthi 2 months, 1 week ago
A higher value would generally be preferable to improve availability. Among the options, D. 40 would provide the highest degree of isolation and availability during maintenance or updates. Therefore, you should configure the platformUpdateDomainCount property with a value of 40
upvoted 1 times

Kopparthi 2 months, 1 week ago
A higher value would generally be preferable to improve availability. Among the options, D. 40 would provide the highest degree of isolation and availability during maintenance or updates. Therefore, you should configure the platformUpdateDomainCount property with a value of 40
upvoted 1 times

CloudixExamTopics 2 months, 2 weeks ago
To ensure high availability within an Azure availability set, you need to distribute virtual machines across different update domains and fault domains. Update domains represent groups of virtual machines that can undergo maintenance or updates together, minimizing potential downtime.

For the best balance between high availability and maintenance, it's recommended to set the platformUpdateDomainCount property to 20. This ensures that your virtual machines are spread across 20 update domains. The platform imposes a default limit of 20 update domains in each availability set.
upvoted 1 times

kamalpur 3 months, 2 weeks ago

Explained the concept of availability sets, fault domain, and update domain in azure.
https://www.youtube.com/watch?v=9UX7Qq0u2oU

upvoted 2 times

Anjan1980 4 months, 1 week ago

Selected Answer: B

The correct answer is B. 20.

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

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

upvoted 1 times

Abhinav1 6 months, 1 week ago

Selected Answer: B

"Each availability set can be configured with up to three fault domains and twenty update domains."
source: https://learn.microsoft.com/en-us/azure/virtual-machines/availability-set-overview

upvoted 1 times

Dinya_jui 6 months, 2 weeks ago

3 fault domain
20 update domain

upvoted 1 times

obhanvarun 6 months, 3 weeks ago

To ensure that the maximum number of virtual machines remain accessible during fabric updates or maintenance, you need to configure the platformUpdateDomainCount property to a value that is equal to or greater than the number of virtual machines you want to deploy.

The platformUpdateDomainCount property specifies the number of update domains used by the availability set. An update domain is a group of virtual machines and underlying physical hardware that can be updated or rebooted at the same time. By spreading virtual machines across multiple update domains, you can ensure that only a subset of virtual machines are impacted by any given update or maintenance event.

Therefore, to deploy as many virtual machines as possible while ensuring high availability, you should configure the platformUpdateDomainCount property to a value of 40, which is the maximum allowed value for this property. This ensures that the virtual machines are spread across 40 update domains, providing maximum redundancy and availability during fabric updates and maintenance.

So the answer is D. 40.

upvoted 2 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: B

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

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

upvoted 1 times

NK203 8 months, 1 week ago

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

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

B. 20, since update domain counts max is 20

upvoted 1 times

fkaracan 8 months, 2 weeks ago

Selected Answer: B

To ensure that as many virtual machines as possible remain accessible during maintenance, you should configure the value of the platformUpdateDomainCount property to be at least 5. The maximum value you can configure for this property is 20.

upvoted 1 times

menlapwa 9 months ago

Selected Answer: B

The answer itself already mentioned 20, don't know why it says 40

upvoted 1 times

guschess 9 months ago

Selected Answer: B

3 fault domain
20 update domain

upvoted 2 times

Question #8 *Topic 1*

DRAG DROP -

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

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

You decide to deploy a function app.

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

Select and Place:

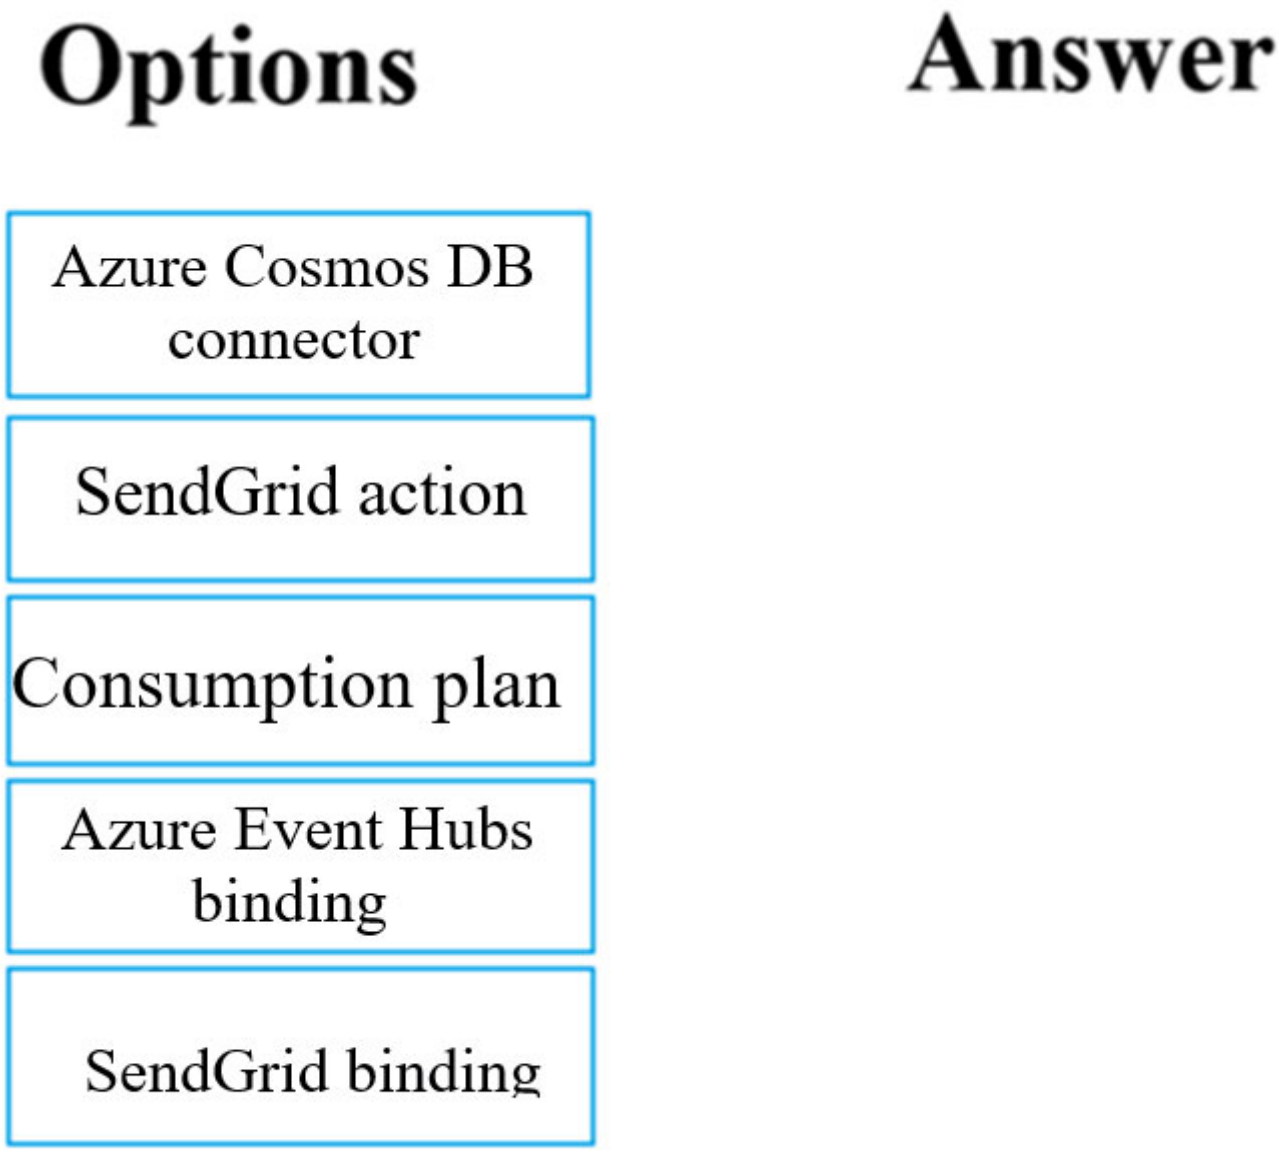

Correct Answer:

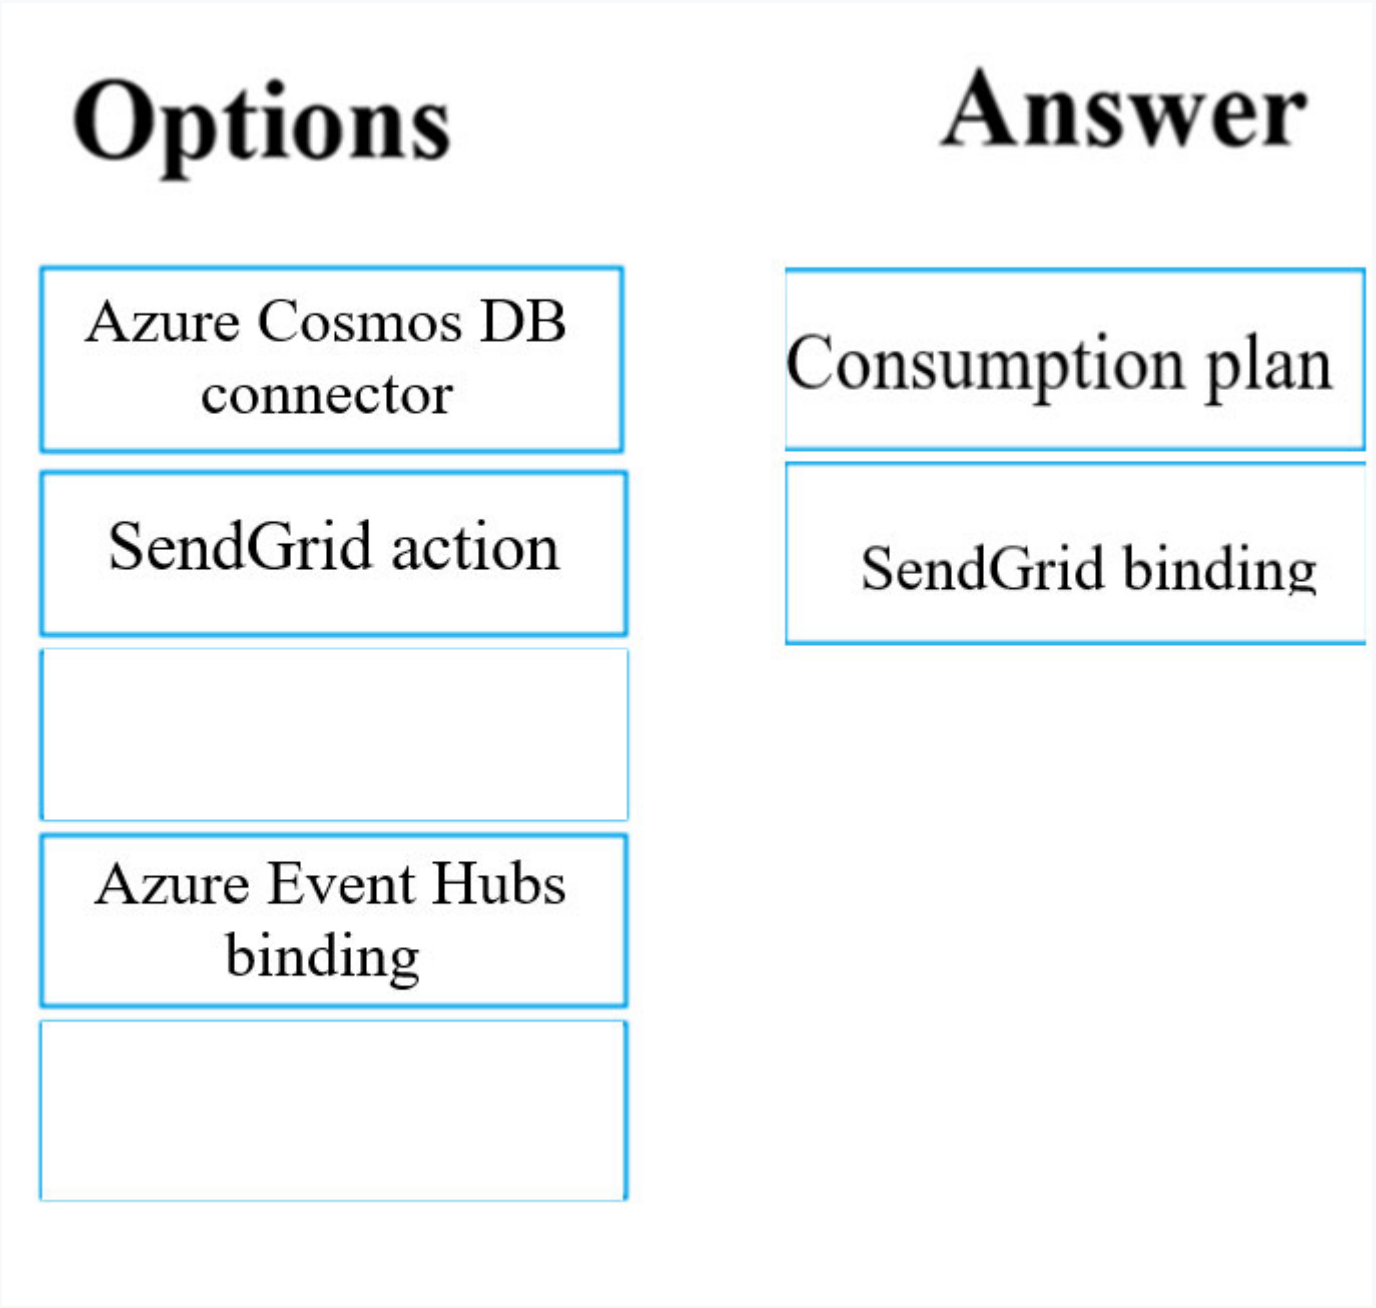

BhavikaSNN Highly Voted 1 year, 6 months ago

Answer is correct: Consumption plan will reduce the cost and SendGrid is used to send emails from azure functions.
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-sendgrid?tabs=in-process%2Cfunctionsv2&pivots=programming-language-csharp

upvoted 15 times

NhatTan 1 year, 1 month ago

Can anyone tell me why we cannot use Azure Event Hubs instead of SendGrid bindings?

upvoted 2 times

tom112 5 months, 3 weeks ago

event hub might be able to be used because it's an event, and from IoT - But generally, Event hub is for a massive amount of processing from many IoT devices. Sounds expensive, right?

upvoted 7 times

aruni_mishra 1 year ago

to send email use SendGrid bindings in Azure Function.
ps: we talk about email and not events.

upvoted 4 times

selimkayali 10 months, 1 week ago

I am not fully agree, reply from IoT device is an event imo. And you can use Event Grid for event based application. Event Hubs has a different use case like data ingestion.

upvoted 2 times

Evo_Morales 1 year, 6 months ago

Received this question last week.

upvoted 5 times

MikeAWS Most Recent 1 month ago

IMO:
1. Consumption plan, and
2. SendGrid binding

https://www.exam-answer.com/azure-cosmos-db-sql-api-notification-solution

upvoted 1 times

Aneeqa 1 month, 1 week ago

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

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

upvoted 1 times

zreaf29 3 months, 2 weeks ago

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

upvoted 1 times

zreaf29 2 months ago

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

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

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

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

upvoted 3 times

Molnmannen1337 3 months ago

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

upvoted 3 times

HQ1227 3 months, 4 weeks ago

what is different between sendgrid binding and sendgrid action?

upvoted 2 times

LuchoGonzalez 4 months, 2 weeks ago

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

upvoted 2 times

deathRac3 5 months ago

Question was there for me on 29th May 2023

upvoted 4 times

forgetfulalligator 8 months, 1 week ago

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

upvoted 4 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 2 times

Alluru 1 year ago

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

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

upvoted 1 times

Chang401 1 year ago

answer should be consumption plan and SendGrid binding. as Azure Cosmos DB connector is for logic apps not for functions. event hub is not used for IOT data. IOT HUB is used for that.

upvoted 4 times

elequiel 1 year ago

Question on day 20/10/2022

upvoted 4 times

gds03 1 year ago

isn't Event hub, since info is been sent by IOT devices? I guess should be Event Hub binding

upvoted 1 times

Chang401 1 year ago

Microsoft recommends using Azure IoT Hub to connect IoT devices to Azure

upvoted 2 times

cronatale 1 year, 1 month ago

Why don't use Azure Cosmo DB connector to trigger the Function App execution?

upvoted 3 times

Expecho 1 year ago

That is a logic apps thing.

upvoted 4 times

nvtienanh 1 year, 1 month ago

- Consumption Plan: Compute cost is reduced
- SendGrid binding (out): Send Email

But how to trigger function app when new information is received from IoT devices?
It's similar with https://www.examtopics.com/discussions/microsoft/view/10101-exam-az-301-topic-3-question-6-discussion/
The answer is not clear for me.

upvoted 3 times

garbas 4 weeks ago

https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-functions
"Currently, the Azure Functions trigger for Azure Cosmos DB is supported for use with the API for NoSQL only."

Maybe the question is outdated, but if the API is SQL, Azure Event Hubs is the only way to trigger the Azure Function when the IoT event is received (from the given options).

So the answer should include Consumption plan (reduce costs), Azure Event Hubs binding (event trigger) and SendGrid binding (mail sending).

upvoted 1 times

finnishr 1 year, 1 month ago

Correct

upvoted 2 times

Question #9 *Topic 1*

This question requires that you evaluate the underlined text to determine if it is correct.
You company has an on-premises deployment of MongoDB, and an Azure Cosmos DB account that makes use of the MongoDB API.
You need to devise a strategy to migrate MongoDB to the Azure Cosmos DB account.
You include the Data Management Gateway tool in your migration strategy.
Instructions: Review the underlined text. If it makes the statement correct, select `No change required.` If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change required

B. mongorestore

C. Azure Storage Explorer

D. AzCopy

Correct Answer: *B*
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/mongodb-migrate https://docs.mongodb.com/manual/reference/program/mongorestore/

Community vote distribution
B (71%) A (29%)

ndh103 Highly Voted 2 years, 4 months ago
I cannot see the 'underlined text' in the question.
upvoted 27 times

pedz99 2 years, 3 months ago
my gut tells me that it's Data Management Gateway since it's the only one that looks out of place
upvoted 10 times

kozchris 1 year, 8 months ago
Wrong, the answer is B - mongorestore. DMG can not have Cosmos DB - Mongo as a sync (target). Refer to the link you provided.
upvoted 1 times

jay158 Highly Voted 2 years, 3 months ago
Answer: B --mongorestore
https://docs.microsoft.com/en-us/azure/cosmos-db/mongodb-pre-migration
upvoted 26 times

ZodiaC 2 years, 3 months ago
1000000000000000000% TRUE!
upvoted 2 times

jay158 2 years, 3 months ago
Azure Database Migration Service [is for online Migration] but it is not a choice
https://docs.microsoft.com/en-us/azure/dms/tutorial-mongodb-cosmos-db-online
upvoted 4 times

AzureLearning 1 year, 9 months ago
The Data management gateway is a client agent that you must install in your on-premises environment to copy data between cloud and on-premises data stores.

https://docs.microsoft.com/en-us/azure/data-factory/v1/data-factory-data-management-gateway
upvoted 2 times

Mol_ Most Recent 1 month, 1 week ago
Selected Answer: B
Got this on the Microsoft documentation: "The MongoDB native tools are a set of binaries that facilitate data manipulation on an existing MongoDB instance. Since Azure Cosmos DB exposes an API for MongoDB, the MongoDB native tools are able to insert data into Azure Cosmos DB. The focus of this doc is on migrating data out of a MongoDB instance using mongoexport/mongoimport or mongodump/mongorestore. Since the native tools connect to MongoDB using connection strings, you can run the tools anywhere, however we recommend running these tools within the same network as the MongoDB instance to avoid firewall issues."
upvoted 1 times

Dianahu 3 months ago
No change required => https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-on-premises-mongodb-connector and https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-move-data-between-onprem-and-cloud all the links supporting option with

mongorestore do not mention data management gateway. For data management gateway you only need to specify schema and the rest is done by DMG and its pipeline with a copy activity.

upvoted 1 times

viking2 4 months ago

Nothing is underlined

upvoted 2 times

Alluru 7 months, 1 week ago

Data Management Gateway is not supported for Azure Cosmos DB - Mongo API
https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-data-movement-activities#supported-data-stores-and-formats
So you can use native tools to perform an offline (one-time) migration of databases from an on-premises or cloud instance of MongoDB to Azure Cosmos DB's API for MongoDB.
https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/tutorial-mongotools-cosmos-db#overview-of-data-migration-using-mongodb-native-tools

upvoted 1 times

JhonR 6 months, 2 weeks ago

The question doesn`t mention DataFactory, which is where your first link points.
Here we can see DMG is listed among the migration tools: https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/pre-migration-steps.

So the answer is A.

upvoted 1 times

adilkhan 7 months, 2 weeks ago

mongorestore chatGPT

upvoted 2 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: B

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

upvoted 4 times

wrak 8 months, 2 weeks ago

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

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

A - No change required

upvoted 1 times

DonH 4 months, 2 weeks ago

without any explanation this answer is useless.

upvoted 1 times

fkaracan 8 months, 2 weeks ago

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

upvoted 3 times

warchoon 8 months, 2 weeks ago

I think "MongoDB API" is a key

upvoted 1 times

guschess 9 months ago

Selected Answer: A

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

upvoted 1 times

armpro 9 months, 3 weeks ago

Selected Answer: A

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

upvoted 2 times

warchoon 8 months, 2 weeks ago

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

upvoted 1 times

lormar72 10 months, 3 weeks ago

i think is A the DMG do the job

upvoted 1 times

kkop 12 months ago

Selected Answer: B

mongorestore imports your BSON file dump into Azure Cosmos DB.
Here, you can enhance the level of your knowledge about MongoDB migrations:

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

upvoted 4 times

Alluru 1 year ago

For datastores with <10 GB monogoDB native tool is used for migration from on-premises data store to cloud data store.
https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/pre-migration-steps#pre-migration-logistics-planning

upvoted 1 times

Alluru 1 year ago

I agree with Kozchris, the answer is B - mongorestore. DMG can not have Cosmos DB - Mongo as a supported sink (target).
https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-data-movement-activities#supported-data-stores-and-formats

upvoted 1 times

Question #10 *Topic 1*

You are developing an e-Commerce Web App.
You want to use Azure Key Vault to ensure that sign-ins to the e-Commerce Web App are secured by using Azure App Service authentication and Azure Active
Directory (AAD).
What should you do on the e-Commerce Web App?

A. Run the az keyvault secret command.

B. Enable Azure AD Connect.

C. Enable Managed Service Identity (MSI).

D. Create an Azure AD service principal.

Correct Answer: *C*
A managed identity from Azure Active Directory allows your app to easily access other AAD-protected resources such as Azure Key Vault.
Reference:
https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity https://docs.microsoft.com/en-us/samples/azure-samples/app-service-msi-keyvault-dotnet/keyvault-msi-appservice-sample/

Community vote distribution
C (92%) 8%

ZodiaC Highly Voted 2 years, 3 months ago
100% Correct
upvoted 12 times

ucsdmiami2020 1 year, 10 months ago
Agreed C. Quoting the provided Microsoft docs URL references,
" Managed identities for Azure resources allow for giving Azure services an automatically managed identitfy in Azure Active Directory (Azure AD).
upvoted 4 times

Molte 1 year, 10 months ago
Why C and not D?
upvoted 2 times

chingdm 1 year, 8 months ago
"Azure Key Vault provides a way to store credentials and other secrets with increased security. But your code needs to authenticate to Key Vault to retrieve them. Managed identities for Azure resources help to solve this problem by giving Azure services an automatically managed identity in Azure Active Directory (Azure AD). You can use this identity to authenticate to any service that supports Azure AD authentication, including Key Vault, without having to display credentials in your code."
https://docs.microsoft.com/en-us/azure/key-vault/general/tutorial-net-create-vault-azure-web-app
upvoted 7 times

junkz Most Recent 1 month, 4 weeks ago
since the statement sais "secured by using Azure App Service authentication and Azure Active Directory (AAD)."
it stand to reason the nswer would be D. when configuring app service authentication with Microsoft identity provider, an app registration is required. which, under the hood, creates an app service principal (the ObjectID from the app registration, not to be confused with the applicationID of the app registration). so although the wording of the D option is missfortunate, it's still true.
upvoted 1 times

SirMarv 3 months, 3 weeks ago
Selected Answer: C
C is correct
upvoted 1 times

forgetfulalligator 8 months, 1 week ago
Selected Answer: C
Managed Service Identity (MSI) is a feature in Azure that allows you to securely authenticate an Azure service to other Azure services without having to manage credentials. By enabling MSI on the Azure App Service hosting the e-Commerce Web App, you can create a trust relationship between the App Service and Azure Key Vault. This allows the e-Commerce Web App to authenticate with Azure Active Directory (AAD) and securely retrieve secrets from the Key Vault.
upvoted 4 times

uffuchsi 8 months, 2 weeks ago
C - Enable Managed Service Identity

upvoted 1 times

examinator007 9 months ago

Please update the question: "Managed identities for Azure resources is the new name for the service formerly known as Managed Service Identity (MSI)." https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview

upvoted 4 times

alexein74 9 months, 1 week ago

D. Create an Azure AD service principal.

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

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

upvoted 3 times

Mattt 7 months ago

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

upvoted 1 times

OPT_001122 11 months ago

Selected Answer: C

C. Enable Managed Service Identity (MSI).

upvoted 2 times

KingChuang 11 months ago

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

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

upvoted 1 times

ericci 11 months, 3 weeks ago

Selected Answer: A

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

upvoted 1 times

sca88 1 year ago

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

upvoted 2 times

finnishr 1 year, 1 month ago

Correct!

upvoted 1 times

PhilLI 1 year, 9 months ago

Selected Answer: C

"Azure AD service principals" are created for Apps registered in Azure AD.
Whereas WebApps, VMs, hubs etc all can get a Managed Identity.

upvoted 4 times

xahah22222 2 years ago

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

upvoted 1 times

Sukon_Desknot 2 years, 2 months ago

This is a bit tricky checkout

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

upvoted 4 times

Question #11 *Topic 1*

This question requires that you evaluate the underlined text to determine if it is correct.
Your Azure Active Directory Azure (Azure AD) tenant has an Azure subscription linked to it.
Your developer has created a mobile application that obtains Azure AD access tokens using the OAuth 2 implicit grant type.
The mobile application must be registered in Azure AD.
You require a redirect URI from the developer for registration purposes.
Instructions: Review the underlined text. If it makes the statement correct, select `No change is needed.` If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change required.

B. a secret

C. a login hint

D. a client ID

Correct Answer: *A*

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

Reference:

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

Community vote distribution

A (100%)

jvyas Highly Voted 2 years ago

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

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

upvoted 28 times

qwerty112233 Highly Voted 1 year, 1 month ago

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

upvoted 7 times

JustHereToLearn 9 months, 1 week ago

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

upvoted 3 times

nihilek_09 Most Recent 4 weeks ago

TEST TEST

upvoted 1 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: A

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

upvoted 2 times

wrak 8 months, 2 weeks ago

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

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 1 times

guschess 9 months ago

Selected Answer: A

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

upvoted 3 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: A

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

"No change is required"

upvoted 2 times

finnishr 1 year, 1 month ago

Correct

upvoted 1 times

upadhyayavi 1 year, 6 months ago

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

upvoted 2 times

Prashant_Salokhe 1 year, 7 months ago

This type of questions are misleading. No underline text.

upvoted 3 times

SivajiTheBoss 1 year, 7 months ago

100% Answer is Correct:
Redirect URI from the developer for registration purposes will solve the purpose as the
mobile application registered in Azure AD and obtains Azure AD access tokens using the OAuth 2 implicit grant type.

upvoted 1 times

Mev4953 1 year, 9 months ago

Which text is underlined text?

upvoted 4 times

huislaw 1 year, 9 months ago

I don't see it too

upvoted 2 times

Laueri 1 year, 9 months ago

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

upvoted 6 times

Patchfox 1 year, 10 months ago

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

upvoted 4 times

Patchfox 1 year, 10 months ago

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

upvoted 3 times

MiraA 2 years ago

Register an application with the Microsoft identity platform - add a redirect URI:
https://docs.microsoft.com/en-us/azure/active-directory/develop/quickstart-register-app#add-a-redirect-uri

upvoted 2 times

MiraA 2 years ago

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

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

upvoted 1 times

Omar102 2 years, 1 month ago

Client ID is also a required.

upvoted 5 times

DV007 1 year, 9 months ago

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

upvoted 8 times

ucsdmiami2020 1 year, 10 months ago

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

"Redirect URI (optional) - We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is

optional and it can be changed later, but a value is required for most authentication scenarios."

On the other hand,

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

upvoted 1 times

noip 2 years, 2 months ago

i did not get this type of questions :/

upvoted 4 times

weasel97 2 years, 2 months ago

If the solution meets the requirements, you choose "No change is needed". If another thing (B,C or D) must be done, choose that option.

upvoted 1 times

Question #12 *Topic 1*

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

A. EnabledForDeployment

B. EnablePurgeProtection

C. EnabledForTemplateDeployment

D. EnableSoftDelete

Correct Answer: *BD*
Reference:
https://docs.microsoft.com/en-us/powershell/module/azurerm.keyvault/new-azurermkeyvault https://docs.microsoft.com/en-us/azure/key-vault/key-vault-ovw-soft-delete

Community vote distribution
BD (100%)

forgetfulalligator Highly Voted 8 months, 1 week ago

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

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

upvoted 8 times

viking2 Most Recent 4 months ago

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

upvoted 4 times

UkadikeToch 4 months, 2 weeks ago

BD

EnabledForDeployment is a property of Azure Key Vault that allows you to control whether the Key Vault can be deployed as part of an Azure Resource Manager (ARM) template deployment.
When EnabledForDeployment is set to true, it means that the Key Vault can be included in ARM templates and deployed using Azure Resource Manager.

EnablePurgeProtection is a property of Azure Key Vault's soft delete feature.
When EnablePurgeProtection is set to true, it ensures that deleted Key Vault entities (such as secrets, keys, or certificates) cannot be permanently deleted immediately.

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

EnableSoftDelete is a property of Azure Key Vault that enables the soft delete feature for the Key Vault.
Soft delete allows you to recover deleted entities (secrets, keys, or certificates) within a specified retention period, even after they have been deleted.

upvoted 3 times

AZAdam22 5 months, 3 weeks ago

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

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

Mis-type, meant to be B, and D

upvoted 1 times

alexein74 9 months, 1 week ago

B. EnablePurgeProtection
D. EnableSoftDelete

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

EnablePurgeProtection prevents the key vault from being permanently deleted before the soft-delete retention period has elapsed. EnableSoftDelete allows deleted vault and its contents to be retained and recoverable for the specified number of days.

upvoted 3 times

HafizSalmanMalik 10 months, 1 week ago

Correct options are
B. EnablePurgeProtection
D. EnableSoftDelete

upvoted 3 times

Topic9 10 months, 1 week ago

B & D.

upvoted 1 times

elequiel 11 months, 2 weeks ago

Selected Answer: BD

Correct.

upvoted 2 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: BD

correct - B and D

upvoted 1 times

Magis 12 months ago

Selected Answer: BD

Correct.

upvoted 2 times

helnik 1 year ago

correct b, d

upvoted 1 times

finnishr 1 year, 1 month ago

The answer is correct!!!

upvoted 4 times

Question #13 *Topic 1*

HOTSPOT -

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

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

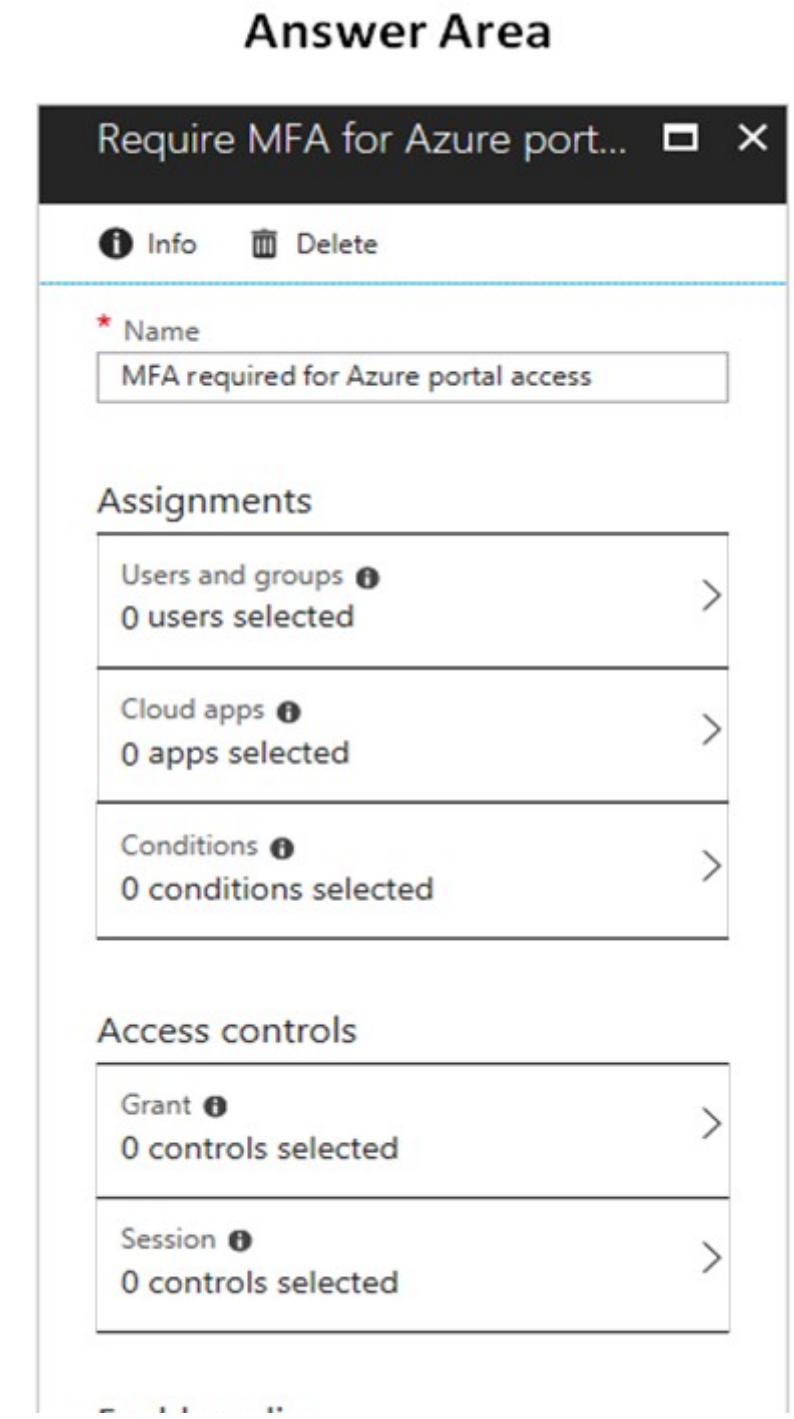

Correct Answer:

Answer Area

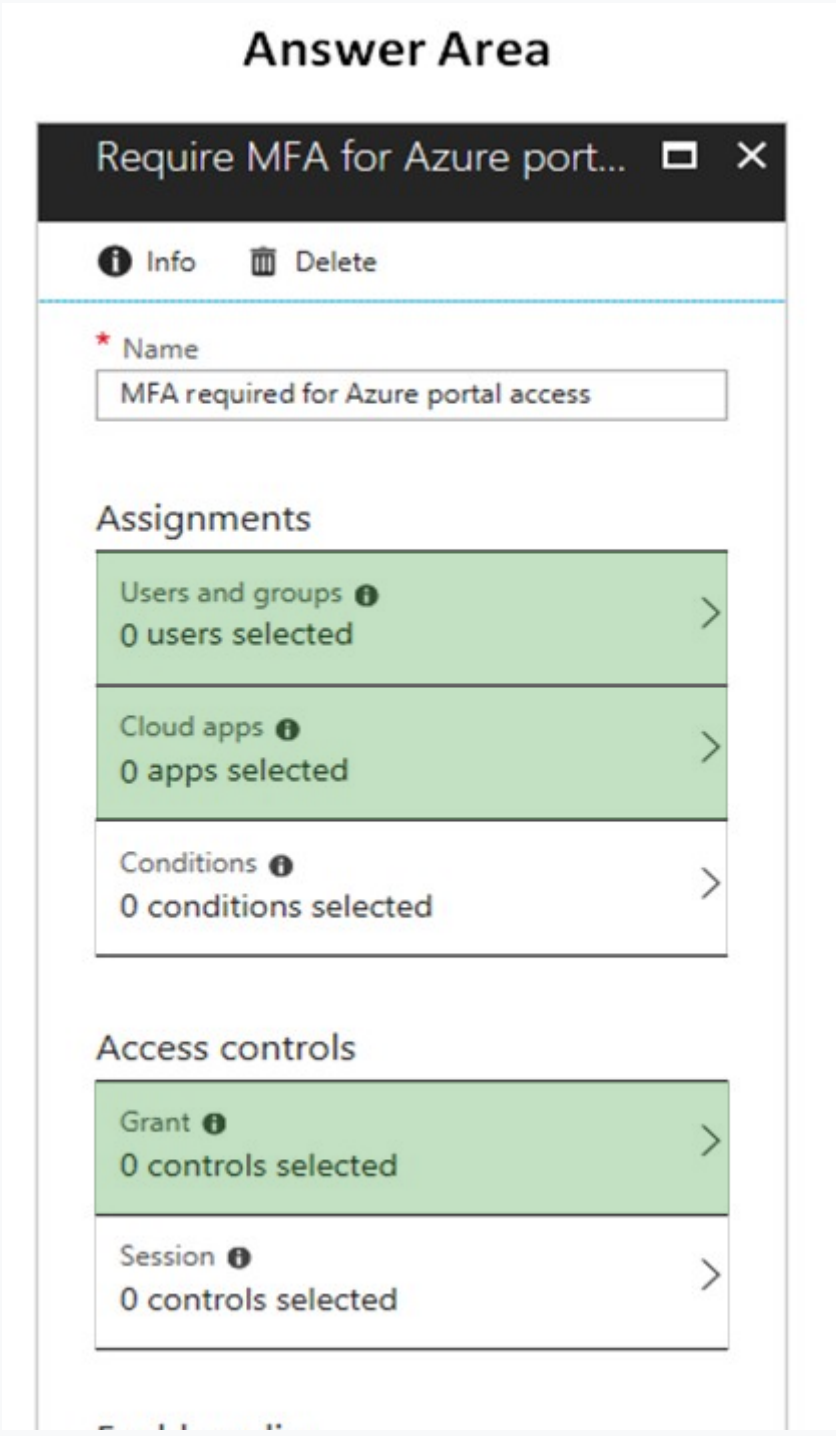

Box 1:

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

Box 2:

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

Box 3:

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

Reference:

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

[Removed] Highly Voted 1 year, 1 month ago

Microsoft wants a developer to remember their unfortunate choice to put the requirement (condition) for MFA not in the "conditions" tab, but in the "grants"?
And the people go along with these without complaining and just trying to remember this totally useless information?

upvoted 50 times

koyakei 3 months, 1 week ago

I think same, MFA must be at condition tab.

upvoted 1 times

finnishr Highly Voted 1 year, 1 month ago

The answer is correct! Please refer to:

https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/howto-conditional-access-policy-azure-management

upvoted 11 times

Nestiei Most Recent 3 weeks, 6 days ago

The answer is correct.
It's explicit here:

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

upvoted 1 times

kvtborad 2 months, 3 weeks ago

I have this question on 6th August. it was passed with 904. chose the given answer.

upvoted 1 times

656823 2 months, 3 weeks ago

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

upvoted 1 times

sbsd 8 months, 1 week ago

The answer seems correct: https://youtu.be/FxW8vLxAjSk?t=463

upvoted 4 times

forgetfulalligator 8 months, 1 week ago

Assignments: Include all users or a specific group that you want to apply the policy to.
Cloud apps or actions: Select "Microsoft Azure Management" or "Microsoft Azure portal".
Access controls: Configure the policy to require multi-factor authentication.

upvoted 2 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 1 times

OPT_001122 11 months ago

The answer is correct

upvoted 1 times

Question #14 *Topic 1*

You manage an Azure SQL database that allows for Azure AD authentication.
You need to make sure that database developers can connect to the SQL database via Microsoft SQL Server Management Studio (SSMS). You also need to make sure the developers use their on-premises Active Directory account for authentication. Your strategy should allow for authentication prompts to be kept to a minimum.
Which of the following should you implement?

A. Azure AD token.

B. Azure Multi-Factor authentication.

C. Active Directory integrated authentication.

D. OATH software tokens.

Correct Answer: *C*
Azure AD can be the initial Azure AD managed domain. Azure AD can also be an on-premises Active Directory Domain Services that is federated with the Azure
AD.
Using an Azure AD identity to connect using SSMS or SSDT
The following procedures show you how to connect to a SQL database with an Azure AD identity using SQL Server Management Studio or SQL Server Database
Tools.
Active Directory integrated authentication
Use this method if you are logged in to Windows using your Azure Active Directory credentials from a federated domain.
1. Start Management Studio or Data Tools and in the Connect to Server (or Connect to Database Engine) dialog box, in the Authentication box, select Active
Directory - Integrated. No password is needed or can be entered because your existing credentials will be presented for the connection.

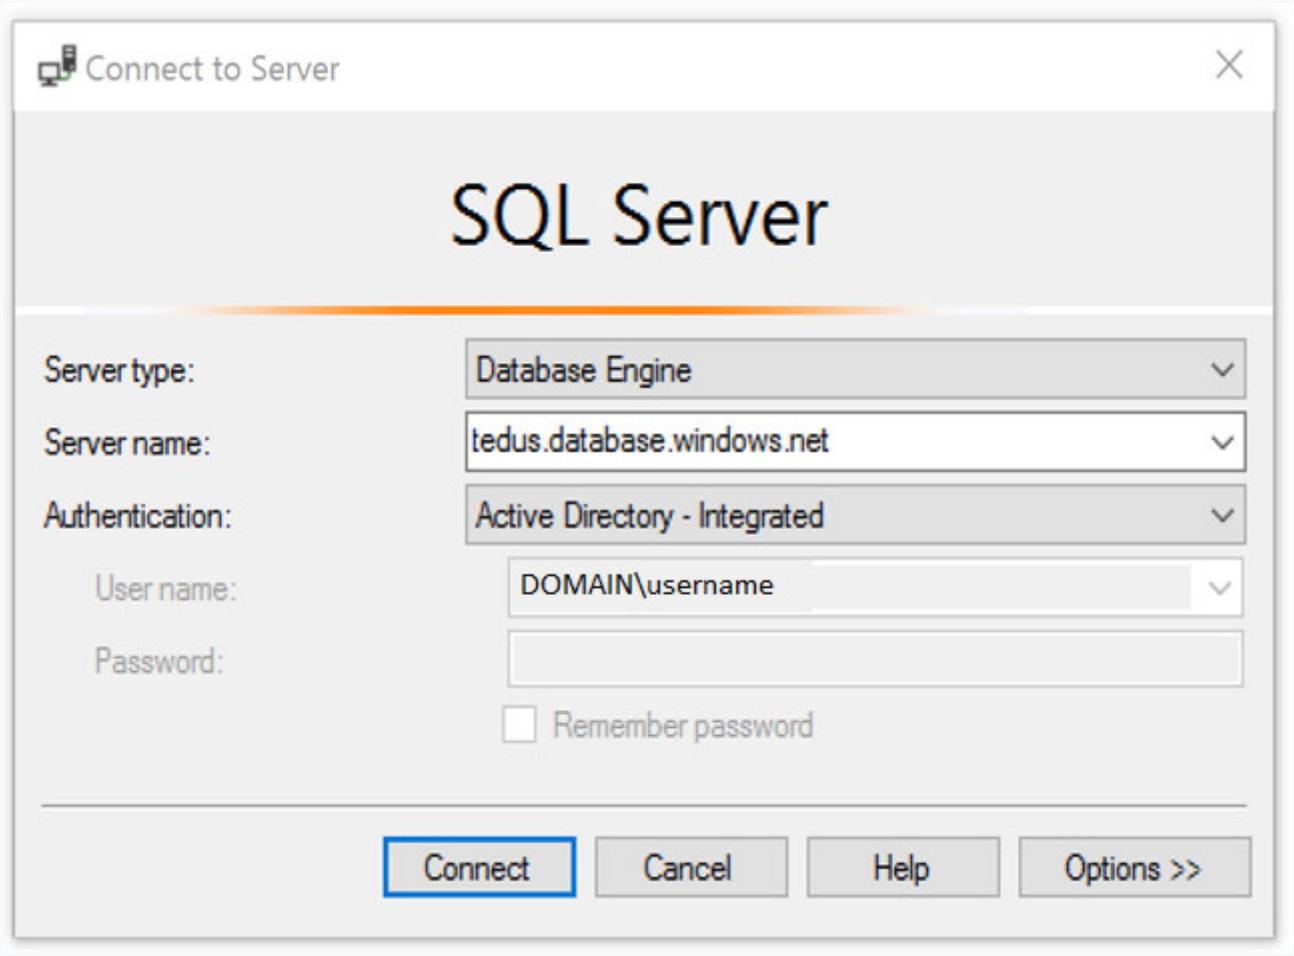

2. Select the Options button, and on the Connection Properties page, in the Connect to database box, type the name of the user database you want to connect to.
(The AD domain name or tenant IDزâ€ option is only supported for Universal with MFA connection options, otherwise it is greyed out.)

Community vote distribution
C (100%)

MrXBasit Highly Voted 2 years, 3 months ago
Correct
upvoted 26 times

uffuchsi Highly Voted 8 months, 1 week ago

Selected Answer: C

Received this in my exam day (22/02/2023). Selected C. Score 927.

upvoted 8 times

TusharChandak Most Recent 3 months, 1 week ago

Selected Answer: C

Correct

upvoted 1 times

UkadikeToch 4 months, 2 weeks ago

Active Directory integrated authentication allows developers to use their on-premises AD account to authenticate and connect to the Azure SQL database.
With this method, SSMS uses the logged-in Windows credentials of the user to authenticate against the database.
By configuring SSO between on-premises AD and Azure AD, developers can achieve a seamless authentication experience without additional prompts.

upvoted 3 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: C

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

upvoted 2 times

uffuchsi 8 months, 2 weeks ago

C - Active Directory integrated authentication

upvoted 1 times

OPT_001122 11 months ago

Selected Answer: C

C. Active Directory integrated authentication.

upvoted 2 times

sca88 1 year ago

I think is not part of the 204 exam...

upvoted 4 times

Net_IT 8 months, 2 weeks ago

I stopped wondering whether a certain question should be on a specific exam, because MS doesn't care about that. They just assume you read and remember all availble documentation on their website regarding any subject they (even vagely) mention in the lessons that are supposed to be sufficient to pass (which they are iobviously not).

upvoted 2 times

[Removed] 1 year, 1 month ago

Are we sure this is a question that comes in a developer exam?

upvoted 3 times

Question #15 *Topic 1*

You are developing an application to transfer data between on-premises file servers and Azure Blob storage. The application stores keys, secrets, and certificates in Azure Key Vault and makes use of the Azure Key Vault APIs.
You want to configure the application to allow recovery of an accidental deletion of the key vault or key vault objects for 90 days after deletion.
What should you do?

A. Run the Add-AzKeyVaultKey cmdlet.

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

C. Implement virtual network service endpoints for Azure Key Vault.

D. Run the az keyvault update --enable-soft-delete false CLI.

Correct Answer: *B*
When soft-delete is enabled, resources marked as deleted resources are retained for a specified period (90 days by default). The service further provides a mechanism for recovering the deleted object, essentially undoing the deletion.
Purge protection is an optional Key Vault behavior and is not enabled by default. Purge protection can only be enabled once soft-delete is enabled.
When purge protection is on, a vault or an object in the deleted state cannot be purged until the retention period has passed. Soft-deleted vaults and objects can still be recovered, ensuring that the retention policy will be followed.
The default retention period is 90 days, but it is possible to set the retention policy interval to a value from 7 to 90 days through the Azure portal. Once the retention policy interval is set and saved it cannot be changed for that vault.
Reference:
https://docs.microsoft.com/en-us/azure/key-vault/general/overview-soft-delete

Community vote distribution
B (100%)

anu_ezio Highly Voted 1 year ago

Selected Answer: B
Correct Answer
upvoted 8 times

viking2 Most Recent 4 months ago

Selected Answer: B
B is correct. Soft-delete and purge protection are both needed, and B shows the correct syntax to achieve this
upvoted 1 times

forgetfulalligator 8 months, 1 week ago

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

The --enable-purge-protection flag turns on purge protection, which prevents the permanent deletion of objects until the soft-delete retention period has expired. This ensures that deleted objects can be recovered within the retention period.
upvoted 3 times

uffuchsi 8 months, 2 weeks ago

B. Purge defines how long a key vault entry is 'alive' before it can be deleted. Soft delete allows you to recover mistakenly deleted key vault entries.
upvoted 1 times

alexein74 9 months, 1 week ago

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

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

Soft delete allows deleted key vaults and their contents to be retained and recoverable for the specified number of days after deletion.
Purge protection prevents key vault from being permanently deleted before the soft-delete retention period has elapsed.
upvoted 2 times

HafizSalmanMalik 10 months, 1 week ago

Correct option is
Answer: B

upvoted 1 times

Topic9 10 months, 1 week ago

Correct answer

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: B

correct

upvoted 2 times

ahmad123456789 11 months, 3 weeks ago

correct

upvoted 2 times

Ganeshd7 1 year ago

correct

upvoted 1 times

RachidBinAhmed 1 year, 1 month ago

correct

upvoted 1 times

finnishr 1 year, 1 month ago

Correct

upvoted 3 times

kampatra 1 year, 1 month ago

Correct

upvoted 1 times

Question #16 *Topic 1*

HOTSPOT -

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

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

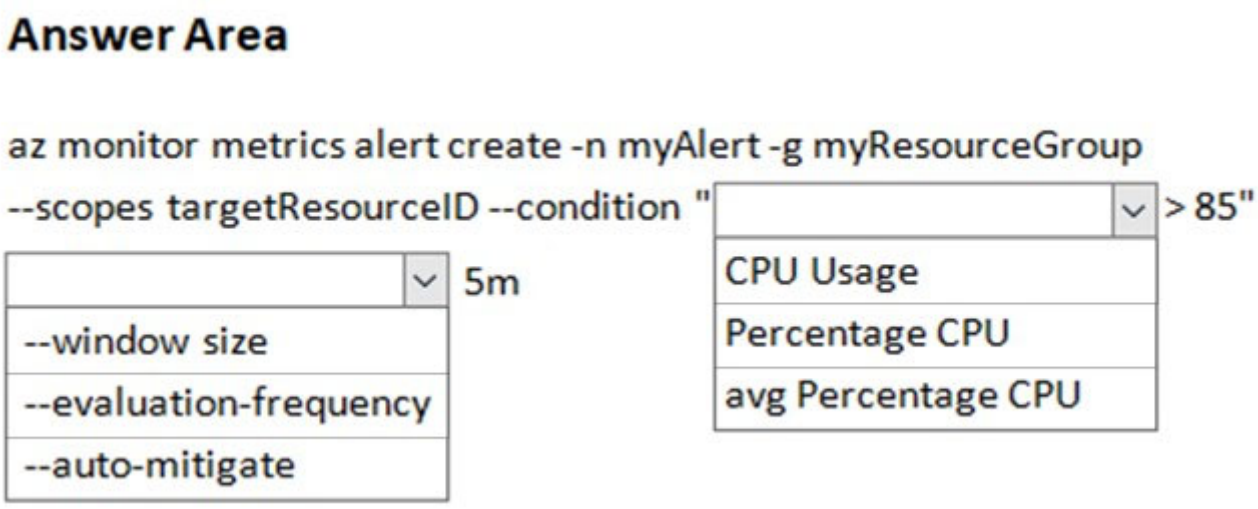

Correct Answer:

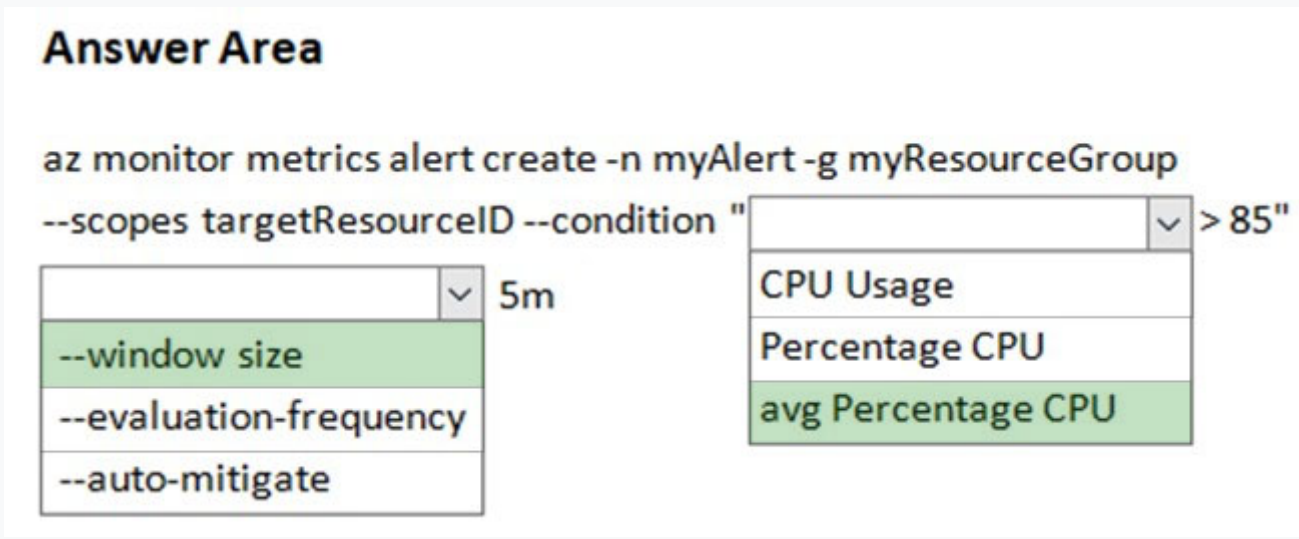

Reference:

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

2bruhornot2bruh Highly Voted 11 months, 2 weeks ago

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

upvoted 29 times

OPT_001122 10 months, 4 weeks ago

According to 2bruhornot2bruh explanation, Windows size is the correct answer.

upvoted 3 times

CellCS Highly Voted 1 year, 6 months ago

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

upvoted 18 times

yunqing Most Recent 4 weeks, 1 day ago

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

upvoted 1 times

zreaf29 2 months ago

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

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

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

upvoted 4 times

Camarade_Emile 3 months, 1 week ago

This was part of my exam

upvoted 2 times

uffuchsi 8 months, 2 weeks ago

1. avg Percentage CPU, 2. window size

upvoted 1 times

Topic9 10 months, 1 week ago

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

upvoted 2 times

narenazure 10 months, 1 week ago

Box 2: --window-size
Time over which to aggregate metrics in "##h##m##s" format.
https://learn.microsoft.com/en-us/cli/azure/monitor/metrics/alert?view=azure-cli-latest

upvoted 1 times

POOOJAAAAAAAAAA 10 months, 1 week ago

Box 2: Window Size

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

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

what is the correct value for box2?

upvoted 1 times

[Removed] 1 year, 1 month ago

Sure. I know how to use this from the portal and now I have to remember the unfortunate name someone in Microsoft chose for the evaluation time?

upvoted 3 times

DiveyK 1 year, 1 month ago

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

upvoted 3 times

RachidBinAhmed 1 year, 1 month ago

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

upvoted 1 times

srinz 1 year, 1 month ago

Given answer is correct

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

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

upvoted 2 times

finnishr 1 year, 1 month ago

Correct

upvoted 2 times

kampatra 1 year, 1 month ago

Box 1: avg Percentage CPU > 85
Box 2: --evaluation-frequency 5m

upvoted 6 times

OPT_001122 11 months, 3 weeks ago

why evaluation-frequency 5m, still not clear

upvoted 3 times

JohnSmith93 1 year ago

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

upvoted 1 times

vadimprokopchuk 1 year, 1 month ago

Box 1: avg Percentage CPU > 85
Box 2: --evaluation-frequency 5m

upvoted 3 times

Question #17 *Topic 1*

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.
Establish if the solution satisfies the requirements.
You are configuring a web app that delivers streaming video to users. The application makes use of continuous integration and deployment.
You need to ensure that the application is highly available and that the users' streaming experience is constant. You also want to configure the application to store data in a geographic location that is nearest to the user.
Solution: You include the use of Azure Redis Cache in your design.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Community vote distribution

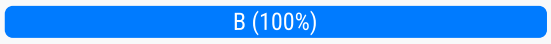

NStanhope Highly Voted 2 years, 3 months ago

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

upvoted 35 times

ucsdmiami2020 1 year, 10 months ago

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

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

upvoted 11 times

StephanieMii 2 years, 3 months ago

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

upvoted 3 times

ranjitklive 2 years, 2 months ago

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

upvoted 11 times

Veks 1 year, 7 months ago

Hmm, Azure Cache for Redis 6.0 includes Redis streams now...
I agree that CDN is correct because u have to ask your self who is writing these questions and why.
I'm sure that Redis is the solution that meets the goal, but the problem is, is it a solution form Microsoft exam question writer :)

upvoted 4 times

Santhosh_Kumar Most Recent 2 months, 1 week ago

Selected Answer: B

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

upvoted 1 times

Fabiann 2 months, 2 weeks ago

This is an open question :
I am agree about the use of Azure CDN about the videos management.

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

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

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

upvoted 1 times

longnguyendh 3 months ago

The answer should be Azure Content Delivery Network (CDN)

upvoted 1 times

sayanchq 4 months, 1 week ago

Should be CDN .. Hence No

upvoted 1 times

UkadikeToch 4 months, 2 weeks ago

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

upvoted 2 times

uffuchsi 8 months, 1 week ago

Received this in my exam today (22/02/2023). Selected B. Score 927.

upvoted 3 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: B

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

upvoted 2 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 1 times

guschess 9 months ago

Selected Answer: B

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

upvoted 1 times

alexein74 9 months, 1 week ago

B. No

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

upvoted 2 times

garbas 2 weeks ago

Redis Cache has geographic replication:
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-active-geo-replication

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

upvoted 1 times

Sukon_Desknot 9 months, 1 week ago

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

upvoted 1 times

Topic9 10 months, 1 week ago

CDN is right fit here

upvoted 1 times

Nedum1983 11 months, 2 weeks ago

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

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

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: B

use CDN

upvoted 1 times

philomath 1 year ago

CDN I think.

upvoted 1 times

Chief2110 1 year, 1 month ago

Selected Answer: B

Correct, CDN!

upvoted 1 times

Question #18 *Topic 1*

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.
Establish if the solution satisfies the requirements.
You are configuring a web app that delivers streaming video to users. The application makes use of continuous integration and deployment.
You need to ensure that the application is highly available and that the users' streaming experience is constant. You also want to configure the application to store data in a geographic location that is nearest to the user.
Solution: You include the use of an Azure Content Delivery Network (CDN) in your design.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*
Reference:
https://docs.microsoft.com/en-in/azure/cdn/

Community vote distribution
A (100%)

ZodiaC Highly Voted 2 years, 3 months ago
1000000000000000000% CORRECT!
upvoted 16 times

d0bermannn 2 years, 2 months ago
you can safely delete 17 zeroes)
upvoted 26 times

BishopeL 2 years, 2 months ago
Just type "CORRECT"
upvoted 12 times

uffuchsi Most Recent 8 months, 1 week ago
Received this in my exam day (22/02/2023). Selected A. Score 927.
upvoted 2 times

uffuchsi 8 months, 2 weeks ago
A. Yes - Since this will need to Cache video content, use CDN
upvoted 1 times

guschess 9 months ago
Selected Answer: A
Check the former question
upvoted 1 times

alexein74 9 months, 1 week ago
Selected Answer: A
For a web app that delivers streaming video to users, you would need to use a combination of technologies that provide high availability, constant streaming experience and store data in a geographic location that is nearest to the user.
Azure Content Delivery Network (CDN) can be used to deliver streaming video content to users with low latency by caching the video at edge locations that are nearest to the user. Azure Media Services can be used to encode, package, and stream video content. Azure Media Services also provides built-in redundancy and failover options to ensure high availability.
Another option is to store the video files in Azure Blob storage, which provides a high-availability and low-latency storage solution that can be replicated to multiple regions.
upvoted 4 times

murdokio 11 months, 1 week ago
Selected Answer: A
agreed
upvoted 2 times

Nedum1983 11 months, 2 weeks ago
Yes! I agreed.
upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: A

Correct

upvoted 1 times

Magis 12 months ago

Selected Answer: A

Correct.

upvoted 2 times

sfofana 1 year ago

CDN IS CORRECT!

upvoted 1 times

ranjitklive 2 years, 2 months ago

What is the difference between Azure CDN and Azure Media Services?
https://docs.microsoft.com/en-us/azure/media-services/latest/media-services-overview

upvoted 3 times

Question #19 *Topic 1*

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.
Establish if the solution satisfies the requirements.
You are configuring a web app that delivers streaming video to users. The application makes use of continuous integration and deployment.
You need to ensure that the application is highly available and that the users' streaming experience is constant. You also want to configure the application to store data in a geographic location that is nearest to the user.
Solution: You include the use of a Storage Area Network (SAN) in your design.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Community vote distribution

B (100%)

forgetfulalligator Highly Voted 8 months, 1 week ago

Selected Answer: B

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

upvoted 12 times

jvyas Highly Voted 2 years ago

Instead you use CDN

upvoted 11 times

deathRac3 Most Recent 5 months, 1 week ago

Selected Answer: B

Its correct

upvoted 1 times

uffuchsi 8 months, 1 week ago

Received this in my exam day (22/02/2023). Selected B. Score 927.

upvoted 3 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 1 times

guschess 9 months ago

Selected Answer: B

CDN instead

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. No

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

upvoted 4 times

HafizSalmanMalik 10 months, 1 week ago

Correct answer is by using CDN

upvoted 1 times

Topic9 10 months, 1 week ago

CDN here

upvoted 1 times

BrainPowder 1 year ago

Ok CDN, makes sense

upvoted 2 times

Migsmiku 1 year ago

Selected Answer: B

Should use CDN

upvoted 1 times

upadhyayavi 1 year, 6 months ago

What about ensuring the avaliability of the application??

upvoted 1 times

Veks 1 year, 7 months ago

What storage u use has nothing to do with what your app is doing. OK, you will not use cold storage, of course, but faster or little slower disks don't have a big impact on the solution that has to provide a lot of content over internet and fast. This question is not about disk (storage) latency, but about internet access to content and its latency.
Definitely NO.
https://www.backblaze.com/blog/whats-the-diff-nas-vs-san/

upvoted 4 times

Vrushalij 2 years, 2 months ago

No is correct answer

upvoted 4 times

ranjitklive 2 years, 2 months ago

Does Azure storage uses SAN behind-the-scenes? The closes option in Azure for the requirement given is Azure Media Services.
https://docs.microsoft.com/en-us/azure/media-services/latest/media-services-overview

upvoted 2 times

RikinPatel 1 month, 2 weeks ago

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

upvoted 1 times

Question #20 *Topic 1*

You develop a Web App on a tier D1 app service plan.
You notice that page load times increase during periods of peak traffic.
You want to implement automatic scaling when CPU load is above 80 percent. Your solution must minimize costs.
What should you do first?

A. Enable autoscaling on the Web App.

B. Switch to the Premium App Service tier plan.

C. Switch to the Standard App Service tier plan.

D. Switch to the Azure App Services consumption plan.

Correct Answer: *C*
Configure the web app to the Standard App Service Tier. The Standard tier supports auto-scaling, and we should minimize the cost. We can then enable autoscaling on the web app, add a scale rule and add a Scale condition.
Reference:
https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-autoscale-get-started https://azure.microsoft.com/en-us/pricing/details/app-service/plans/

Community vote distribution
C (100%)

examTaker455 Highly Voted 2 years, 3 months ago
D1 is Shared Tier and does not offer autoscaling. You need to switch to Standard, Premium or Isolated.
upvoted 34 times

ZodiaC 2 years, 3 months ago
So you have to reduce cost som you need standard app tier. 1000% correct!
upvoted 9 times

Gautam47 Highly Voted 2 years, 1 month ago
Tier D1 is a basically shared app service plan, so we need to move standard or premium plan to enable auto scaling. As we need to provide low cost solution, then standard plan will be best for this approach
upvoted 15 times

zreaf29 Most Recent 1 month, 3 weeks ago
The D1 tier plan is a basic tier plan that does not offer autoscaling. For using autoscaling, you have to change D1 basic tier plan to Others(Standard, Premium, Isolated, Consumption). The Standard plan is cheapest plan which has auto scaling. So Answer is C.
upvoted 1 times

seanishok 2 months ago
Selected Answer: C
Standard Tier is the correct answer.
upvoted 1 times

longnguyendh 3 months ago
D1 is basic tier and It does not support auto scaling feature. => we should upgrade to Standard tier
upvoted 1 times

arpannaprap 3 months, 2 weeks ago
Standard tier is the correct one.
upvoted 1 times

viking2 4 months ago
C is correct.
Auto scaling is supported on both Premium and Standard, but Standard is lower cost
upvoted 1 times

LuchoGonzalez 4 months, 2 weeks ago
Enabling autoscaling (option A) can still be a valuable addition to further optimize the resource allocation based on the actual demand. By combining autoscaling with the upgraded Standard App Service tier, you can automatically scale the resources up or down based on CPU load, ensuring optimal performance during peak traffic while minimizing costs during periods of lower traffic.

Therefore, the recommended course of action in this scenario would be to switch to the Standard App Service tier (option C) and then enable autoscaling (option A) to effectively address the page load time issue during peak traffic while considering cost optimization.
upvoted 1 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 3 times

adilkhan 7 months, 2 weeks ago

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

upvoted 1 times

tom112 5 months, 3 weeks ago

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

upvoted 4 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: C

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

upvoted 1 times

wrak 8 months, 2 weeks ago

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

upvoted 3 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

Correction: C. Switch to Standard, not B

upvoted 1 times

guschess 9 months ago

Selected Answer: C

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

upvoted 1 times

santoshsidnal 10 months, 2 weeks ago

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

upvoted 1 times

OPT_001122 11 months ago

Selected Answer: C

C. Switch to the Standard App Service tier plan.

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: C

switch to Standard

upvoted 2 times

Question #21 *Topic 1*

Your company's Azure subscription includes an Azure Log Analytics workspace.
Your company has a hundred on-premises servers that run either Windows Server 2012 R2 or Windows Server 2016, and is linked to the Azure Log Analytics workspace. The Azure Log Analytics workspace is set up to gather performance counters associated with security from these linked servers.
You must configure alerts based on the information gathered by the Azure Log Analytics workspace.
You have to make sure that alert rules allow for dimensions, and that alert creation time should be kept to a minimum. Furthermore, a single alert notification must be created when the alert is created and when the alert is resolved.
You need to make use of the necessary signal type when creating the alert rules.
Which of the following is the option you should use?

A. The Activity log signal type.

B. The Application Log signal type.

C. The Metric signal type.

D. The Audit Log signal type.

Correct Answer: *C*
Metric alerts in Azure Monitor provide a way to get notified when one of your metrics cross a threshold. Metric alerts work on a range of multi-dimensional platform metrics, custom metrics, Application Insights standard and custom metrics.
Note: Signals are emitted by the target resource and can be of several types. Metric, Activity log, Application Insights, and Log.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-metric

Community vote distribution
C (100%)

Pirgos Highly Voted 2 years, 3 months ago
C, https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-logs#configuring-metric-alert-for-logs
upvoted 17 times

MiraA 2 years ago
From that page:
"Metric Alerts for Log provide multiple dimensions, allowing filtering to specific values like Computers, OS Type, etc. simpler; without the need for defining a complex query in Log Analytics."
upvoted 7 times

MiraA 2 years ago
From that page:
"Metric Alerts are stateful - only notifying once when alert is fired and once when alert is resolved; as opposed to Log alerts, which are stateless and keep firing at every interval if the alert condition is met."
upvoted 11 times

MiraA Highly Voted 2 years ago
Signal types:
* Activity Log ... includes service health records along with records on any configuration changes made to the resources (and is available to all Azure resources)
* Audit Log ... contains the history of sign-in activity and audit trail of changes made within a particular tenant
* Metric ... numerical values that are collected at regular intervals and describe some aspect of a system at a particular time
* Application Log ... ?

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

https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-overview
https://docs.microsoft.com/en-us/azure/azure-monitor/essentials/data-platform-metrics
https://docs.microsoft.com/en-us/azure/azure-monitor/logs/data-platform-logs
upvoted 17 times

forgetfulalligator Most Recent 8 months, 1 week ago
Selected Answer: C
The Metric signal type is the correct option for creating alert rules based on performance counters associated with security gathered by the Azure Log Analytics workspace. The Metric signal type allows you to create alerts based on metrics collected by Azure Monitor. This includes metrics collected by Log Analytics from your servers, such as performance counters associated with security. The Metric signal type allows you to set dimensions, which can be used to filter the metrics and reduce the noise of alerts. You can also configure the alert to be fired on multiple criteria

and the alert creation time is kept to a minimum. Furthermore, a single alert notification is created when the alert is created and when the alert is resolved.

upvoted 2 times

wrak 8 months, 2 weeks ago

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

upvoted 2 times

uffuchsi 8 months, 2 weeks ago

B. Metric

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

Correctio: C. Metric

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: C

C. The Metric signal type

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

upvoted 2 times

nhannn 1 year, 1 month ago

Selected Answer: C

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

upvoted 6 times

Anitay 1 year, 1 month ago

C is correct answer

upvoted 1 times

Question #22 *Topic 1*

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

A. Configure the SearchMode property of the SearchParameters class.

B. Configure the QueryType property of the SearchParameters class.

C. Configure the Facets property of the SearchParameters class.

D. Configure the Filter property of the SearchParameters class.

Correct Answer: *B*
The SearchParameters.QueryType Property gets or sets a value that specifies the syntax of the search query. The default is 'simple'. Use 'full' if your query uses the Lucene query syntax.
You can write queries against Azure Search based on the rich Lucene Query Parser syntax for specialized query forms: wildcard, fuzzy search, proximity search, regular expressions are a few examples.
Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.querytype

Community vote distribution
B (100%)

7ack Highly Voted 2 years, 4 months ago
Correct, although I'm not sure Azure search is in the exam...
upvoted 21 times

gunz123 Highly Voted 2 years, 3 months ago
Answer is correct

https://docs.microsoft.com/en-us/azure/search/search-query-overview
upvoted 14 times

N9 Most Recent 4 weeks ago
Selected Answer: B
B is correct answer
upvoted 1 times

TusharChandak 3 months, 1 week ago
This is not part of AZ-204 anymore.
upvoted 6 times

LuchoGonzalez 4 months, 2 weeks ago
B. Configure the QueryType property of the SearchParameters class.

The QueryType property of the SearchParameters class allows you to specify the type of query that will be executed against the Azure Search index. By setting the QueryType to "full" or "full+regex", you can enable regular expression support in the search queries.
upvoted 2 times

forgetfulalligator 8 months, 1 week ago
Selected Answer: B
The QueryType property of the SearchParameters class specifies the type of query to use when searching the index. By default, it is set to Simple, which means that the query is interpreted as a simple query string. However, you can set it to Full, which means that the query is interpreted as a full Lucene query string, allowing for regular expressions to be used in the search query
upvoted 1 times

NombreFalso 8 months, 2 weeks ago
huh, I'll ignore these
upvoted 1 times

wrak 8 months, 2 weeks ago
out of the scope for AZ-204 exam; please remove/hide this question from the Question Bank

upvoted 5 times

uffuchsi 8 months, 2 weeks ago

D. Filter property or B. QueryType property

upvoted 1 times

examinator007 9 months ago

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

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. Configure the QueryType property of the SearchParameters class.

To allow the application to search the index by using regular expressions, you should configure the QueryType property of the SearchParameters class in your .NET Core MVC application. The QueryType property allows you to specify the type of query to use when searching the index, including regular expression queries.
You can set the QueryType property to "full" to enable full-text search or "regex" to enable regular expression search.

upvoted 1 times

halfway 9 months, 2 weeks ago

Selected Answer: B

QueryType

upvoted 1 times

OPT_001122 1 year, 1 month ago

https://www.examtopics.com/exams/microsoft/az-204/view/6/

same question #28 has answer D.
which one is correct

upvoted 1 times

Chang401 1 year ago

ask is different in both the questions.

upvoted 2 times

thedevguy 1 year ago

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

upvoted 2 times

OPT_001122 11 months, 3 weeks ago

yes Thanks.

upvoted 1 times

jay158 2 years, 3 months ago

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

upvoted 1 times

jay158 2 years, 3 months ago

my bad should be B

upvoted 7 times

Question #23 *Topic 1*

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

A. the Enterprise Integration Pack (EIP)

B. the Logic App Code View

C. the API Connections

D. the Logic Apps Designer

Correct Answer: *B*
Edit JSON - Azure portal -
1. Sign in to the Azure portal.
2. From the left menu, choose All services. In the search box, find "logic apps", and then from the results, select your logic app.
3. On your logic app's menu, under Development Tools, select Logic App Code View.
4. The Code View editor opens and shows your logic app definition in JSON format.
Reference:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-overview https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions

Community vote distribution
B (91%) 9%

abdou1987 Highly Voted 2 years, 3 months ago
the Logic App has not been part of the exam since March
upvoted 19 times

MiraA Highly Voted 2 years ago
I think being a developer I can use both options - Logic App Code View or Logic Apps Designer. It depends on the change required, usage of a source control, the project complexity, CI/CD...

Note: Some Azure Logic Apps capabilities, such as defining parameters and multiple triggers in logic app definitions, are available only in JSON, not the Logic Apps Designer. So for these tasks, you must work in Code View or another editor.
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions
upvoted 10 times

ConanBarb Most Recent 1 month ago
Selected Answer: B
That's an ignorant question. It totally depends on lots of things - not the least what type of change we are talking about. I would venture to state that it is virtually impossible, or at least takes at least 10x the time, to make a general big changes in code view such as adding new Actions, Connector Actions, a Loop or something. Thats what you do in the designer with just a few button clicks and getting it right from start.

Small and specific changes yes, such as parameters, certain variable settings, etc
upvoted 1 times

UkadikeToch 4 months, 2 weeks ago
The the logic app code view is in the logic apps designer. So i think the more accurate answer is THE LOGIC APP CODE VIEW
upvoted 2 times

moota 5 months, 1 week ago
Selected Answer: B
ChatGPT: B: B. the Logic App Code View

Logic Apps in Azure provide a Code View that allows you to directly edit the JSON definition of your Logic App. This can be useful if you need to make advanced changes that are not easily done in the designer view, or if you are more comfortable working directly with the underlying code. The Logic Apps Designer is more for creating and configuring workflows in a visual manner, but the Code View is the more direct method to update the definitions.
upvoted 3 times

NiebieskiEskimos 6 months, 2 weeks ago
Selected Answer: B
Answer is correct
upvoted 1 times

adilkhan 7 months, 2 weeks ago

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

upvoted 1 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: B

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

upvoted 2 times

wrak 8 months, 2 weeks ago

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

upvoted 5 times

uffuchsi 8 months, 2 weeks ago

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

upvoted 1 times

Chang401 8 months, 3 weeks ago

answer should be Logic App Code View.
To work with workflow definitions in JSON, open the Code View editor when working in the Azure portal or in Visual Studio, (go tot he below link and just above the first note its mentioned)
https://learn.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions

upvoted 1 times

Osip21 9 months ago

Selected Answer: D

D. the Logic Apps Designer.

upvoted 1 times

fkaracan 8 months, 2 weeks ago

why? at least explain your answer...

upvoted 1 times

Molnmannen1337 3 months ago

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

upvoted 2 times

murdokio 11 months, 1 week ago

Selected Answer: B

We need to see definitions! similar as methods

upvoted 1 times

AsuraX 11 months, 2 weeks ago

Selected Answer: B

Not all definitions can be changed in Logic Apps Designer (e. g. defining parameters) and you must pick THE BEST solution. So remember: definitions == Logic App Code View.
More info: https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions

upvoted 2 times

ning 2 years, 2 months ago

No idea, why Logic App Designer cannot work???

upvoted 3 times

coffecold 1 year ago

Code Viewer is part of Logic App Designer. So it could work. But best answer remains code viewer

upvoted 1 times

tomomo1219 2 years, 2 months ago

I think it's because "You are a developer at your company." Logic App Designer is for non-developers.

upvoted 8 times

NombreFalso 8 months, 2 weeks ago

Oof, sick burn

upvoted 1 times

BishopeL 2 years, 2 months ago

"You are a developer" hence, the Code Viewer should be suitable for you.

https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions

upvoted 4 times

Question #24 *Topic 1*

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.
Establish if the solution satisfies the requirements.
You are developing a solution for a public facing API.
The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end.
You must configure back-end authentication for the API Management service instance.
Solution: You configure Basic gateway credentials for the Azure resource.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
API Management allows to secure access to the back-end service of an API using client certificates.
Reference:
https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-backend-entity

Community vote distribution
B (65%) A (35%)

lexowe9241 Highly Voted 2 years, 1 month ago
Hmm, there is "Authenticate with Basic policy" for api-management https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#Basic
So maybe Yes is correct answer
upvoted 19 times

PhilLI 1 year, 9 months ago
Agree. My research lead me also to this page.
Authentication policies:
- Authenticate with Basic -> Authenticate with a backend service using Basic authentication.
- Authenticate with client certificate -> Authenticate with a backend service using client certificates.
- Authenticate with managed identity -> Authenticate with the managed identity for the API Management service.
upvoted 7 times

Mev4953 1 year, 9 months ago
I agree with you.
https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies
upvoted 1 times

basquiat 2 months, 3 weeks ago
Authentication policy sounds quite different from authentication gateway IMHO.
upvoted 1 times

MiraA 2 years ago
The answer is YES (Basic + the HTTP(s) endpoint) based on this screenshot:

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

Target: "Azure Logic App" or "HTTP(s) endpoint"
Gateway credentials: "None" or "Basic" or "Client cert"
upvoted 3 times

MiraA 2 years ago
My own correction... :-(

The answer is NO.
The solution mentions "You configure Basic gateway credentials for the Azure resource." and there is no such option for generic Azure resource.
upvoted 9 times

warchoon 8 months, 2 weeks ago
Another tricky question about settings UI, that will soon be deprecated :)
upvoted 1 times

koolexam 1 year, 10 months ago
But you can do that at collection level in APIM.

upvoted 1 times

leonidn Highly Voted 1 year, 9 months ago

Selected Answer: B

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

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

upvoted 13 times

Mick1z8 Most Recent 5 days, 17 hours ago

Selected Answer: B

There are two components we're interested in:
a. APIM, and
b. App Service (resource)

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

upvoted 1 times

Mick1z8 5 days, 18 hours ago

Selected Answer: A

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

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

upvoted 1 times

Vladimir_Gajinov 1 month, 3 weeks ago

Selected Answer: A

The correct answer is YES.
Use the authentication-basic policy to authenticate with a backend service using Basic authentication.

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

upvoted 3 times

JH81 4 months, 1 week ago

Selected Answer: B

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

upvoted 1 times

BikashKatwal 4 months, 3 weeks ago

B. No

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

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

upvoted 3 times

Mattt 7 months ago

B. No.

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

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

Client Certificate authentication
Token-based authentication
OAuth 2.0 authentication
These options provide secure authentication and access control for the API Management service instance and its associated APIs.

upvoted 5 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: A

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

upvoted 3 times

forgetfulalligator 8 months, 1 week ago

Opps sorry meant to select B

upvoted 1 times

uffuchsi 8 months, 2 weeks ago

B. Would you not instead use Anon and authenticate pass through JWT token to the backend?
upvoted 2 times

warchoon 8 months, 2 weeks ago
So it will work much faster than certs.
upvoted 1 times

alexein74 9 months, 1 week ago
Selected Answer: B
B. No

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

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

In summary, configuring Basic gateway credentials for the Azure resource does not meet the goal of securing back-end authentication for the public facing API.
upvoted 2 times

fkaracan 8 months, 2 weeks ago
not everytime chatgpt generates the correct answer. you need to do your own research before paste here...
upvoted 3 times

mitu7 10 months, 1 week ago
Selected Answer: B
ChatGPT:
No, the solution of configuring Basic gateway credentials for the Azure resource does not meet the goal of configuring back-end authentication for the API Management service instance.

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

To properly configure back-end authentication for the API Management service instance, you need to choose the appropriate authentication method and specify the credentials accordingly. You can find more information about configuring back-end authentication in the Azure API Management documentation.
upvoted 2 times

mitu7 10 months, 1 week ago
ChatGPT:
No, the solution of configuring Basic gateway credentials for the Azure resource does not meet the goal of configuring back-end authentication for the API Management service instance.

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

To properly configure back-end authentication for the API Management service instance, you need to choose the appropriate authentication method and specify the credentials accordingly. You can find more information about configuring back-end authentication in the Azure API Management documentation.
upvoted 1 times

Alluru 12 months ago
Answer - No, because Solution says You configure Basic gateway credentials for the Azure resource. Instead it should be You configure Basic gateway credentials for the HTTP(s) endpoint.
upvoted 4 times

Chang401 1 year ago
Selected Answer: B
NO because we can authenticate using basic-auth but not using basic gateway creds.
upvoted 1 times

[Removed] 1 year, 1 month ago
Selected Answer: A
Basic authentication is supported. Nowhere it says any requirement that says otherwise. People commenting that the question says at the higher level it is not supported etc are really going to remember that kind of details in a broad exam as this?
upvoted 1 times

nhannn 1 year, 1 month ago
Selected Answer: A
I choose A because of this https://learn.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#Basic

upvoted 1 times

Question #25 *Topic 1*

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.
Establish if the solution satisfies the requirements.
You are developing a solution for a public facing API.
The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end.
You must configure back-end authentication for the API Management service instance.
Solution: You configure Client cert gateway credentials for the HTTP(s) endpoint.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
The API back end is hosted in an Azure App Service instance. It is an Azure resource and not an HTTP(s) endpoint.
Reference:
https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-backend-entity

Community vote distribution
A (67%) B (33%)

jay158 Highly Voted 2 years, 3 months ago
This is scenario questions.
If backend is accepts HTTP(S)
Then Basic AUTH or Certificate will work.
so Client Certificate + HTTP(s) YES
upvoted 33 times

r3verse 11 months ago
APIM supports both, but that is only part of the question.

Does App Service itself support basic auth? NO not really (only a hacky way).
Does app Service Support certificate auth? YES, built in the Azure portal directly as a setting for app Service.

So my answer would be YES for client certificate.
upvoted 2 times

aradice Highly Voted 2 years, 4 months ago
yes? https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates
upvoted 13 times

MiraA 2 years ago
The answer is YES (client cert + the HTTP(s) endpoint) based on this screenshot:
https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication

Target: "Azure Logic App" or "HTTP(s) endpoint"
Gateway credentials: "None" or "Basic" or "Client cert"
upvoted 5 times

aradice 2 years, 4 months ago
https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#ClientCertificate
upvoted 4 times

warchoon 8 months, 2 weeks ago
It's the only link that alive. All others are dead and redirect to another pages or target deprecated sources.
upvoted 1 times

BikashKatwal Most Recent 4 months, 3 weeks ago
The solution mentioned, which is configuring Client cert gateway credentials for the HTTP(s) endpoint, does not meet the goal of configuring back-end authentication for the API Management service instance. Client certificate authentication is used for authenticating clients accessing the API Management service, not for authenticating the back-end service itself.

To configure back-end authentication for the API Management service instance hosted in an Azure App Service, you have several options, such as:

API key authentication: Generate an API key and configure it in the API Management service to authenticate requests sent to the back-end service.

OAuth 2.0 authentication: Configure OAuth 2.0 authentication between the API Management service and the back-end service, allowing clients to

authenticate using OAuth 2.0 tokens.

Certificate-based authentication: If your back-end service supports certificate-based authentication, you can configure the API Management service to authenticate requests using client certificates.

Therefore, the correct answer is:

B. No
upvoted 4 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: A

https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates
upvoted 1 times

uffuchsi 8 months, 2 weeks ago

A. Yes
upvoted 1 times

AAcademy 9 months ago

I would say "YES". Secure backend services using client certificate authentication in Azure API Management: https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates
upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. No

Configuring client cert gateway credentials for the HTTP(s) endpoint would not provide a suitable solution for back-end authentication for the public facing API. Client certificate authentication is based on the possession of a certificate and private key, which can be installed on the client device. However, this approach is not suitable for public facing APIs, as it requires the client to have a unique certificate, which is not easy to manage, and it's not a scalable solution for public facing APIs.
upvoted 4 times

shacent 9 months, 1 week ago

chatGPT: This solution does not meet the goal of configuring back-end authentication for the API Management service instance. Client certificate gateway credentials are used for client-side authentication, which is not the same as back-end authentication. Back-end authentication is used to authenticate the API Management service instance with the back-end service hosted in the Azure App Service instance.
For back-end authentication, you can use Azure Active Directory (AAD) authentication, Azure AD B2C, or OAuth 2.0 authentication to authenticate the API Management service instance with the back-end service.
upvoted 2 times

deepbh 10 months ago

A (Yes) is correct.

https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates
"API Management allows you to secure access to the backend service of an API using client certificates." ... "You should have your backend service configured for client certificate authentication. To configure certificate authentication in the Azure App Service, refer to this article."

Then see: https://learn.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth?tabs=azurecli#enable-client-certificates
upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: A

Client certificate is supported. Even if I select an App service while creating the API it shows the back-end authentication option as Https endpoint and not as Azure resource.
upvoted 1 times

SaintBahamut 1 year, 6 months ago

Selected Answer: A

Answer is Yes, its obvious if you check it yourself
Design Tab->Backend panel->HTTP(s) endpoint->And you have None, Basic and Certificate
upvoted 3 times

azure_dunce 1 year, 6 months ago

I think the key phrase in this question is "App Service". I am voting No since the API is hosted in an app service and not APIM.

https://docs.microsoft.com/en-us/azure/app-service/overview-authentication-authorization
upvoted 1 times

leonidn 1 year, 9 months ago

Selected Answer: A

This is one of the possible options
https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication
upvoted 3 times

Mev4953 1 year, 9 months ago

The given link is deprecated.

upvoted 1 times

phvogel 1 year, 12 months ago

The <authentication-cert> policy statement just lets you specify the id or thumbprint of the certificate to use (optionally the body and password to decrypt the body) -- it doesn't let you specify an endpoint (HTTP or HTTPS). I'll vote No.

upvoted 1 times

Omar102 2 years, 1 month ago

Only ClientCert, Basic is not supported anymore

upvoted 2 times

koolexam 1 year, 10 months ago

What makes you say so? It is supported.

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

upvoted 2 times

BishopeL 2 years, 2 months ago

The answer is 'NO'.

It is an Azure resource and not an HTTP(s) endpoint.

upvoted 3 times

Question #26 *Topic 1*

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.
Establish if the solution satisfies the requirements.
You are developing a solution for a public facing API.
The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end.
You must configure back-end authentication for the API Management service instance.
Solution: You configure Basic gateway credentials for the HTTP(s) endpoint.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
API Management allows to secure access to the back-end service of an API using client certificates. Furthermore, the API back end is hosted in an Azure App
Service instance. It is an Azure resource and not an HTTP(s) endpoint.
Reference:
https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-backend-entity

Community vote distribution
B (50%) A (50%)

jay158 Highly Voted 2 years, 3 months ago
This is scenario questions.
If backend is accepts HTTP(S)
Then Basic AUTH or Certificate will work.
so Basic + HTTPS Yes
upvoted 31 times

MiraA 2 years ago
Target: "Azure Logic App" or "HTTP(s) endpoint"
Gateway credentials: "None" or "Basic" or "Client cert"

https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication
upvoted 6 times

ndchris2003 Highly Voted 2 years, 3 months ago
The correct answer is Yes.

https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#ClientCertificate
upvoted 10 times

garbas Most Recent 1 week, 6 days ago
Selected Answer: A
API Management gateway supports basic authentication.
You have a RESTful service, so there is an HTTP(s) endpoint.
App Service doesn't have built-in support fot basic authentication, but you can use the security features of the framework the service is built on (https://learn.microsoft.com/en-us/azure/app-service/overview-authentication-authorization#why-use-the-built-in-authentication)

The question only says "you must configure authentication". It doen't says it has to be the best solution or the most secure.
Maybe Microsoft is expecting No in this question, but the solution meets the goal, so the answer should be Yes.
upvoted 1 times

juanHuizar 1 month, 1 week ago
No, configuring Basic gateway credentials for the HTTP(s) endpoint does not fully meet the goal of configuring backend authentication for the API Management service instance, especially if this API is public-facing and needs to be secure.

Basic gateway credentials refer to HTTP Basic Authentication, where a client sends a username and password with each request.
upvoted 1 times

TheMCT 1 month, 3 weeks ago
Selected Answer: B
Configuring Basic gateway credentials for the HTTP(s) endpoint in API Management is not typically used for back-end authentication. Basic gateway credentials are used to secure access to the API Management gateway itself, allowing clients to authenticate when making requests to the API Management layer. It does not secure the communication between API Management and the back-end Azure App Service.

upvoted 1 times

BikashKatwal 4 months, 3 weeks ago

The solution mentioned, which is configuring Basic gateway credentials for the HTTP(s) endpoint, does meet the goal of configuring back-end authentication for the API Management service instance.

Configuring Basic gateway credentials means that the API Management service will authenticate requests sent to the back-end service using basic authentication. Basic authentication requires clients to include a username and password in the request headers.

By configuring Basic gateway credentials for the HTTP(s) endpoint, you can enforce authentication for incoming requests and ensure that only authorized clients can access the API back end hosted in the Azure App Service instance.

Therefore, the correct answer is:

A. Yes

upvoted 2 times

Net_IT 8 months, 2 weeks ago

Selected Answer: A

Answer should be YES according to this link: https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication

upvoted 1 times

Michael2023 9 months, 1 week ago

Selected Answer: B
Because the Backend is an App Service the Microsoft security recommendations to protect back-end resources are: User/Application Identity and/or require client certificate authentication

APIM would support Basic & Client, but it also depends on which kind of backend application you are using

https://learn.microsoft.com/en-us/azure/app-service/security-recommendations#identity-and-access-management

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. No

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

upvoted 3 times

deepbh 10 months ago

The backend is App Service. Afaik it doesn't allow basic auth, so your only choice is client cert or managed identity.

It shouldn't matter if you have the option to enable basic auth at the frontend settings, you won't be able to make it work at the backend.

So I'd say NO.

upvoted 4 times

Lucifer14 10 months ago

In 28-12-2022 exam

upvoted 1 times

Lucifer14 10 months ago

In 28-12-2022 exam, but in bit diff way.

upvoted 2 times

narenazure 10 months, 1 week ago

Selected Answer: A

It should be YES.
Refer the followings : https://learn.microsoft.com/en-us/azure/api-management/media/api-management-howto-mutual-certificates/apim-client-cert-enable-select.png

upvoted 1 times

r3verse 11 months ago

Selected Answer: B

APP SERVICE ITSELF DOESNT HAVE BUILT IN SUPPORT FOR BASIC AUTHENTICATION, IT DOESNT MATTER THEN IF YOU FOUND A LINK SAYING API MANAGEMENT SUPPORTS BASIC AUTH THEN, DOES IT? B B B B

upvoted 2 times

hubekpeter 11 months, 2 weeks ago

Selected Answer: B

This is azure native resource. Managed identity is preferred over anything else, doesn't matter if it's mutual authentication, basic etc. !!! The question emphasizes on a basic iam comprehension.

upvoted 1 times

adilkhan 11 months, 2 weeks ago

https://learn.microsoft.com/en-us/azure/api-management/api-management-authentication-policies Should be YES

upvoted 1 times

DusanS 1 year, 1 month ago

I think it is referred to self hosted gateway. There when you create new one you have section configure and option to set certificate.

upvoted 1 times

Question #27 *Topic 1*

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.
Establish if the solution satisfies the requirements.
You are developing a solution for a public facing API.
The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end.
You must configure back-end authentication for the API Management service instance.
Solution: You configure Client cert gateway credentials for the Azure resource.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*
API Management allows to secure access to the back-end service of an API using client certificates.
Reference:
https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-backend-entity

Community vote distribution
B (58%) A (42%)

jay158 Highly Voted 2 years, 3 months ago
This is scenario questions.
If backend is accepts HTTP(S)
Then Basic AUTH or Certificate will work.
so Certificate + Azure Resource NO
https://www.youtube.com/watch?v=HQ0U7lwP93o
upvoted 23 times

MiraA 2 years ago
Target: "Azure Logic App" or "HTTP(s) endpoint"
Gateway credentials: "None" or "Basic" or "Client cert"

https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication
upvoted 6 times

PhilLI 1 year, 9 months ago
https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth
So I think it's YES
upvoted 3 times

warchoon 8 months, 2 weeks ago
No. Your page is not about backend. It's about CLIENT certificate authentication
upvoted 1 times

warchoon 8 months, 2 weeks ago
Sorry my joke. But it is true.
upvoted 1 times

TheMCT Most Recent 1 month, 3 weeks ago
Selected Answer: B
The answer is; NO
Configuring "Client cert gateway credentials" for the Azure resource suggests that you are configuring client certificate authentication for the API Management gateway itself when accessing Azure resources, but it may not necessarily secure the communication between API Management and the back-end Azure App Service.
upvoted 3 times

JH81 4 months, 1 week ago
Selected Answer: A
One of the questions is for securing with Basic and the other with a client cert. Clearly MS is saying Basic = bad and client cert = good.
upvoted 3 times

BikashKatwal 4 months, 3 weeks ago
A. Yes

The solution meets the goal of configuring back-end authentication for the API Management service instance. By configuring Client cert gateway

credentials for the Azure resource, you can authenticate and secure the communication between the API Management gateway and the back-end service hosted in the Azure App Service instance.

Client certificate authentication involves using X.509 client certificates to verify the identity of the client requesting access to the API Management gateway. This authentication method ensures that only clients with valid certificates can access the back-end service.

By configuring Client cert gateway credentials, you establish a secure and authenticated connection between the API Management gateway and the Azure App Service instance, thereby meeting the requirement for back-end authentication.

upvoted 3 times

basquiat 2 months, 3 weeks ago

If I got it right the link below shows how to achieve that. https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates

upvoted 1 times

deathRac3 5 months, 1 week ago

Selected Answer: A

Looks right

upvoted 1 times

forgetfulalligator 8 months, 1 week ago

Selected Answer: A

When client certificate authentication is used for the API Management gateway, a client must provide a client certificate in the header of each request to access the API. The client certificate is then validated by the gateway before forwarding the request to the back-end API. This can provide a high level of authentication and help prevent unauthorized access to the back-end API.

Client certificate authentication is a more secure mechanism than Basic authentication, as it uses a secure communication channel and cryptographic keys to authenticate clients. Additionally, it provides better scalability, as it does not require the gateway to maintain a large number of user credentials.

upvoted 3 times

mabdo 8 months, 3 weeks ago

Are the website answers the official ones? This is getting confusing!!!

upvoted 3 times

joshfry 8 months, 1 week ago

No. They don't have access to official answers. They try their best, but always read the discussion.

upvoted 2 times

Michael2023 9 months, 1 week ago

Selected Answer: A -> yes

as long as a certificate is used to can authenticate to the backend via "Custom URL" or "Azure resource" (Function App, Logic App, Web App)... in both cases you provide an URL to which the requests are forwared to.
From the App Service side you should use HTTPS only or User/System Managed Identity to authenticate to the App Service.
So that means, the only solutions for this scenario would be using a certificate for HTTP(s) endpoint and resources.

https://learn.microsoft.com/en-us/azure/app-service/security-recommendations#identity-and-access-management
https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates

upvoted 2 times

warchoon 8 months, 2 weeks ago

Logic App is ok. But it is not the App Service instance. Others use gateway Https endpoint settings.

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. No

Configuring client cert gateway credentials for the Azure resource would not provide a suitable solution for back-end authentication for the public facing API. Client certificate authentication is based on the possession of a certificate and private key, which can be installed on the client device. However, this approach is not suitable for public facing APIs, as it requires the client to have a unique certificate, which is not easy to manage, and it's not a scalable solution for public facing APIs.

upvoted 3 times

deepbh 10 months ago

Selected Answer: B

Azure Logic App and HTTP(S) Endpoint are the two choices on the portal page. Whenever the question is about using app service at the backend, and it talks about configuring auth for the "azure resource" instead of HTTP(S), it's an automatic NO.

upvoted 3 times

HellBaby 10 months ago

mTLS is used for AuthZ, and not AuthN - is like presenting your ID to the policeman, but not giving him the house keys.

So the answer is for sure NO!

upvoted 1 times

rotimislaw 10 months ago

Selected Answer: A

I vote for A

upvoted 1 times

HaoTPV 11 months ago

Selected Answer: A

A is correct

upvoted 1 times

HaoTPV 11 months ago

Selected Answer: A

A is correct

upvoted 1 times

mum_in_tech 11 months, 1 week ago

YES. API Management Service > Backends > Azure REsource (App Service instance), and setup the url > Client Certificate

upvoted 1 times

hpjjphhpj 1 year, 1 month ago

Selected Answer: B

oNLY FOR HTTP

upvoted 1 times

SaintBahamut 1 year, 6 months ago

Selected Answer: B

This can be configured for an https endpoint (and logic apps in the past?)
So - No

upvoted 2 times

Question #28 *Topic 1*

You are developing a .NET Core MVC application that allows customers to research independent holiday accommodation providers.
You want to implement Azure Search to allow the application to search the index by using various criteria to locate documents related to accommodation venues.
You want the application to list holiday accommodation venues that fall within a specific price range and are within a specified distance to an airport.
What should you do?

A. Configure the SearchMode property of the SearchParameters class.

B. Configure the QueryType property of the SearchParameters class.

C. Configure the Facets property of the SearchParameters class.

D. Configure the Filter property of the SearchParameters class.

Correct Answer: *D*
The Filter property gets or sets the OData $filter expression to apply to the search query.
Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.querytype

Community vote distribution
D (100%)

7ack Highly Voted 2 years, 4 months ago
Azure search was in AZ-203, but it's not in the AZ-204 list of measured skills. Is this even still relevant?
upvoted 19 times

Bogdan75 1 year, 7 months ago
You're right, it's not on the list of measured skills: https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B
upvoted 3 times

hstml Highly Voted 2 years, 1 month ago
The solution D - Filter is correct.
The thing is that you filter the queried information by a price range and distance to airport. The QueryType is about the syntax of what you can pass into the search.
Read the explanations in the docs: https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters?view=azure-dotnet
upvoted 12 times

MiraA 2 years ago
Geo-search example:
https://docs.microsoft.com/en-us/azure/search/search-query-simple-examples#example-6-geo-search
upvoted 4 times

PhilLI 1 year, 9 months ago
So [A] should be the answer then?
https://docs.microsoft.com/en-us/azure/search/query-simple-syntax
The searchMode parameter is relevant in this example. Whenever boolean operators are on the query, you should generally set searchMode=all to ensure that all of the criteria is matched. Otherwise, you can use the default searchMode=any that favors recall over precision.
upvoted 1 times

PhilLI 1 year, 9 months ago
I see I'm wrong ...
upvoted 1 times

aruni_mishra 1 year, 9 months ago
correct link:
https://docs.microsoft.com/en-us/azure/search/search-query-simple-examples#example-6-geospatial-search
upvoted 1 times

forgetfulalligator Most Recent 8 months, 1 week ago
Selected Answer: D
The Filter property is used to narrow down search results by specifying one or more filter expressions to apply to the search query. A filter expression is a Boolean expression that is used to match specific criteria, such as a specific price range or proximity to an airport.
upvoted 2 times

wrak 8 months, 2 weeks ago

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

upvoted 3 times

examinator007 9 months ago

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

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: D

D. Configure the Filter property of the SearchParameters class.

To allow the application to list holiday accommodation venues that fall within a specific price range and are within a specified distance to an airport, you should configure the Filter property of the SearchParameters class in your .NET Core MVC application. The filter property allows you to filter the search results by applying conditions to specific fields in the indexed documents. You can use this property to specify the price range and distance range of the accommodation venues.

upvoted 1 times

KingChuang 10 months, 4 weeks ago

Selected Answer: D

My Answer:D

Ref:https://learn.microsoft.com/en-us/azure/search/search-query-simple-examples#example-5-range-filters

upvoted 1 times

Chang401 1 year ago

Selected Answer: D

similar type of question is there on examtopics az-204. but the ask is different in both the questions.

upvoted 2 times

PhilLI 1 year, 9 months ago

Selected Answer: D

checkout MiraA's link

upvoted 3 times

ning 2 years, 1 month ago

querytype cannot work for geo search, it is D, filter with odata can work with geo data

upvoted 2 times

BishopeL 2 years, 2 months ago

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.filter?view=azure-dotnet#Microsoft_Azure_Search_Models_SearchParameters_Filter

upvoted 1 times

BishopeL 2 years, 2 months ago

Correct answer is D.

upvoted 3 times

hussamAlHunaiti 2 years, 2 months ago

Wrong answer, the correct answer is B- Querytype property

upvoted 2 times

wolf_lu 2 years, 3 months ago

choose B ,querytype

upvoted 2 times

Question #29 *Topic 1*

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

A. the Enterprise Integration Pack (EIP)

B. the Logic App Code View

C. the API Connections

D. the Logic Apps Designer

Correct Answer: *A*
For business-to-business (B2B) solutions and seamless communication between organizations, you can build automated scalable enterprise integration workflows by using the Enterprise Integration Pack (EIP) with Azure Logic Apps.
Reference:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-overview https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions

Community vote distribution
D (75%) B (23%)

Nokaido Highly Voted 1 year, 1 month ago
Selected Answer: D
For definitions use the Code View, for the Workflows use the Designer I think
upvoted 30 times

Ming0808 12 months ago
yes. Test in Azure
upvoted 3 times

ryvel Highly Voted 1 year, 5 months ago
Since the question says developer, I think B(the Logic App Code View) should be the answer
upvoted 11 times

dddddd111 Most Recent 1 week, 1 day ago
chatGPT chose B and D.

You can indeed use both the Logic App Code View (option B) and the Logic Apps Designer (option D) to edit workflows for an existing Logic App. The choice between them depends on your preference and the specific requirements of your task.

Logic App Code View (option B): This allows you to edit the Logic App's underlying JSON definition directly. It provides more granular control and is suitable for developers who prefer working with JSON or need to make intricate changes to the Logic App's structure.

Logic Apps Designer (option D): This provides a visual interface to design Logic Apps using pre-built triggers, actions, and connectors. It's a more user-friendly option for those who prefer a graphical approach and don't need to work directly with the JSON definition.

Both options are valid tools for editing workflows, and the choice depends on your familiarity with JSON, your specific task requirements, and personal preference.
upvoted 1 times

garbas 1 week, 5 days ago
Selected Answer: D
Both Code View and Designer can be used to edit the workflow.
But it's easier to edit the workflow in the designer, so I think that is what we *should* use.
We shouldn't use the more complicated tool just because we are developers.
upvoted 1 times

HimGarg 4 weeks ago
Edit Workflow, so Designer
upvoted 1 times

heatfan900 1 month ago
D. the Logic Apps Designer
upvoted 1 times

RikinPatel 1 month, 2 weeks ago

May be "Enterprise Integration Pack (EIP)" correct ans as per this portal because we don't know context of this que. in exam may be given some case study before this que.

upvoted 1 times

sen_su 1 month, 3 weeks ago

How to understand the given explanation "For business-to-business (B2B) solutions and seamless communication between organizations, you can build automated scalable enterprise integration workflows by using the Enterprise Integration Pack (EIP) with Azure Logic Apps." to edit the workflow for the logic App?

upvoted 1 times

CellCS 3 months, 2 weeks ago

Logic App Code View because your are developer. designer is for "designer" who do not need has coding skills

upvoted 1 times

tiagodfsouza 3 months, 2 weeks ago

Selected Answer: A

Shoud be A

upvoted 1 times

viking2 4 months ago

Selected Answer: D

There is nothing in the question about EIP.
I'm a developer and I still use the Designer.
I'd answer D

upvoted 2 times

deathRac3 5 months, 1 week ago

Selected Answer: D

I am developer and I prefer designer view for quick edit

upvoted 3 times

MikeAWS 1 month ago

it is not what you prefer, it is how should be the correct and right way to do this!

upvoted 1 times

deathRac3 5 months, 1 week ago

Selected Answer: D

I am .NET developer, still I like to use Designer in logic apps. Saves time

upvoted 1 times

Mattt 7 months ago

D. the Logic Apps Designer.

To edit the workflows for an existing Logic App, the Logic Apps Designer is the most appropriate tool to use. The Logic Apps Designer is a visual interface for creating and editing Logic Apps. It provides a drag-and-drop interface for building workflows using a wide range of connectors and actions.

Using the Logic Apps Designer, you can easily modify existing Logic Apps by adding, removing, or modifying connectors, actions, and conditions. You can also configure triggers and inputs, and test the Logic App to ensure that it works as expected.

The Enterprise Integration Pack (EIP) is a set of connectors and tools that can be used to integrate on-premises systems with cloud services. It is not specifically designed for editing Logic App workflows.

The Logic App Code View is a text-based view of the Logic App definition. While it can be used to modify the Logic App definition directly, it requires a solid understanding of the underlying JSON schema.

API Connections are used to connect to external services and systems, but they are not used for editing Logic App workflows.

upvoted 1 times

jimmyyml 8 months, 1 week ago

Selected Answer: D

The correct answer to this question is D, the Logic Apps Designer. The Logic Apps Designer is a visual, low-code interface that allows developers to create, edit, and manage Logic Apps without having to write code. It provides a drag-and-drop interface that enables developers to add and configure connectors, actions, and triggers to create a workflow. The Logic App Code View is an advanced feature that enables developers to edit the underlying code of a Logic App, but it is not required for basic editing of workflows. The Enterprise Integration Pack (EIP) is a set of connectors that provides additional capabilities for connecting to on-premises systems and services, but it is not required for editing workflows. The API Connections are used to securely connect to APIs and services, but they do not provide editing capabilities for Logic App workflows.

upvoted 2 times

wrak 8 months, 2 weeks ago

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

upvoted 4 times

warchoon 8 months, 2 weeks ago

A, B and D are usable in different scenarios. So i think that the question is not complete.

upvoted 1 times

Question #30 *Topic 1*

DRAG DROP -

You are a developer for a company that provides a bookings management service in the tourism industry. You are implementing Azure Search for the tour agencies listed in your company's solution.

You create the index in Azure Search. You now need to use the Azure Search .NET SDK to import the relevant data into the Azure Search service.

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

Select and Place:

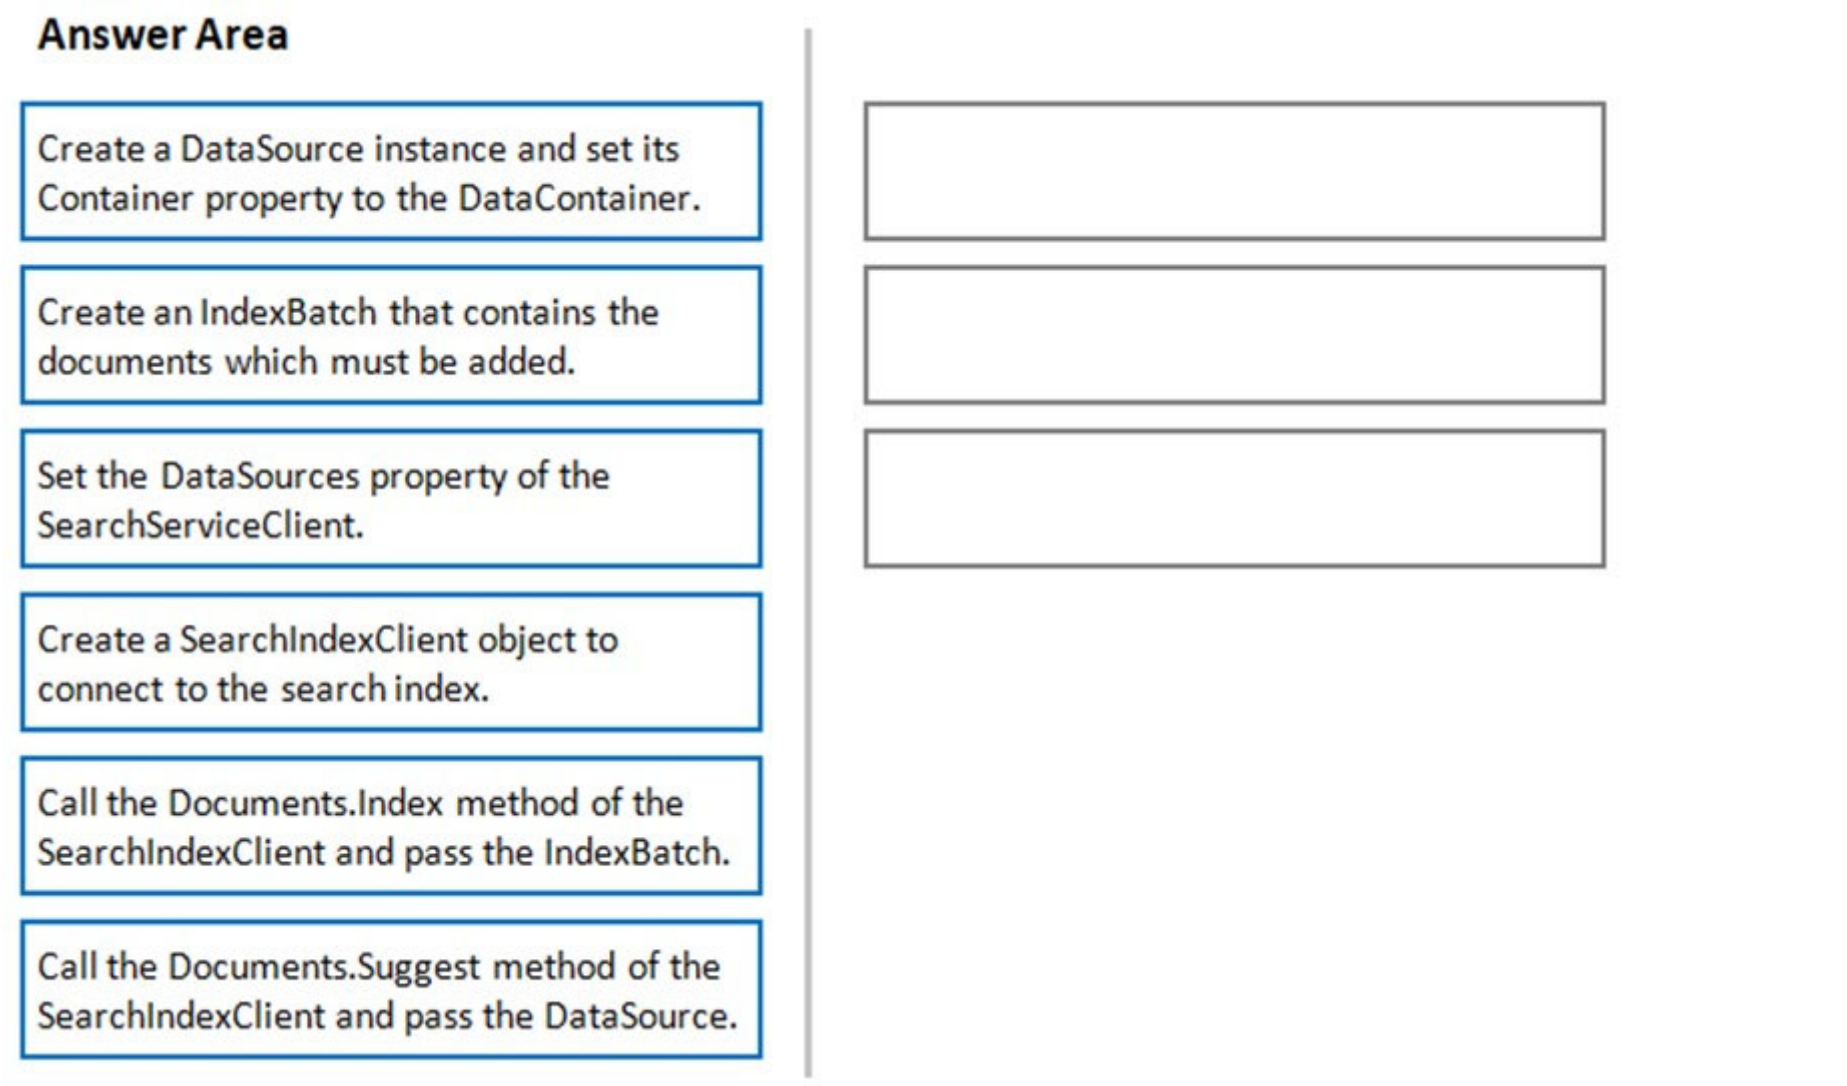

Correct Answer:

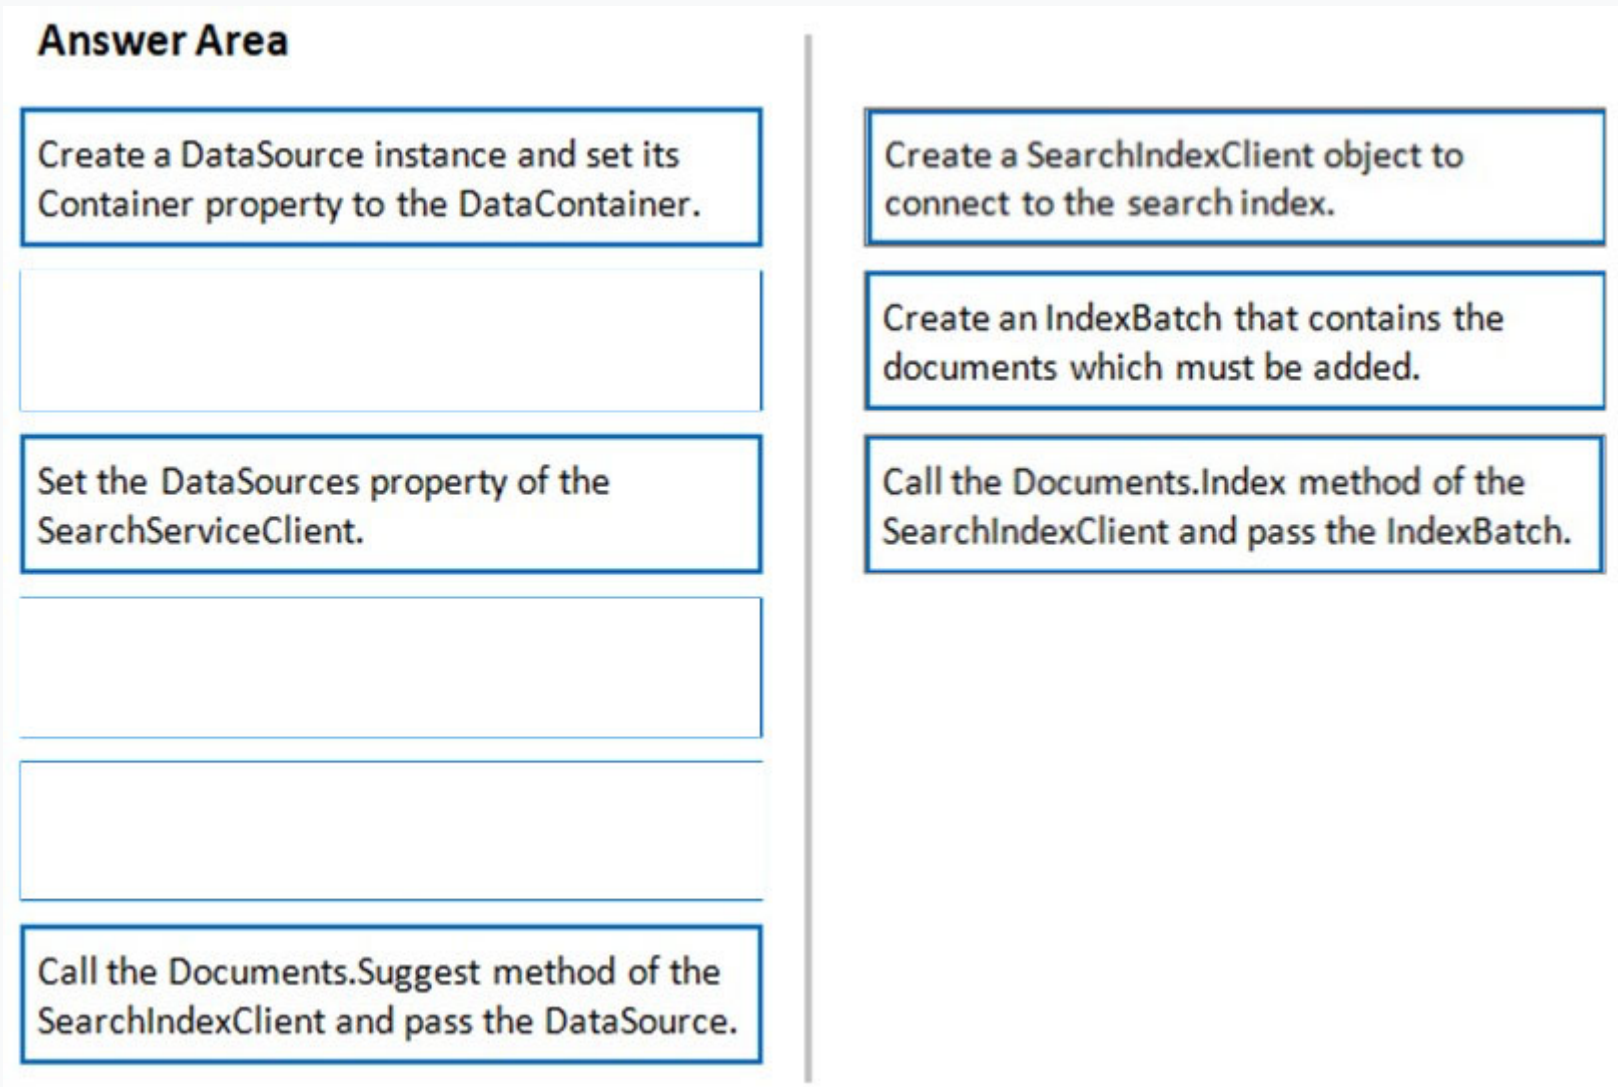

1. The index needs to be populated. To do this, we will need a SearchIndexClient. There are two ways to obtain one: by constructing it, or by calling

Indexes.GetClient on the SearchServiceClient. Here we will use the first method.

2. Create the indexBatch with the documents

Something like:
var hotels = new Hotel[];
{
new Hotel()
{
HotelId = "3",
BaseRate = 129.99,
Description = "Close to town hall and the river"
}
};
ג€¦
var batch = IndexBatch.Upload(hotels);
3. The next step is to populate the newly-created index
Example:
var batch = IndexBatch.Upload(hotels);
try
{
indexClient.Documents.Index(batch);
}
Reference:
https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk

[Removed] Highly Voted 1 year, 1 month ago

Azure search is out of scope for AZ-204

upvoted 23 times

Zinat Most Recent 1 month ago

azure search is deprecated

upvoted 1 times

moota 5 months, 1 week ago

It's also possible the indexClient comes after creating the batch.

// Assume that you have a collection of documents called "documents"
List<MyDocumentType> documents = GetDocuments(); // Replace MyDocumentType with the actual type of your documents

// 1. Create an IndexBatch that contains the documents which must be added.
IndexBatch<MyDocumentType> batch = IndexBatch.Upload(documents);

// 2. Create a SearchIndexClient object to connect to the search index.
// Assume that searchIndexName is the name of your Azure Search index
SearchIndexClient indexClient = serviceClient.Indexes.GetClient(searchIndexName);

upvoted 1 times

tom112 5 months, 3 weeks ago

it's SearchClient.IndexDocuments(batch) method, actually.

upvoted 1 times

NombreFalso 8 months, 2 weeks ago

😫😫😫 Bad question

upvoted 1 times

wrak 8 months, 2 weeks ago

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

upvoted 4 times

warchoon 8 months, 2 weeks ago

Looks correct
https://learn.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk#create-an-index

upvoted 1 times

Discovery_ 12 months ago

Azure Search doesn't seem to be in the latest study guide. Best to avoid this question for less cram.

upvoted 3 times

coffecold 1 year, 1 month ago

Can't be true : uploading the data by putting it hardcoded in the source? Would suggest the solutions that use a datasource. After all you had to extract the data from the "company's solution"

upvoted 2 times

warchoon 8 months, 2 weeks ago

It's a simple prototype. TDD, baby.

upvoted 1 times

sghaha 1 year, 6 months ago

https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk

upvoted 2 times

Question #31 *Topic 1*

You are developing an application that applies a set of governance policies for internal and external services, as well as for applications.
You develop a stateful ASP.NET Core 2.1 web application named PolicyApp and deploy it to an Azure App Service Web App. The PolicyApp reacts to events from
Azure Event Grid and performs policy actions based on those events.
You have the following requirements:
- Authentication events must be used to monitor users when they sign in and sign out.
- All authentication events must be processed by PolicyApp.
- Sign outs must be processed as fast as possible.

What should you do?

A. Create a new Azure Event Grid subscription for all authentication events. Use the subscription to process sign-out events.

B. Create a separate Azure Event Grid handler for sign-in and sign-out events.

C. Create separate Azure Event Grid topics and subscriptions for sign-in and sign-out events.

D. Add a subject prefix to sign-out events. Create an Azure Event Grid subscription. Configure the subscription to use the subjectBeginsWith filter.

Correct Answer: *D*

Reference:

https://docs.microsoft.com/en-us/azure/event-grid/subscription-creation-schema

Community vote distribution

C (66%) D (34%)

ning Highly Voted 2 years, 2 months ago

Only C is mentioned both topic and subscription, which are two critical parts for event grid, so I will go C, anyway, this question is very poorly worded

upvoted 33 times

BishopeL Highly Voted 2 years, 2 months ago

For your interest, please see https://docs.microsoft.com/en-us/azure/event-grid/event-filtering

Correct answer is D.

upvoted 12 times

edengoforit 1 year, 9 months ago

D doesn't contain any action of Sign-in. I am more inclined to C

upvoted 7 times

ucsdmiami2020 1 year, 10 months ago

I agree with D per the following text found in Microsoft docs URL https://docs.microsoft.com/en-us/azure/event-grid/concepts

"A subscription tells Event Grid which events on a topic you're interested in receving. When creating the subscription, you provide an endpoint for handling the vent. You can filter the events that are sent to the endpoint. You can filter by event type or subject pattern."

upvoted 2 times

Mick1z8 Most Recent 1 week, 5 days ago

Selected Answer: C

Filtering is something that you can use regardless (which adds a small overhead though). But the question says "sign-outs should be as fast as possible" which means creating a dedicated topic only for "sign-outs" is the best possible solution.

upvoted 1 times

INGAKULA 1 month, 2 weeks ago

The correct answer is C

upvoted 1 times

MVCChicago 2 months ago

Answer is D

upvoted 2 times

LxstInNxght 2 months, 2 weeks ago

Selected Answer: C

It's C

upvoted 1 times

jakobste 2 months, 2 weeks ago

I think i have seen this before and the answer was C on that practice test.

upvoted 1 times

xmashoodx 2 months, 3 weeks ago

"chatgpt"

The requirement is to process authentication events, specifically sign-in and sign-out, and make sure that sign-outs are processed as quickly as possible. You want to be able to filter events specifically for this purpose.

Option D will fulfill these requirements. By adding a subject prefix to sign-out events and creating an Azure Event Grid subscription that is configured to use the subjectBeginsWith filter, you can specifically target and process the sign-out events quickly. The subjectBeginsWith filter allows for efficient filtering of events, which aligns with the requirement of processing sign-outs as fast as possible.

Therefore, the correct answer is:

D. Add a subject prefix to sign-out events. Create an Azure Event Grid subscription. Configure the subscription to use the subjectBeginsWith filter.

upvoted 3 times

basquiat 2 months, 3 weeks ago

I tend to agree, especially bc D is the only alternative that explores a more advanced capability of Event Grid, such as subject filtering. And this is a core thing for MS exams. However, this is an awful question, so many topics mentioned at once let things messy.

upvoted 1 times

rameeze 2 months, 3 weeks ago

are you preparing for the exam from here??

upvoted 2 times

macobuzi 2 months, 2 weeks ago

Hey, you too preparing for the exam? So many vague answers to these questions but i couldn't find anywhere better to prepare.

upvoted 2 times

sthmishra 2 months, 1 week ago

i am preparing from this site too, please mention any other source if you know

upvoted 1 times

Vaibhab 3 months, 3 weeks ago

Answer is D:
By adding a subject prefix to sign-out events, you can differentiate them from other authentication events. Then, you can create an Azure Event Grid subscription and configure it to use the subjectBeginsWith filter. This filter allows you to specify that you only want to process events with a subject that begins with the specified prefix, in this case, the sign-out events.

upvoted 1 times

viking2 3 months, 4 weeks ago

Selected Answer: C

Separate topics with a subscription for each allows both sign-ins and sign-outs to be handled, and a subscriber to sign-outs don't need to wait for any sign-in events to be processed first.
My answer is C

upvoted 1 times

[Removed] 4 months, 1 week ago

shit question

upvoted 2 times

MachindraGholap 4 months, 3 weeks ago

Option C :
Option D does not cover the the sign-in event . there is also requirement of processing signout event as fast as possible. creating separate subscription will be my choice.

upvoted 3 times

bantawit 5 months ago

Selected Answer: C

To meet the requirements of monitoring user sign-ins and sign-outs, processing all authentication events in PolicyApp, and ensuring fast processing of sign-outs, you should:

C. Create separate Azure Event Grid topics and subscriptions for sign-in and sign-out events.

By creating separate Azure Event Grid topics and subscriptions for sign-in and sign-out events, you can have dedicated channels to receive and process each type of event independently. This approach allows PolicyApp to handle authentication events efficiently and respond quickly to sign-out events without being affected by other types of events.

Therefore, option C is the most appropriate choice for this scenario.

upvoted 3 times

kabbas 5 months ago

I am more inclined towards D as you do not have to worry about sign ins
upvoted 1 times

viking2 3 months, 4 weeks ago
But you do need to worry about sign-ins, as both sign-ins and sign-outs need to be handled. Separate topics and a subscription for each, so C
upvoted 1 times

zulqarnainali 5 months, 1 week ago
C is the correct answer in my opinion
upvoted 1 times

deathRac3 5 months, 1 week ago
Selected Answer: C
To process sign-out ASAP, you should create separate subscriptions. So, C.
upvoted 2 times

tom112 5 months, 3 weeks ago
Selected Answer: C
"A topic is used for a collection of related events. To respond to certain types of events, subscribers decide which topics to subscribe to."
So, by adding more topics, you make other collections - it would work as having two different queues in this question. Then at the code level of the app, it can try to subscribe sing-out topic first.
upvoted 2 times

Question #32 *Topic 1*

HOTSPOT -

You are developing a C++ application that compiles to a native application named process.exe. The application accepts images as input and returns images in one of the following image formats: GIF, PNG, or JPEG.

You must deploy the application as an Azure Function.

You need to configure the function and host json files.

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

NOTE: Each correct selection is worth one point.

Hot Area:

function.json

```
{
  [ ▼ ]
    "direction": "out",
    "name" : "result"
}
```

- "type": "http"
- "platform": "gcm"
- "datatype": "stream"
- "path": "process.exe"

host.json

```
[ ▼ ]
    "defaultExecutablePath": "process.exe"
},
[ ▼ ]
}
```

- "customHandler": { "description": {
- "languageWorker": { "path": {
- "extensions": {"worker": {
- "extensionBundle": {

- "enableForwardingHttpRequest": true
- "enableForwardingHttpRequest": false

Correct Answer:

function.json

```
{
  [ “type”: “http” | “platform”: “gcm” | “datatype”: “stream” | “path”: “process.exe” ]
  “direction”: “out”,
  “name” : “result”
}
```

host.json

```
[ “customHandler”: { “description”: { | “languageWorker”: { “path”: { | “extensions”: {“worker”: { | “extensionBundle”: { ]
    “defaultExecutablePath”: “process.exe”
},
[ “enableForwardingHttpRequest”: true | “enableForwardingHttpRequest”: false ]
}
```

Box 1: "type": "http"

Box 2: "customHandler": { "description":{

A custom handler is defined by configuring the host.json file with details on how to run the web server via the customHandler section.

The customHandler section points to a target as defined by the defaultExecutablePath.

Example:

"customHandler": {

"description": {

"defaultExecutablePath": "handler.exe"

Box 3: "enableForwardingHttpRequest": false

Incorrect:

For HTTP-triggered functions with no additional bindings or outputs, you may want your handler to work directly with the HTTP request and response instead of the custom handler request and response payloads. This behavior can be configured in host.json using the enableForwardingHttpRequest setting.

At the root of the app, the host.json file is configured to run handler.exe and enableForwardingHttpRequest is set to true.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers

seb142 Highly Voted 12 months ago

On my exam 2022-11-02

upvoted 22 times

kay000001 3 months, 2 weeks ago

NOTHING infuriates me more than statements like this.

upvoted 3 times

OPT_001122 11 months, 2 weeks ago

Thanks for mentioning the date

upvoted 12 times

[Removed] Highly Voted 1 year, 1 month ago

I just hope the real questions are not this deep to know if enableForwardingHttpRequest is true or false. I use Azure functions every day but I never used a custom handler and I hope most people do not have the burning need to use this when there are so many other solutions to handle this.

upvoted 20 times

ENGs Most Recent 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

jakobste 2 months, 2 weeks ago

This was on my exam in august 2023.

upvoted 7 times

AbidooKing 3 months, 2 weeks ago

On my exam 2023-07-13- went with "enableForwardingHttpRequest": true. Score 917

upvoted 5 times

Gednehm 4 months ago

https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers#, the last one should be true.

upvoted 3 times

moota 5 months, 1 week ago

- Your Azure Function will be a custom handler since Azure Functions doesn't directly support C++ as a runtime.
- enableForwardingHttpRequest is set to true to allow the custom handler to receive the HTTP request data

upvoted 4 times

surprise0011 6 months, 2 weeks ago

I will go with 'yes'

In the function.json file, enableForwardingHttpRequest should be set to true for the Azure Function to accept HTTP requests and forward them to the C++ application for processing.

This is because the C++ application is compiled as a native application and cannot directly receive HTTP requests. By setting enableForwardingHttpRequest to true, the Azure Function acts as a proxy and forwards incoming HTTP requests to the C++ application.

So the correct option is:
"enableForwardingHttpRequest": true

upvoted 4 times

Videira 8 months ago

On my exam 2023-02-25

upvoted 7 times

RonanMatos 8 months ago

shit question

upvoted 2 times

forgetfulalligator 8 months, 1 week ago

Everything looks gucci except enableforwardhttprequest got to be true. If it was false, our customer handler can receiving any incoming HTTP requests

upvoted 2 times

AAcademy 8 months, 3 weeks ago

enableForwardingHttpRequest should be true

When enableForwardingHttpRequest is true, the behavior of HTTP-only functions differs from the default custom handlers behavior in these ways:

The HTTP request does not contain the custom handlers request payload. Instead, the Functions host invokes the handler with a copy of the original HTTP request.
The Functions host invokes the handler with the same path as the original request including any query string parameters.
The Functions host returns a copy of the handler's HTTP response as the response to the original request.

Source: https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers

upvoted 5 times

Asten 9 months ago

The third selection should be true
----'At the root of the app, the host.json file is configured to run handler.exe and enableForwardingHttpRequest is set to true'

upvoted 4 times

Esward 9 months, 1 week ago

HTTP, Custom handler, true are the correct answers

upvoted 5 times

Knightie 10 months ago

enable forward should be true because you are not returning the image stream, instead if you take in an http request to load a .exe process, the image should be generated somewhere and you will then need to forward to that location to get the image to complete the process functionally.

upvoted 1 times

mdg3501 10 months, 3 weeks ago

got this on 2022-12-7

upvoted 4 times

Chinks098 10 months, 1 week ago

what were your answer for enableForwardingHttpRequest ?

upvoted 3 times

hubekpeter 11 months, 2 weeks ago

enableForwardingHttpRequest Should be true, you have an http stream output declared in the example.

upvoted 4 times

Question #33 *Topic 1*

HOTSPOT
-

You are developing an Azure Static Web app that contains training materials for a tool company. Each tool's training material is contained in a static web page that is linked from the tool's publicly available description page.

A user must be authenticated using Azure AD prior to viewing training.

You need to ensure that the user can view training material pages after authentication.

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

NOTE: Each correct selection is worth one point.

Answer Area

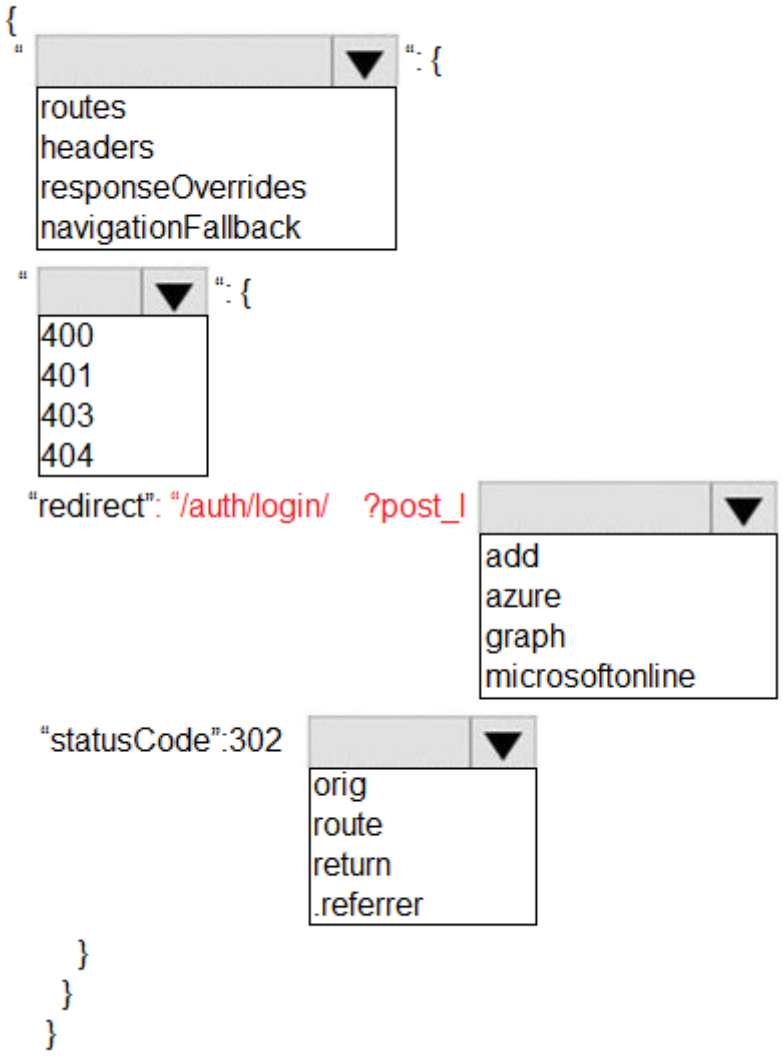

Correct Answer:

Answer Area

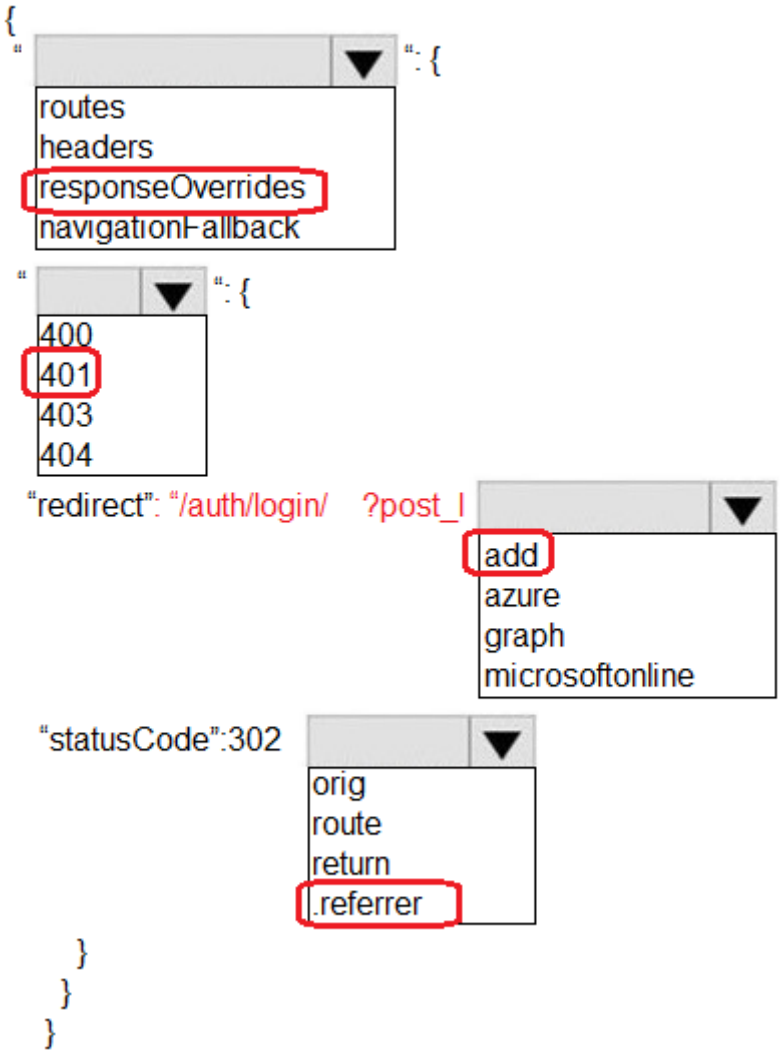

[Removed] Highly Voted 7 months, 4 weeks ago

Got it in the exam 7/03/23. Went with the given answer

upvoted 7 times

chettir01 Highly Voted 9 months, 3 weeks ago

That's correct :
https://learn.microsoft.com/en-us/azure/static-web-apps/configuration

upvoted 6 times

JustHereToLearn 9 months, 1 week ago

How that answer is correct? The first two are okay but add does not make sense it should be "aad" see reference
https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#restrict-access-to-entire-application
Also not sure what the last answer is but it is surely not referrer as there no such thing I could find in the docs.

upvoted 2 times

alejary 9 months ago

I think this should look like this (.referrer will be replaced with the URL of original request):

"responseOverrides": {
"401": {
"statusCode": "302",
"redirect": "/.auth/login/aad?post_login_redirect_uri=.referrer"
}
}

upvoted 22 times

warchoon 8 months, 2 weeks ago

looks like after scan typo

upvoted 3 times

p2006 Most Recent 1 month ago

Got on 9/25/2023
responseoverrides
401
aad
.referrer

upvoted 2 times

RuffBoii 1 month, 1 week ago

Had this on my exam today. Went with the given answer.

upvoted 1 times

mihailos 2 months ago

Got it in exam 28/08/23. Went with proposed answer. Scored 912

upvoted 2 times

ReyPirata 2 months, 1 week ago

This was on the exam (08/20/2023). Went with proposed. Scored 925

upvoted 2 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 4 times

deathRac3 5 months ago

Question was there for me on 29th May 2023

upvoted 4 times

narenazure 6 months, 2 weeks ago

https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#restrict-access-to-entire-application
"responseOverrides": {
"401": {
"statusCode": 302,
"redirect": "/.auth/login/aad"
}

upvoted 4 times

ucskips 7 months, 4 weeks ago

add is meant to be aad. For Azure Active Directory

Response overrides make a custom response instead of error codes 400, 401, 403, 404

upvoted 3 times

Asten 8 months, 3 weeks ago

Given answer is correct. The last selection should directly connected to the third one, not after 302.
https://learn.microsoft.com/en-us/azure/static-web-apps/authentication-authorization

upvoted 2 times

ucskips 7 months, 4 weeks ago

I think 302 is correct there. 302 means after redirect found

upvoted 1 times

Esward 9 months, 1 week ago

it seems given answers correct!!

upvoted 1 times

Esward 9 months, 1 week ago

https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#response-overrides

upvoted 1 times

hugodagg 9 months, 3 weeks ago

Correct answer

upvoted 1 times

Question #34 *Topic 1*

HOTSPOT
-

You are authoring a set of nested Azure Resource Manager templates to deploy Azure resources. You author an Azure Resource Manager template named mainTemplate.json that contains the following linked templates: linkedTemplate1.json, linkedTemplate2.json.

You add parameters to a parameters template file named mainTemplate.parameters,json. You save all templates on a local device in the C:\templates\ folder.

You have the following requirements:

- Store the templates in Azure for later deployment.
- Enable versioning of the templates.
- Manage access to the templates by using Azure RBAC.
- Ensure that users have read-only access to the templates.
- Allow users to deploy the templates.

You need to store the templates in Azure.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
az [ts create / storage account create / storage account update / blueprint artifact template create] \
  --name templateStore \
   --version "1.0" \
   --resource-group templatesRG \
   --location "eastus" \

  --template-file [mainTemplate.json / linkedTemplate1.json / linkedTemplate2.json / mainTemplate.parameters.json] "\
   "C:\templates\
    --tags Dept=HumanResources Environment=Production
```

Correct Answer:

Answer Area

```
az [ts create ▼] \
    --name templateStore \
    --version "1.0" \
    --resource-group templatesRG \
    --location "eastus" \

    --template-file [mainTemplate.json ▼] "\
        "C:\templates\
         --tags Dept=HumanResources Environment=Production
```

ts create
storage account create
storage account update
blueprint artifact template create

mainTemplate.json
linkedTemplate1.json
linkedTemplate2.json
mainTemplate.parameters.json

narenazure Highly Voted 9 months, 3 weeks ago

correct
https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/template-specs-create-linked?tabs=azure-cli

upvoted 14 times

Esward 9 months, 1 week ago

I agree

upvoted 1 times

Dotar 9 months, 2 weeks ago

I agree the second option is because it needs the path to main template no to the parameters.

upvoted 2 times

ReyPirata Most Recent 2 months, 1 week ago

This was on the exam (08/20/2023). Went with proposed. Scored 925

upvoted 4 times

ehrenbe7 3 months, 1 week ago

Why not use az storage account create?
We need to create a storage account first before mentioning the name of the template file..

upvoted 1 times

macobuzi 2 months, 2 weeks ago

Because the code is not about creating a storage account. Pay attention to the part: --name TemplateStore

upvoted 2 times

viking2 3 months, 4 weeks ago

Correct.
The parameters given to the command match those of the az ts command, but not those of az blueprint, so the "ts create" and the main template file are correct answers.

upvoted 2 times

nedlo 7 months, 2 weeks ago

No, the az ts command is not for managing Azure Blueprints.

The az ts command is used for managing Azure Time Series Insights, which is a service that allows you to store, visualize, and query large amounts of time-series data in the cloud. With az ts command, you can create and manage Azure Time Series Insights environments, as well as interact with the data stored in those environments.

Azure Blueprints, on the other hand, is a service that allows you to define a repeatable set of Azure resources that adhere to organizational standards, patterns, and requirements. With Azure Blueprints, you can create a blueprint that specifies the resources and configurations needed to meet your organization's standards and deploy it to multiple subscriptions and environments.

To manage Azure Blueprints, you can use the Azure portal, Azure PowerShell, Azure CLI, or Azure REST API. The Azure CLI command for managing Azure Blueprints is az blueprint, which allows you to create, manage, and assign blueprints to Azure environments.

upvoted 1 times

SprintingLlama 7 months, 1 week ago

This Microsoft article contradicts your point
https://learn.microsoft.com/en-us/cli/azure/ts?view=azure-cli-latest

'az ts' allows you to manage template specs at subscription or resource group scope.

upvoted 4 times

rgullini 7 months ago

This whole text is wrong. Please, ignore this answer

upvoted 13 times

Mattt 7 months ago

It looks like the wrong generated answer by ChatGPT

upvoted 1 times

ucskips 7 months, 4 weeks ago

Just a question. Could account storage have been used instead?

upvoted 1 times

chettir01 9 months, 3 weeks ago

correct
https://learn.microsoft.com/en-us/cli/azure/ts?view=azure-cli-latest

upvoted 1 times

Question #35 *Topic 1*

HOTSPOT
-

You are developing a service where customers can report news events from a browser using Azure Web PubSub. The service is implemented as an Azure Function App that uses the JSON WebSocket subprotocol to receive news events.

You need to implement the bindings for the Azure Function App.

How should you configure the binding? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
{
  "bindings": [
      {
      "type": "                ",
      "direction": "in",
      "name": "data",
      "eventName": "message",
      "eventType": "                "
      }
  ]
}
```

"type" dropdown: user, system, message, connected, webPubSubTrigger, webPubSubConnection

"eventType" dropdown: user, system, message, connected, webPubSubTrigger, webPubSubConnection

Correct Answer:

Answer Area

```
{
  "bindings": [
      {
      "type": "                ",
      "direction": "in",
      "name": "data",
      "eventName": "message",
      "eventType": "                "
      }
  ]
}
```

"type" dropdown: user, system, message, connected, **webPubSubTrigger**, webPubSubConnection

"eventType" dropdown: **user**, system, message, connected, webPubSubTrigger, webPubSubConnection

surprise0011 Highly Voted 6 months, 2 weeks ago
received 2023-17-04 went with given answer, score 926
upvoted 17 times

adilkhan 6 months, 1 week ago
were all Qs from exam topics?
upvoted 6 times

HafizSalmanMalik Highly Voted 9 months, 2 weeks ago
Given answer is correct.
https://learn.microsoft.com/en-us/azure/azure-web-pubsub/reference-functions-bindings?tabs=javascript
upvoted 8 times

Esward 9 months, 1 week ago
Agreed
upvoted 1 times

mihailos Most Recent 2 months ago
Got it in exam 28/08/23. Went with proposed answer. Scored 912
upvoted 3 times

xmashoodx 2 months, 3 weeks ago
what does eventType "user" do ?
upvoted 4 times

Jarynpl 22 hours, 32 minutes ago
EventType can be either user or system. EventType says who triggers the function. The question says: 'where ->customers<- can report news events from a browser using Azure Web PubSub'. For me customer is a person, not a system. So the answer is user.
upvoted 1 times

deathRac3 5 months ago
Question was there for me on 29th May 2023
upvoted 3 times

szumik 8 months, 1 week ago
had this question on exam today 22/02/23, answered wrongly as I saw it for the first time, but seems it's still on the syllabus
upvoted 2 times

adilkhan 6 months, 1 week ago
were all Qs from exam topics?
upvoted 3 times

red0400 8 months ago
pls tell me, is it right answer or not as given?
upvoted 3 times

narenazure 9 months, 3 weeks ago
the answer is correct but is Azure Web PubSub on the syllabus?
upvoted 2 times

AzSimuz 8 months, 2 weeks ago
I wondered the same, I cannot find in the syllabus as of Feb 2023.
upvoted 1 times

pseudocode2023 9 months, 3 weeks ago
Correct
https://learn.microsoft.com/en-us/azure/azure-web-pubsub/reference-functions-bindings?tabs=javascript#example
upvoted 3 times

chettir01 9 months, 3 weeks ago
Crrect
https://learn.microsoft.com/en-us/azure/azure-web-pubsub/reference-functions-bindings?tabs=javascript
upvoted 3 times

Question #36 *Topic 1*

HOTSPOT
-

You are building a software-as-a-service (SaaS) application that analyzes DNA data that will run on Azure virtual machines (VMs) in an availability zone. The data is stored on managed disks attached to the VM. The performance of the analysis is determined by the speed of the disk attached to the VM.

You have the following requirements:

• The application must be able to quickly revert to the previous day's data if a systemic error is detected.
• The application must minimize downtime in the case of an Azure datacenter outage.

You need to provision the managed disk for the VM to maximize performance while meeting the requirements.

Which type of Azure Managed Disk should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Solution
Disk type	Premium SSD / Standard SSD / Standard HDD
Redundancy	Geo-redundant storage (GRS) / Zone-redundant storage (ZRS) / Locally-redundant storage (LRS)

Correct Answer:

Answer Area

Requirement	Solution
Disk type	**Premium SSD** / Standard SSD / Standard HDD
Redundancy	**Geo-redundant storage (GRS)** / Zone-redundant storage (ZRS) / Locally-redundant storage (LRS)

armpro Highly Voted 9 months, 3 weeks ago

Wrong answer: Managed disk doesn't support GRS

Ans: Premium SSD and ZRS
They are asking for high performance workloads which is supported by Premium tier https://learn.microsoft.com/en-us/azure/virtual-machines/disks-types
Also they are asking for zone redundancy (if datacenter goes down, NOT region outage). Also managed disk doesn't support GRS
https://learn.microsoft.com/en-us/azure/virtual-machines/disks-redundancy

upvoted 51 times

zreaf29 1 month, 1 week ago

I agree, too. Azure managed disk doesn't have GRS redundancy option. GRS is a option of azure storage. But this question request use azure managed disk.

upvoted 1 times

Esward 9 months, 1 week ago

I greed. Premium SSD, ZRS are correct answers

upvoted 5 times

armpro 9 months, 3 weeks ago

Premium SSD support backup. Premium SSD v2 doesn't support backup

upvoted 2 times

narenazure Highly Voted 9 months, 3 weeks ago

Azure managed disks offer two storage redundancy options, zone-redundant storage (ZRS), and locally-redundant storage.
Ans: Premium SSD and ZRS
https://learn.microsoft.com/en-us/azure/virtual-machines/disks-redundancy

upvoted 7 times

heatfan900 Most Recent 1 month ago

WHEN THEY SAY 'Azure datacenter outage' they are implying a data center within a region which is hosted in one of the region's three zones. That alone gives it away that its ZRS with P-SSD.

upvoted 2 times

RaoMOwais 7 months ago

Answer by ChatGpt :
Premium Managed Disk
Zone Redundant Storage (ZRS) replication.

upvoted 1 times

kay000001 3 months, 2 weeks ago

Or you could refer to MS docs and learn something?

upvoted 3 times

Kluk_Kluk 7 months, 3 weeks ago

Premium SSD and Zone-Redundant storage ZRS
Because Azure managed disks offer two storage redundancy options, zone-redundant storage and locally-redundan storage. ZRS provider higher availability for managed disks than LRS.

upvoted 1 times

uffuchsi 8 months, 1 week ago

Maybe I am getting confused by the wording here, but doesn't 'Azure datacenter outage' mean the region is out, and thus you need GRS?

upvoted 1 times

Defa 7 months, 4 weeks ago

From: https://learn.microsoft.com/en-us/azure/architecture/resiliency/recovery-loss-azure-region
Azure is divided physically and logically into units called regions. A region consists of one or more data centers in close proximity. Many regions and services also support availability zones, which can be used to provide more resiliency against outages in a single data center. Consider using regions with availability zones to improve the availability of your solution.
***
So, zones for outage in a datacenter.

upvoted 1 times

guschess 9 months ago

Premium and zone. Azure managed disks offer two storage redundancy options, zone-redundant storage (ZRS), and locally-redundant storage. No geo

upvoted 1 times

adilkhan 9 months, 1 week ago

must be ZRS

upvoted 2 times

Question #37 *Topic 1*

HOTSPOT
-

You are developing an application that includes two Docker containers.

The application must meet the following requirements:

- The containers must not run as root.
- The containers must be deployed to Azure Container Instances by using a YAML file.
- The containers must share a lifecycle, resources, local network, and storage volume.
- The storage volume must persist through container crashes.
- The storage volume must be deployed on stop or restart of the containers.

You need to configure Azure Container Instances for the application.

Which configuration values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Configuration setting	Configuration value
Shared lifecycle	Container group / Container image / Service endpoint / Resource group
Storage volume	Azure file share / Secret / Empty directory / Cloned Git repo

Correct Answer:

Answer Area

Configuration setting	Configuration value
Shared lifecycle	**Container group** / Container image / Service endpoint / Resource group
Storage volume	Azure file share / Secret / Empty directory / **Cloned Git repo**

armpro Highly Voted 9 months, 3 weeks ago

Ans: Container group, EmptyDir

Container group is the only logical answer that can have shared lifecycle https://learn.microsoft.com/en-us/azure/container-instances/container-instances-container-groups?source=recommendations#what-is-a-container-group
Azure files need root permission
Secret is for secrets and read-only
EmtyDir can persist through crash and redeployed on stop and restart

https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-emptydir#emptydir-volume
Cloned Git Repo also does the job but it needs more details like Git URL and stuff which are not mentioned to be available in the question

upvoted 25 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score:904/1000

upvoted 6 times

narenazure 9 months, 3 weeks ago

Data in an emptyDir volume is persisted through container crashes. Containers that are restarted, however, are not guaranteed to persist the data in an emptyDir volume. If you stop a container group, the emptyDir volume is not persisted.

upvoted 4 times

alekomar 8 months, 4 weeks ago

There is no requirement to persist the data in a volume between restarts

upvoted 2 times

Esward 9 months, 1 week ago

Gievn answers are correct!!

upvoted 1 times

dddddd111 Most Recent 1 week, 1 day ago

Ans: Container group and Azure File share.

upvoted 1 times

nayru 1 month ago

Got on 26/09/23 and chose container group and empty dir (score 910)

upvoted 3 times

katrang 1 month, 2 weeks ago

Shared Lifecycle - Container Group
Storage volume - The storage volume must persist through container crashes, so emptyDir meets the requirements

upvoted 1 times

JH81 4 months ago

Got this on 6/28/2023 and passed with 850. Went with Container group, EmptyDir.

upvoted 3 times

aragones 5 months, 2 weeks ago

Got this 2023-05-12.

my cases also:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 4 times

GuruPratap 7 months, 1 week ago

Container Group, Azure fileshare
For Shared lifecycle, the appropriate configuration setting would be "Container Group," which allows multiple containers to be deployed together as a group and share a local network, storage, and lifecycle.
For Storage volume, the appropriate configuration setting would be "Azure file share," which provides a durable and highly available storage solution that can be mounted as a volume in the container instance. This ensures that the storage volume persists through container crashes and can be deployed on stop or restart of the containers.

upvoted 3 times

rgullini 6 months, 2 weeks ago

Azure File Share requires root. It is EmptyDir

upvoted 6 times

dddddd111 1 week, 1 day ago

No, using Azure File Share does not require running containers as root. Azure File Share allows you to mount a shared file system to containers in Azure Container Instances without the need for elevated privileges. You can mount Azure File Share to your containers securely without running them as root, ensuring better security practices.

upvoted 1 times

GuruPratap 7 months, 1 week ago

Chat GPT : Container Group, Azure fileshare

upvoted 4 times

rgullini 6 months, 4 weeks ago

Cannot be Azure File Share as requirement is not to run as root.

https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-azure-files

upvoted 4 times

varinder82 7 months, 1 week ago

Could anyone give the correct answare. As it is quite confusing from the comments and multiple people having diff. options

upvoted 4 times

sachinrikhe 7 months, 2 weeks ago

Container Group is the only logical
But storage volume should be azure file share
"Azure File Share allows the containers to share a persistent volume for storing data, and it can be mounted to the containers as a file system. It also provides a durable and highly available storage solution that persists through container crashes and can be deployed on stop or restart of the containers."
Cloned Git does not provide a mechanism for deploying the application or managing its dependencies.

upvoted 3 times

rgullini 6 months, 2 weeks ago

Azure File Share requires root. It is EmptyDir

upvoted 2 times

dddddd111 1 week, 1 day ago

No, using Azure File Share does not require running containers as root. Azure File Share allows you to mount a shared file system to containers in Azure Container Instances without the need for elevated privileges. You can mount Azure File Share to your containers securely without running them as root, ensuring better security practices.

upvoted 1 times

Priya0703 8 months, 1 week ago

This question was in exam today 20-2-2023

upvoted 3 times

dimsok 8 months, 3 weeks ago

"The storage volume must persist through container crashes" +
"The storage volume must be deployed on stop or restart of the containers." = "Cloned Git Repo"

upvoted 1 times

guschess 9 months ago

Ans: Container group, EmptyDir

upvoted 3 times

narenazure 9 months, 3 weeks ago

https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-azure-files here it is clearly mentioned Azure file share volume mount requires the Linux container run as root . so Azure file share is not an option

upvoted 2 times

Ar_priya 9 months, 3 weeks ago

as the container instances must run as non root user, answer is definitely not file share.

"Azure file share volume mount requires the Linux container run as root ."
https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-azure-files#limitations

upvoted 2 times

Ar_priya 9 months, 3 weeks ago

its not directory either. if container stops, directory data is lost and cant be persistent. so it is gitRepo Volume
https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-gitrepo#gitrepo-volume

upvoted 2 times

NicolasMous 1 month, 3 weeks ago

The link is not outdated:
The link you provided is outdated and does not reflect the current state of Azure Container Instances. According to the latest documentation, Azure file share volume mounts do not require the Linux container to run as root anymore. You can use the runAsUser and runAsGroup properties in the container group YAML file to specify a non-root user and group for the container. This way, you can meet the requirement of not running the container as root while using Azure file share as a persistent storage volume. I hope this clarifies your doubt. 😊

: https://docs.microsoft.com/en-us/azure/container-instances/container-instances-volume-azure-files#mount-an-azure-file-share-as-a-volume : https://docs.microsoft.com/en-us/azure/container-instances/container-instances-container-groups#containeridentity

upvoted 1 times

rgullini 6 months, 2 weeks ago

We don´t need the data to remain on restart. The containers are re-deployed.

upvoted 2 times

imanonion 9 months, 3 weeks ago

Why isn't this Azure file share? Clone git repo is only available for Linux machines and this question did not specify Linux as a requirement..

upvoted 2 times

imanonion 9 months, 3 weeks ago

oops my bad..seems like all the four options are only available for Linux lol.
Check this link: https://learn.microsoft.com/en-us/azure/container-instances/container-instances-overview#linux-and-windows-containers

Nonetheless, the link also mentions "To retrieve and persist state with Azure Container Instances, we offer direct mounting of Azure Files shares backed by Azure Storage."..shouldn't Azure Files then be the best answer?

upvoted 2 times

tcybu 9 months, 3 weeks ago

I agree with you, it should be Azure File Share, see: https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-azure-files

upvoted 1 times

JustHereToLearn 9 months, 1 week ago

It cannot be Azure File Share as it contradicts with this "Azure file share volume mount requires the Linux container run as root ." which is explicitly mentioned in question not to do.

upvoted 4 times

tcybu 8 months, 3 weeks ago

You're right, I missed the non-root requirement in the description.
So, it can't be the Azure File Share

upvoted 2 times

Topic 2 - Question Set 2

Question #1 *Topic 2*

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:

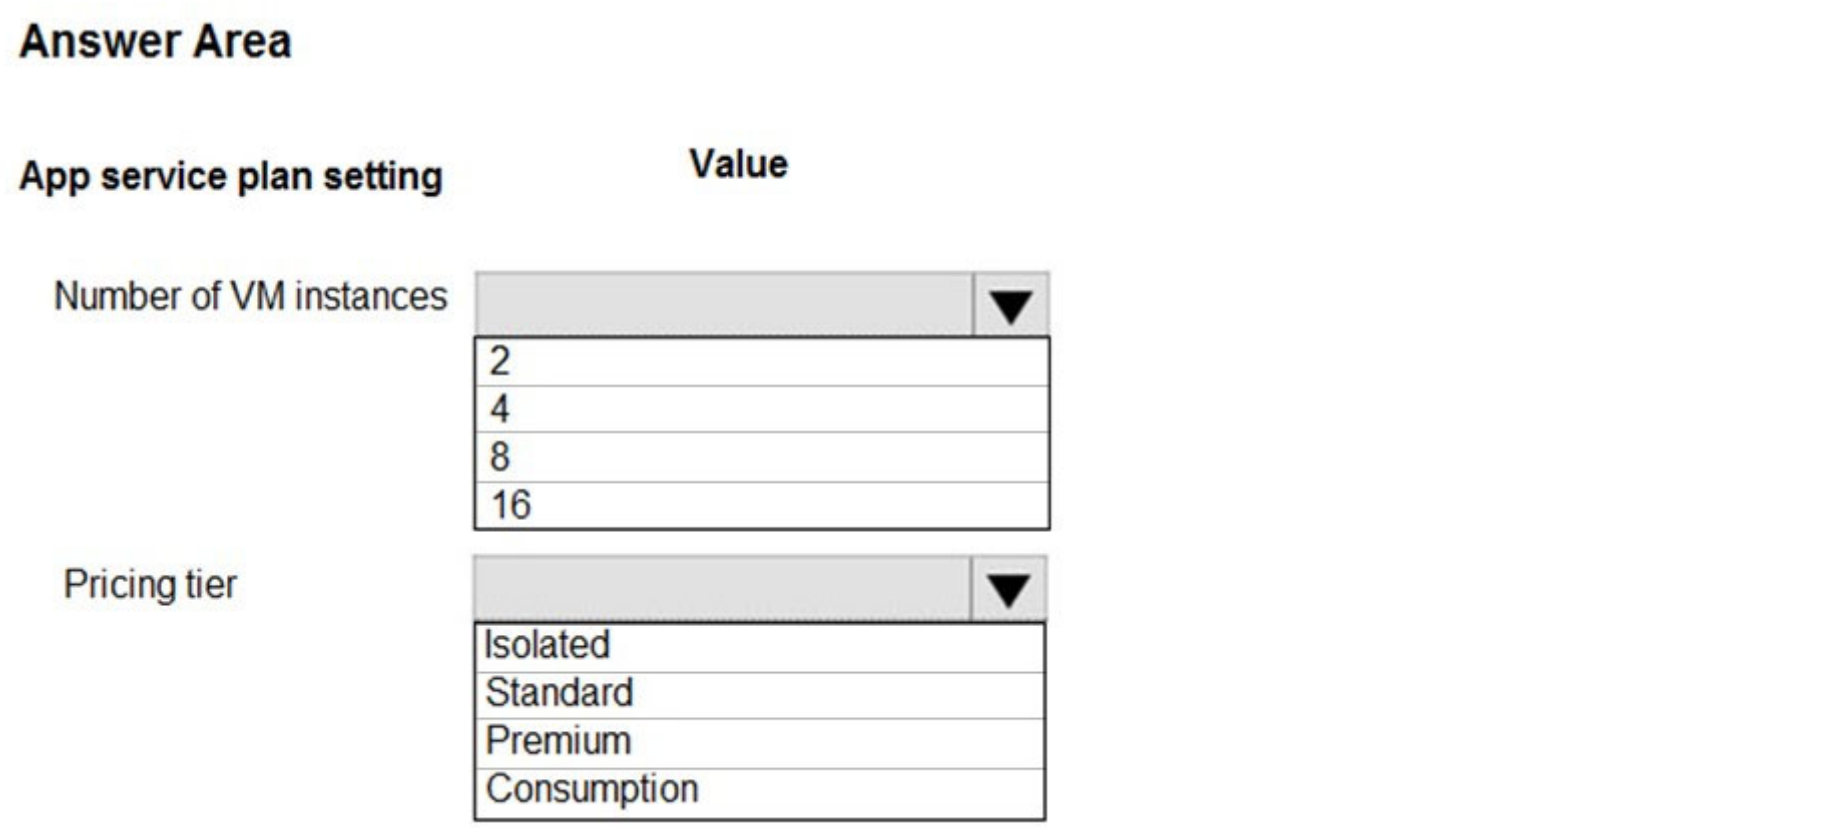

Correct Answer:

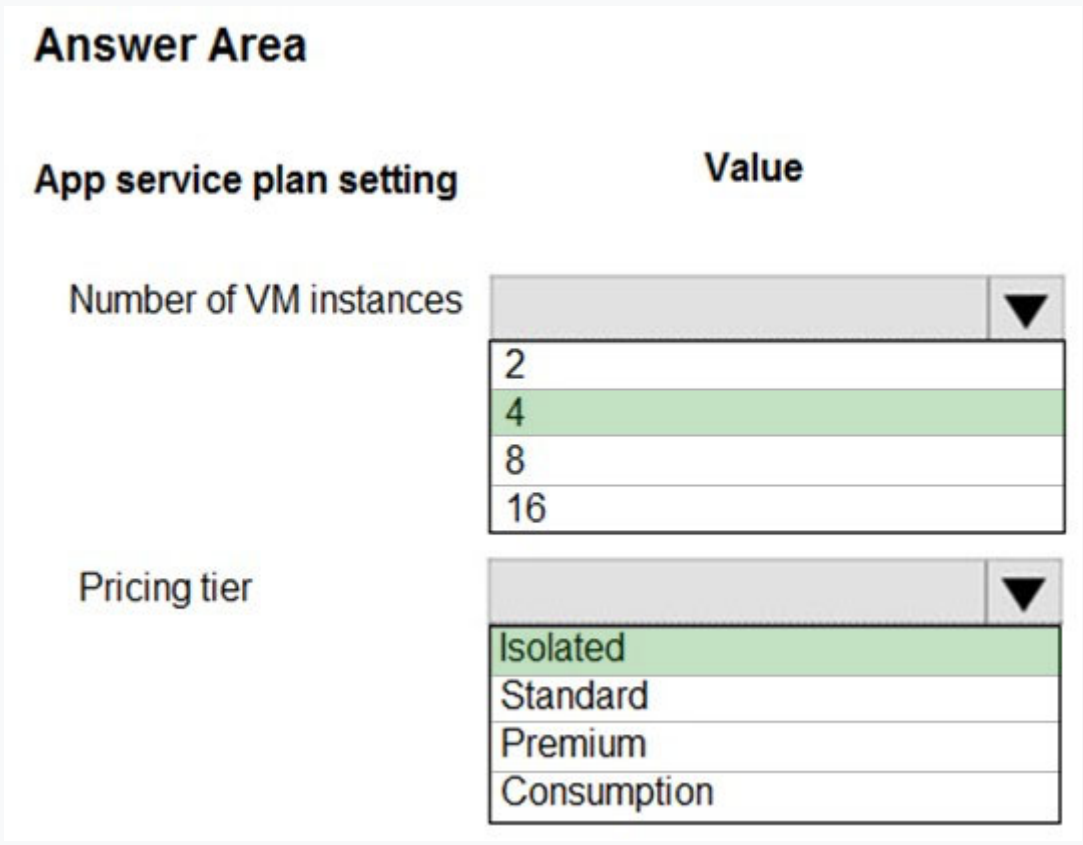

Number of VM instances: 4 -
You are not charged extra for deployment slots.

Pricing tier: Isolated -
The App Service Environment (ASE) is a powerful feature offering of the Azure App Service that gives network isolation and improved scale

capabilities. It is essentially a deployment of the Azure App Service into a subnet of a customer's Azure Virtual Network (VNet).

Reference:

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

mlantonis Highly Voted 2 years, 5 months ago

Box 1: 4
There are four customers that use this service, and each instance of the WebJob processes data for a single customer and must run as a singleton instance. So, the number of VM should be 4. WebJobs is a feature of Azure App Service that enables you to run a program or script in the same instance as a web app. Like running background tasks.

Box 2: Isolated
Azure resources must be located in an isolated network .
In the Isolated tier, the App Service Environment defines the number of isolated workers that run your apps, and each worker is charged. In addition, there's a flat Stamp Fee for the running the App Service Environment itself. Isolated: This tier runs dedicated Azure VMs on dedicated Azure Virtual Networks. It provides network isolation on top of compute isolation to your apps. It provides the maximum scale-out capabilities.

upvoted 108 times

goudigubba Highly Voted 2 years, 8 months ago

Got this in the exam yesterday.

upvoted 19 times

dhosth 2 years, 7 months ago

What is the correct answer?

upvoted 1 times

joanbdm 2 years, 5 months ago

4 & isolated

upvoted 6 times

Soumayard 2 years, 5 months ago

Heyy, did you get a lot of questions like the ones here?

upvoted 2 times

juanckar Most Recent 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 6 times

Devendra_Shukla 3 months, 3 weeks ago

How many questions are from Exam topics?

upvoted 1 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

drgn7676 1 year, 3 months ago

Lorem Ipsum is simply dummy text of the printing and typesetting industry.

upvoted 2 times

BharatVasi 1 year, 4 months ago

Wrote exam yesterday, not even 60% question came from here :-(

upvoted 2 times

adilkhan 9 months, 1 week ago

did you purchase the examtopics account and gone through all 33 pages? please reply

upvoted 3 times

Devendra_Shukla 3 months, 3 weeks ago

Adilkhan how many questions from ET?

upvoted 1 times

Eltooth 1 year, 5 months ago

Going with 4 VM's and Isolated.

upvoted 1 times

veenet 1 year, 7 months ago

the resource must be allocated as Isolate

upvoted 1 times

yifado5713 1 year, 10 months ago

Many questions are not from the dumps so prepare well before appearing for the exam. I passed with 720 score.

upvoted 5 times

sari67 1 year, 10 months ago

I passed the exam today with score 882.All the case studies and questions were the same with these samples in examtopics, except for two or three questions and one case study related to logic app. Good luck!

upvoted 10 times

john4p 1 year, 10 months ago

Nice, thanks for the info.

upvoted 1 times

sangmin214 1 year, 11 months ago

4/Isolated

upvoted 1 times

rohansingh04121985 2 years, 1 month ago

I have cleared the exam, if anyone need paid pdf, then send me an email on azuredeveloper007@gmail.com

upvoted 10 times

prabhjot 2 years, 5 months ago

yes the ans is correct

upvoted 2 times

mlantonis 2 years, 5 months ago

Box 1: 4
There are four customers that use this service, and each instance of the WebJob processes data for a single customer and must run as a singleton instance. So, the number of VM should be 4.

Box 2: Isolated
Azure resources must be located in an isolated network .
In the Isolated tier, the App Service Environment defines the number of isolated workers that run your apps, and each worker is charged. In addition, there's a flat Stamp Fee for the running the App Service Environment itself. Isolated: This tier runs dedicated Azure VMs on dedicated Azure Virtual Networks. It provides network isolation on top of compute isolation to your apps. It provides the maximum scale-out capabilities.

Reference:

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

upvoted 7 times

Spooky7 2 years, 5 months ago

I really have an issue with this question. Why there have to be 4 VMs? Single Azure Plan can be used by multiple applications. They simply reuse existing VMs available within the plan. If the plan have for example 4 VMs then EACH application will use those 4 VMs. WebJob run as a Singleton will run on 1 VM max, but it doesn't mean that each WebJob will need separate VM, right? All 4 may reuse single VM or I understands sth wrong?

upvoted 5 times

basquiat 2 months, 3 weeks ago

I was with the same doubt but think I got it now. The main thing that would justify the 4 VMs in my opinion would be that there are 4 different consumers. Thus, how each one interacts with the VMs might change.

upvoted 1 times

Molte 1 year, 10 months ago

Yeah I thought the same....has nobody an answer for this?

upvoted 2 times

dookebroyle 1 year, 4 months ago

My guess, don't take my word for it, is that it's a Singleton class and only one instance is allowed to be created. After the one instance is created, the instance creation is locked to a single thread. I think the other apps would be unable to create an instance of the WebJob if they are on the same infrastructure.

upvoted 2 times

[Removed] 1 year, 1 month ago

That purely depends on the language and the singleton implementation. Microsoft is a company that does not think through the questions to check if they make sense. Their world revolve around what they know

upvoted 1 times

glam 2 years, 5 months ago

4
Isolated

upvoted 1 times

Bakhodir 2 years, 6 months ago

Each deployment must be tested by using deployment slots prior to serving production data. 8 instances 2 (staging and production) for each customer. How it will be tested in case of 4 instances?

upvoted 2 times

jms309 2 years, 5 months ago

You can have several deployment slots within the same instance so you don't need 8 instances for it but 4. The number of slots you can have will depend on the instance tier you are using. You can check out the next link in order to see it by your own

https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits#app-service-limits

upvoted 3 times

Question #2 *Topic 2*

DRAG DROP -

You are a developer for a software as a service (SaaS) company that uses an Azure Function to process orders. The Azure Function currently runs on an Azure

Function app that is triggered by an Azure Storage queue.

You are preparing to migrate the Azure Function to Kubernetes using Kubernetes-based Event Driven Autoscaling (KEDA).

You need to configure Kubernetes Custom Resource Definitions (CRD) for the Azure Function.

Which CRDs should you configure? To answer, drag the appropriate CRD types to the correct locations. Each CRD type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

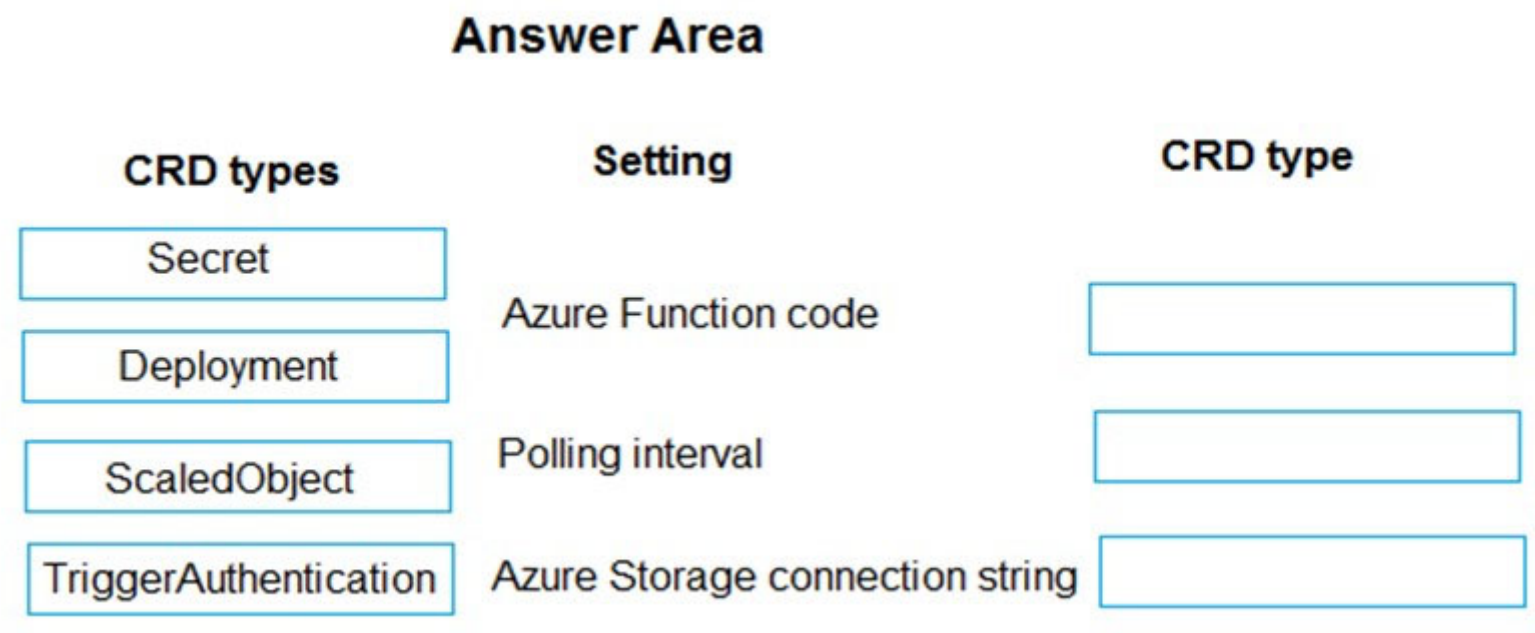

Correct Answer:

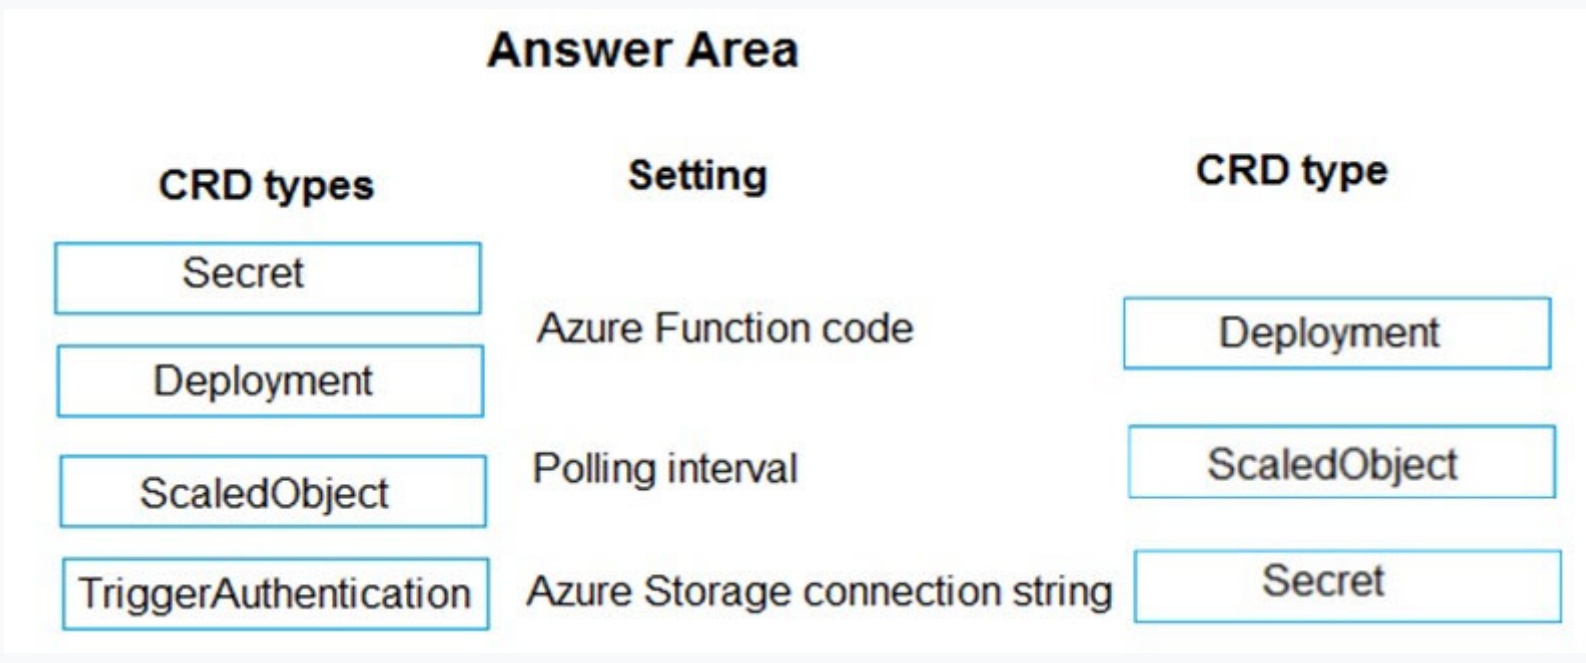

Box 1: Deployment -

To deploy Azure Functions to Kubernetes use the func kubernetes deploy command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.

Box 2: ScaledObject -

With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.

Example of ScaledObject with polling interval

```
apiVersion: keda.k8s.io/v1alpha1
kind: ScaledObject
metadata:
name: transformer-fn
namespace: tt
labels:
deploymentName: transformer-fn
spec:
scaleTargetRef:
deploymentName: transformer-fn
pollingInterval: 5
minReplicaCount: 0
```

```
maxReplicaCount: 100

Box 3: Secret -
Store connection strings in Kubernetes Secrets.
Example: to create the Secret in our demo Namespace:
# create the k8s demo namespace
kubectl create namespace tt
# grab connection string from Azure Service Bus
KEDA_SCALER_CONNECTION_STRING=$(az servicebus queue authorization-rule keys list \
-g $RG_NAME \
--namespace-name $SBN_NAME \
--queue-name inbound \
-n keda-scaler \
--query "primaryConnectionString" \
-o tsv)
# create the kubernetes secret
kubectl create secret generic tt-keda-auth \
--from-literal KedaScaler=$KEDA_SCALER_CONNECTION_STRING \
--namespace tt
Reference:
https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/
```

mlantonis Highly Voted 2 years, 5 months ago

Box 1: Deployment
To deploy Azure Functions to Kubernetes use the func kubernetes deploy command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.

Box 2: ScaledObject
With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.

Box 3: Secret
Store connection strings in Kubernetes Secrets.

upvoted 63 times

neerajupadhyay82 Highly Voted 2 years, 9 months ago

kubernetes is marked as out of scope but still there are few questions coming in the exam from this section

upvoted 17 times

jakobste Most Recent 2 months, 1 week ago

I hope this is not on the exam anymore.

upvoted 1 times

wrak 8 months, 2 weeks ago

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

upvoted 3 times

Esward 9 months, 1 week ago

Given answers are correct!!
https://learn.microsoft.com/en-us/azure/azure-functions/functions-kubernetes-keda

upvoted 1 times

[Removed] 1 year, 1 month ago

This question is not AKS, but Functions being run in Kubernetes. So, technically in scope

upvoted 4 times

nuiowmokka 1 year, 2 months ago

Is Kubernetes part of AZ-204, I did not study Kubernetes?

upvoted 3 times

VACAS 1 year, 2 months ago

I dont know but it doesn't appers here https://docs.microsoft.com/en-us/certifications/exams/az-204

upvoted 4 times

dookebroyle 1 year, 4 months ago

Is there any Microsoft documentation on this? I can't find it.

upvoted 1 times

ldiw02 1 year, 4 months ago

https://keda.sh/
upvoted 2 times

Dev666 1 year, 9 months ago
Box 1: Deployment -
Box 2: ScaledObject -
Box 3: Secret
upvoted 2 times

finnishr 2 years, 2 months ago
1. Azure Function code - Deployment
-To deploy Azure Functions to Kubernetes use the func kubernetes deploy command
2. Polling interval - ScaledObject
- This is the interval to check each trigger on. By default KEDA will check each trigger source on every ScaleObject every 30 seconds.
3. Azure Storage connection string - Secret
- Store connection string in Kubernetes secret

Source for ScaledObject: https://keda.sh/docs/1.4/concepts/scaling-deployments/
upvoted 6 times

glam 2 years, 5 months ago
Box 1: Deployment -
Box 2: ScaledObject -
Box 3: Secret -
upvoted 3 times

Kiranvnv 2 years, 11 months ago
https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/serverless-workloads-with-keda/ is the right link
upvoted 5 times

Leandromellor 2 years, 11 months ago
is this question in the exam?
upvoted 4 times

notjon 2 years, 11 months ago
yes it is, got it yesterday when taking the 204
upvoted 15 times

dirtygooback 2 years, 11 months ago
Yes. There are several Kubernetes questions. I found this odd since all the study material from MS states they use the container registry.
upvoted 8 times

Question #3 *Topic 2*

HOTSPOT -

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

Variable name	Value
$gitrepo	https://github.com/Contos/webapp
$webappname	Webapp1103

You need to automatically deploy code from GitHub to the newly created web app.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
az group create --location westeurope --name myResourceGroup
```

[▼] `--name $webappname --resource-group myResourceGroup --sku FREE`

- az webapp
- az appservice plan create
- az webapp deployment
- az group delete

[▼] `--name $webappname --resource-group myResourceGroup`

- az webapp create
- az appservice plan create
- az webapp deployment
- az group delete

[▼]

- --repo-url $gitrepo --branch master --manual-integration
- git clone $gitrepo
- --plan $webappname

[▼] `source config --name $webappname`

- az webapp
- az appservice plan create
- az webapp deployment
- az group delete

`--resource-group myResourceGroup` [▼]

- --repo-url $gitrepo --branch master --manual-integration
- git clone $gitrepo
- --plan $webappname

Correct Answer:

Answer Area

```
az group create --location westeurope --name myResourceGroup
```

[▼] `--name $webappname --resource-group myResourceGroup --sku FREE`

- az webapp
- **az appservice plan create**
- az webapp deployment
- az group delete

[▼] `--name $webappname --resource-group myResourceGroup`

- **az webapp create**
- az appservice plan create
- az webapp deployment
- az group delete

[▼]

- **--repo-url $gitrepo --branch master --manual-integration**
- git clone $gitrepo
- **--plan $webappname**

[▼] `source config --name $webappname`

- az webapp
- az appservice plan create
- **az webapp deployment**
- az group delete

`--resource-group myResourceGroup` [▼]

- **--repo-url $gitrepo --branch master --manual-integration**
- git clone $gitrepo
- --plan $webappname

Box 1: az appservice plan create

The azure group creates command successfully returns JSON result. Now we can use resource group to create a azure app service plan

Box 2: az webapp create -

Create a new web app..

Box 3: --plan $webappname -

..with the serviceplan we created in step 1.

Box 4: az webapp deployment -

Continuous Delivery with GitHub. Example:

az webapp deployment source config --name firstsamplewebsite1 --resource-group websites--repo-url $gitrepo --branch master --git-token $token

Box 5: --repo-url $gitrepo --branch master --manual-integration

Reference:

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

kemtin Highly Voted 2 years, 11 months ago

Given answer is correct, got this on my test yesterday

upvoted 67 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000

upvoted 7 times

Wanni Highly Voted 2 years, 11 months ago

#!/bin/bash

# Replace the following URL with a public GitHub repo URL
gitrepo=https://github.com/Azure-Samples/php-docs-hello-world
webappname=mywebapp$RANDOM

Create a resource group.
az group create --location westeurope --name myResourceGroup

```
# Create an App Service plan in `FREE` tier.
az appservice plan create --name $webappname --resource-group myResourceGroup --sku FREE

# Create a web app.
az webapp create --name $webappname --resource-group myResourceGroup --plan $webappname

# Deploy code from a public GitHub repository.
az webapp deployment source config --name $webappname --resource-group myResourceGroup \
--repo-url $gitrepo --branch master --manual-integration

# Copy the result of the following command into a browser to see the web app.
echo http://$webappname.azurewebsites.net
```

upvoted 27 times

ENGs Most Recent 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

SangeethaRamasamy 1 month, 2 weeks ago

Got this in 14/09/2023 exam. This exam topic questions are helpful to pass the exam. But we need to understand the scenarios clearly so that we can able to answer the question if its twisted slightly.

upvoted 2 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23

upvoted 2 times

Priya0703 8 months, 1 week ago

Got this question on 20-02-2023 exam.

upvoted 1 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 5 times

AbdulMannan 1 year, 1 month ago

Got this question on 30-Sep-2022 exam.
Answer is correct. Passed with 870 score.

upvoted 6 times

coffecold 1 year, 1 month ago

https://learn.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-github

upvoted 3 times

serpevi 1 year, 1 month ago

Got this in 09/22 , went with the most voted answers, score 927.

upvoted 4 times

ischava780225 1 year, 3 months ago

This is comming in the exam,

upvoted 1 times

Baskman 1 year, 7 months ago

Got this in the exam 03/22

upvoted 2 times

Alasmindas 1 year, 8 months ago

Got this in 03/22 , went with the given answers, the answers are right

upvoted 2 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with given answer.

upvoted 5 times

Mev4953 1 year, 9 months ago

Got this in the exam 01/22

upvoted 5 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with given answer.

upvoted 4 times

rick_cschen 1 year, 11 months ago

have this on exam

upvoted 4 times

Question #4 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure
Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Trigger the photo processing from Blob storage events.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
You need to catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.
Note: Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow.
Events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener.
However, the processing must start in less than one minute.
Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

Community vote distribution
B (67%) A (33%)

YahyaSonmez Highly Voted 3 years ago
Hi All,
The answer (B) is correct. Because, the trick is in the "less than one minute" detail.
You can read about "..10-minute delay in processing new blobs.." in "3-Minimizing latency" description.
Microsoft says: ".....Use Event Grid instead of the Blob storage trigger for the following scenarios:"
1-Blob-only storage accounts: Blob-only storage accounts are supported for blob input and output bindings but not for blob triggers.
2-High-scale: High scale can be loosely defined as containers that have more than 100,000 blobs in them or storage accounts that have more than 100 blob updates per second.
3-Minimizing latency: If your function app is on the Consumption plan, there can be up to a ##10-minute delay in processing new blobs## if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled. You can also use an Event Grid trigger with your Blob storage account. For an example, see the Event Grid tutorial.

REFENCE: https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp#event-grid-trigger

I wish you a good day.
upvoted 150 times

Spooky7 2 years, 5 months ago
I've experienced this in one project I was participating in. If the amount of blobs becomes high enough (like 50k+) blob trigger becomes completely unreliable. It may not trigger at all or delay can be greater than 1 hour. It had nothing to do with Consumption Plan. Switching to Event Grid resolved that issue. So blob trigger DOESN'T guarantee processing within 1 minute, so anwer (B) is correct
upvoted 28 times

matejka 2 years, 11 months ago
I believe you are right with the point 3, but there is no information about the function plan. If it is a consumption plan then correct answer is NO, if it is on Premium then YES can be correct too. But I would agree with you that correct answer is probably NO.
upvoted 4 times

coffecold 1 year, 1 month ago
I agree with you, there is nothing about the chosen function plan. I would vote for B because I think Microsoft wants you to use Event grid instead of Microsoft.Azure.WebJobs.BlobTrigger
upvoted 1 times

Juanlu 3 years ago

Totally agree, YahyaSonmez !
upvoted 1 times

Zidimirite 2 years, 6 months ago
You are right, since this is a MS exam they always want to show off the variations that you can do with their platform and them specifying v2 is definitely a tell... But to get around the latency you could also turn on "Always on", which looking at that alone you could say A is right. This is not a great question.
upvoted 4 times

[Removed] Highly Voted 3 years, 4 months ago
Answer is not at all clear from the explanation. It seems like answer should be yes instead of no.
upvoted 26 times

xRiot007 3 years, 4 months ago
The answer is correct. "No". You need an Azure Function for processing.
upvoted 9 times

Camios 3 years, 3 months ago
Except the question says what triggers the processing, not what does the processing. Question/answer is poorly written.
upvoted 35 times

angelsrp 3 years, 3 months ago
you can capture an event without processing the data related to the event. Given answer is correct
upvoted 3 times

[Removed] Most Recent 3 months, 3 weeks ago
Correct answer is A
upvoted 1 times

DonH 4 months, 2 weeks ago
Just as information I just had this question on my AZ204 exam - 16-jun-2023.
I barely made it (with only 767 points) so I can't inform anyony if this answer is correct or not, just stating that this is an actual exam question.
upvoted 3 times

Bababar 5 months, 4 weeks ago
Correct Answer : B
Not necessary to convert the account, instead move photo
processing to an Azure Function triggered from the blob upload..
Azure Storage events allow applications to react to events.
Common Blob storage event scenarios include image or video
processing, search indexing, or any file- oriented workflow.
Note: Only storage accounts of kind StorageV2 (general purpose
v2) and BlobStorage support event integration. Storage (general
purpose v1) does not support integration with Event Grid.
upvoted 1 times

CODE_STS 8 months ago
Got this in the exam today! Feb 28, 2023
upvoted 3 times

fkaracan 8 months, 1 week ago
I think "##must start## in less than one minute" is the keyword here. Thus answer is highly B.
upvoted 1 times

ABHI_SHEK_31 8 months, 3 weeks ago
Yes, the solution of triggering the photo processing from Blob storage events meets the goal of starting the photo processing in less than one minute after a photo is uploaded. Azure Storage Blob storage provides event notifications for blob creation, which can be used to trigger the photo processing. This allows for the processing to start as soon as a photo is uploaded, thereby meeting the requirement of starting the process in less than one minute.
upvoted 1 times

Rosu123 9 months, 2 weeks ago
The question states that you have developed a service. So why wait for the Blob storage to trigger when the service intercepts the event requested before even reaching the blob. So the answer is B.
upvoted 1 times

carlosghosn 11 months, 1 week ago
Got this in the exam today ! Nov 25, 2022
upvoted 2 times

hubekpeter 11 months, 2 weeks ago
Selected Answer: A
Please stop with that 10 minutes non sense. They mentioned event based blob storage which have low latency https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=in-process%2Cextensionv5&pivots=programming-language-python so the answer I A.
upvoted 3 times

[Removed] 1 year, 1 month ago

All well, but Microsoft documentations says "Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions ..." So, the choice implicitly mean Event grid is there. And it does not say anything about a Function app or the plan used. If it did say consumption plan, I could have agreed the answer is B. But it does not. Just another bad question by Microsoft in the guess-whats-running-in-question-creators-mind game.

upvoted 2 times

gmishra88 1 year ago

But also consider that the blob-trigger they say is not the events coming through event grid. So, the answer is correct. This is not a good way to use blob-trigger

upvoted 1 times

Knightie 1 year, 2 months ago

If you want 1 min, you need to start the processing right when a photo is uploaded, so event subscription is correct. It is a Blob Event

Azure Storage events allow applications to react to events, such as the creation and deletion of blobs. It does so without the need for complicated code or expensive and inefficient polling services. The best part is you only pay for what you use.

Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener. Event Grid provides reliable event delivery to your applications through rich retry policies and dead-lettering.

See the Blob storage events schema article to view the full list of the events that Blob storage supports.

Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow. Asynchronous file uploads are a great fit for events. When changes are infrequent, but your scenario requires immediate responsiveness, event-based architecture can be especially efficient.

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

upvoted 2 times

N9 1 year, 2 months ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

dookebroyle 1 year, 4 months ago

Has anyone received this question on the exam and know absolutely sure what the answer is? Seems like there's a 50/50 opinion on this threat, although the people saying the answer is yes have not voted

upvoted 2 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.

upvoted 1 times

TangAnna 1 year, 6 months ago

The answer is YES. Blog storage Event(implemented by Event Grid) can be a good solution.

upvoted 1 times

Question #5 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.
You need to ensure that scripts run and resources are available before a swap operation occurs.
Solution: Update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.
Does the solution meet the goal?

A. No

B. Yes

Correct Answer: *A*
Specify custom warm-up.
Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostName="[app hostname]" />
<add initializationPage="/Home/About" hostName="[app hostname]" />
</applicationInitialization>
</system.webServer>
Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps

Community vote distribution
B (82%) A (18%)

GMartinez Highly Voted 1 year, 5 months ago
Selected Answer: B
This should work:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up
upvoted 23 times

d79 6 months, 2 weeks ago
The anwser might be right, but the question is wrong. Auto-swap needs to be set on the source slot, not the target slot (production in this question)
upvoted 2 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000
upvoted 3 times

Lauri_n Highly Voted 10 months, 3 weeks ago
Selected Answer: A
I think the correct answer really might be A (No). It's not because of the last part about the warm-up being wrong. The incorrect part is before that: "You enable auto swap on the Production deployment slot.". The auto swap is enabled on the source slot, not the destination slot.

https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots#configure-auto-swap
upvoted 6 times

AAcademy 9 months, 1 week ago
The target slot (production) is swapped with the source slot (staging). You have to warm up the staging slot with the new code and then switch it with the production slot. This way, the production slot will used the previously warmed staging slot.
upvoted 1 times

shekhar11 Most Recent 1 month, 1 week ago
On my exam in sept23 2023
upvoted 1 times

Aneeqa 1 month ago

what is the correct ans?

upvoted 1 times

Tarajee 1 month, 2 weeks ago

On my exam in 2023 sept10

upvoted 1 times

longnguyendh 2 months, 4 weeks ago

It should be yes. Keyword is applicationInitialization

upvoted 1 times

deathRac3 5 months, 1 week ago

Selected Answer: B

sequence was wrong

upvoted 1 times

p2006 6 months, 4 weeks ago

Selected Answer: B

https://learn.microsoft.com/en-us/answers/questions/980267/what-do-we-need-to-warmup-before-swapping-apps-ser

upvoted 1 times

Priya0703 8 months, 1 week ago

Got this question series on 20-02-2023 exam.

upvoted 3 times

Poops 10 months ago

Selected Answer: B

Should be B (Yes), that's the way to implement warmup according to MS:

https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up

upvoted 1 times

gds03 1 year, 2 months ago

You need to ensure that scripts run and resources are available before a swap operation occurs.

App Initialization has nothing to do with have sure that resources are available.

upvoted 2 times

rohitpatil113 1 year, 2 months ago

Selected Answer: A

https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#Auto-Swap

upvoted 2 times

d_vaibhav 1 year, 2 months ago

Correct answer must be B (applicationinitialization tag is way of implementing custom warm-up)

upvoted 2 times

Brewer 1 year, 2 months ago

Could be the trick is in "run the scripts". From the docs seems that only trigger one or more pages is possible, but there is nothing related to "running custom scripts".

Reference: https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#Auto-Swap

upvoted 5 times

Satish_Babu 1 year, 2 months ago

Selected Answer: B

correct answer is B

upvoted 1 times

N9 1 year, 2 months ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

Notorious19 1 year, 3 months ago

Selected Answer: B

b(yes) is correct as per the doc

upvoted 1 times

Dani_ac7 1 year, 4 months ago

Selected Answer: B

B is the answer

upvoted 1 times

Question #6 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.
You need to ensure that scripts run and resources are available before a swap operation occurs.
Solution: Enable auto swap for the Testing slot. Deploy the app to the Testing slot.
Does the solution meet the goal?

A. No

B. Yes

Correct Answer: *B*
Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.
Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot.
Here's a sample web.config fragment.
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostName="[app hostname]" />
<add initializationPage="/Home/About" hostName="[app hostname]" />
</applicationInitialization>
</system.webServer>
Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps

Community vote distribution
A (71%) B (29%)

GMartinez Highly Voted 1 year, 5 months ago
Would say "No"
upvoted 15 times

Mousavi 1 year, 4 months ago
I think the given answer is correct
https://www.pragimtech.com/blog/azure/azure-deployment-slots-auto-swap/
upvoted 3 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000
upvoted 6 times

Spartaco Highly Voted 1 year, 5 months ago
Selected Answer: A
I vote A, No, because for me the solution is updating the web.config file to include the applicationInitialization configuration element.
upvoted 8 times

Tarajee Most Recent 1 month, 2 weeks ago
Got this on 2023sept10
upvoted 1 times

Vladimir_Gajinov 1 month, 3 weeks ago
Selected Answer: A
My vote goes to 'No'.

In the "How to configure auto swap" section, step 3 "Navigate to the **source slot**, in our case staging slot." gives us an answer.
https://www.pragimtech.com/blog/azure/azure-deployment-slots-auto-swap/
upvoted 1 times

basquiat 2 months, 3 weeks ago

All the respect for you guys commenting here, but just left: "The correct answer is X" without saying why or providing documentation kinda sucks.
upvoted 4 times

MakMao 2 months, 4 weeks ago
Given answer is correct:
'Before you configure auto swap for the production slot, consider testing auto swap on an nonproduction target slot'
https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots?tabs=portal
upvoted 4 times

VijayNarayanan 2 months, 3 weeks ago
Thanks MakMao. Confusion is cleared.
upvoted 2 times

JH81 4 months ago
Selected Answer: A
Got this on 6/28/2023 and passed with 850. Went with No.
upvoted 2 times

red0400 8 months, 1 week ago
B. Yes

Enabling auto swap for the Testing slot and deploying the app to the Testing slot will ensure that the app is fully tested and resources are available before the swap occurs to the Production slot. This way, any issues that arise during testing can be addressed before the app is made live in Production.
upvoted 4 times

Mike5 8 months, 3 weeks ago
Selected Answer: A
To meet the goal of ensuring that scripts run and resources are available before a swap operation occurs, you should update the web.config file to include the application Initialization configuration element, as described in my previous answer.
upvoted 1 times

mdg3501 10 months, 3 weeks ago
got this on 2022-12-7
upvoted 3 times

rohitpatil113 1 year, 2 months ago
Selected Answer: B
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#Auto-Swap
upvoted 3 times

Dani_ac7 1 year, 4 months ago
Selected Answer: B
Answer is NO
upvoted 1 times

cemepo9819 1 year, 3 months ago
B is 'yes'.
upvoted 6 times

Energo 7 months, 3 weeks ago
switcheroo moment
upvoted 1 times

dookebroyle 1 year, 4 months ago
Selected Answer: B
No, because if you need scripts to run prior then you need to specify a custom warm up
upvoted 4 times

cemepo9819 1 year, 3 months ago
B is 'yes'.
upvoted 5 times

Satish_Babu 1 year, 4 months ago
Selected Answer: A
No,
No,
because correct solution is updating the web.config file to include applicationinitialization configuration element.
upvoted 3 times

omega11 1 year, 4 months ago
Selected Answer: A
I say A, "No".

You cannot ensure that scripts run and resources are available before a swap operation occurs by configuring auto swap. The applicationInitialization configuration in the web.config does.

https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up

upvoted 5 times

Question #7 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.
You need to ensure that scripts run and resources are available before a swap operation occurs.
Solution: Disable auto swap. Update the app with a method named statuscheck to run the scripts. Re-enable auto swap and deploy the app to the Production slot.
Does the solution meet the goal?

A. No

B. Yes

Correct Answer: *B*
Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.
Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot.
Here's a sample web.config fragment.
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostName="[app hostname]" />
<add initializationPage="/Home/About" hostName="[app hostname]" />
</applicationInitialization>
</system.webServer>
Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps

Community vote distribution
A (87%) 13%

GMartinez Highly Voted 1 year, 5 months ago
Selected Answer: A
Instead, use applicationInitialization
upvoted 19 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000
upvoted 5 times

robin1337 Highly Voted 1 year, 3 months ago
It depends.... has the app an WEBSITE_SWAP_WARMUP_PING_PATH set to /statuscheck??
Seriously, I hold multiple certifications, tricky hands-on ones but this certification is a joke. Microsoft, who is designing those questions?!? It's not just this one, there are so many bad questions.
upvoted 17 times

[Removed] 1 year, 1 month ago
Same experience here. I have many AWS certifications, OCP-11 certification, etc. This is really a joke of a certification. I do not even know what Microsoft expects and more importantly the kind of people ready to live with this kind of questions from Microsoft and trying to remember it.
upvoted 8 times

robin1337 1 year, 3 months ago
Ok, my bad autoswap has to be configured on the Staging Slot. So "You enable auto swap on the Production deployment slot." is not what you would do. Answer is A (NO).
upvoted 6 times

Tarajee Most Recent 1 month, 2 weeks ago
Got this on 2023sept
upvoted 1 times

GuruPratap 7 months, 1 week ago

Yes, the solution should meet the goal of ensuring that scripts run and resources are available before a swap operation occurs. By disabling auto swap, the deployment to the Production slot will not occur until it is manually initiated. Then, the app can be updated with a method named statuscheck to run the necessary scripts and ensure that resources are available. Once the app has been updated, auto swap can be re-enabled and the updated app can be deployed to the Production slot. This approach allows for greater control over when the swap occurs and ensures that the necessary steps are taken before the swap happens.

upvoted 5 times

Saronyo 8 months, 3 weeks ago

Selected B , Application initialisation can be set up to do the statuscheck . auto swap will deploy it to production after warm up

upvoted 3 times

red0400 8 months, 1 week ago

Exactly

upvoted 1 times

kociuba 10 months ago

got in on my exam 30-12-2022
score: 818

upvoted 2 times

Congphan 10 months ago

what is your choice?

upvoted 1 times

mdg3501 10 months, 3 weeks ago

got this on 2022-12-7

upvoted 1 times

micro9000 11 months, 3 weeks ago

The answer is A, you still need to set the WEBSITE_WARMUP_PATH in app settings

WEBSITE_WARMUP_PATH: A relative path on the site that should be pinged whenever the site restarts (not only during slot swaps). Example values include /statuscheck or the root path, /.

https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots#Auto-Swap

upvoted 1 times

[Removed] 1 year, 1 month ago

Deploy the app directly to the Production slot will fail the auto swap. But that is a detail in the question that will be missed by a person who did some reading on the auto swap because he is thinking of app initialization option vs custom_warmup with those properties. Do not be evil, Microsoft.

upvoted 1 times

N9 1 year, 2 months ago

Selected Answer: A

web.config file with applicationinitialization is correct answer. So A is the correct answer.

upvoted 2 times

Dani_ac7 1 year, 4 months ago

Selected Answer: B

applicationInitialization

upvoted 1 times

Lucky_me 1 year, 4 months ago

Selected Answer: B

although the provided explaination in the answer section, is not related to the correct option, however, the solution provided in the question would actually work, and that's is the correct "Manual Way" of making sure the deployment slot is warmed up and running before swap, so the solution would work and correct.

upvoted 3 times

Satish_Babu 1 year, 4 months ago

Selected Answer: A

No,
No,
because correct solution is updating the web.config file to include applicationinitialization configuration element for running initializing resources.

upvoted 1 times

cemepo9819 1 year, 3 months ago

multiple solutions can be correct.

upvoted 4 times

omega11 1 year, 4 months ago

I say A, "No".

The reason why setting up auto swap is to warm up the application to prevent downtime, so deploying to the production slot directly is killing the

purpose.

https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up
upvoted 3 times

Eltooth 1 year, 4 months ago

Selected Answer: A

A seems to be correct answer.
upvoted 2 times

muggins 1 year, 4 months ago

Can we update to the correct answer then? Seems like B is the wrong answer, but is the selected answer.
upvoted 3 times

Spartaco 1 year, 5 months ago

Selected Answer: A

Use applicationInitialization
upvoted 2 times

Question #8 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure
Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Convert the Azure Storage account to a BlockBlobStorage storage account.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
Not necessary to convert the account, instead move photo processing to an Azure Function triggered from the blob upload..
Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file- oriented workflow.
Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

Community vote distribution
B (89%) 11%

oleks Highly Voted 3 years, 2 months ago
B is correct
upvoted 31 times

MiraA 2 years ago
The answer is NO - the Even Grid integration should be used.
https://docs.microsoft.com/en-us/azure/event-grid/resize-images-on-storage-blob-upload-event
upvoted 8 times

Justing_Gao Highly Voted 3 years, 3 months ago
A BlockBlobStorage account is a specialized storage account in the premium performance tier for storing unstructured object data as block blobs or append blobs. Compared with general-purpose v2 and BlobStorage accounts, BlockBlobStorage accounts provide low, consistent latency and higher transaction rates.
upvoted 17 times

syj123 3 years, 2 months ago
so, Do you mean that BlockBlobStorage account is better but not necessary?
upvoted 2 times

kvtborad Most Recent 2 months, 3 weeks ago
I have this question on 6th August. it was passed with 904. chose the NO.
upvoted 3 times

Hogwarts 6 months, 4 weeks ago
Got this in exam today (5 April 2023)
upvoted 2 times

CODE_STS 8 months ago
Got this in the exam today! Feb 28, 2023
upvoted 1 times

ABHI_SHEK_31 8 months, 3 weeks ago
No, converting the Azure Storage account to a BlockBlobStorage storage account does not meet the requirement of starting the photo processing in less than one minute. The conversion of the storage account type would not have any impact on the time it takes to start the photo processing. A different solution such as triggering the photo processing from Blob storage events or using a queue-based solution may be more appropriate to meet the requirement.

upvoted 3 times

kociuba 10 months ago

got it on my exam 30-12-2022
score: 818

upvoted 1 times

SoftSol 10 months, 3 weeks ago

Selected Answer: B

Got this in the exam dec 13, 2022

upvoted 1 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

AZAdam22 1 year, 3 months ago

Selected Answer: B

No - The account type does not affect how the photo processing code is triggered.

upvoted 5 times

dookebroyle 1 year, 4 months ago

Selected Answer: A

No, because you can't convert this account. I think the question is intentionally misleading, wanting you to be distracted by the 1 minute start to process time. Note that the requirements state the processing needs to START in less than 1 minute, not finish in a minute.

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.

upvoted 1 times

Goyo 1 year, 8 months ago

Selected Answer: B

The correct answer is NO

upvoted 1 times

dmeld 1 year, 11 months ago

Answer is No. Blob storage trigger cannot handle the reuqirment less than one minutes. You should use the Event Grid trigger. Hence, it is not useful change the account storage tier from v2 to block blob account.

upvoted 2 times

Chked 2 years, 3 months ago

Answer - NO
"You can't convert an existing standard performance storage account to a block blob storage account with premium performance. To migrate to a premium performance storage account, you must create a premium block blob account, and migrate the data to the new account."
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-performance-tiers

upvoted 10 times

kondapaturi 2 years, 4 months ago

B-No is the correct, Here we need to develop the processing logic , and the best way is to create an Azure Function with a blob storage trigger.

upvoted 2 times

az204dev 1 year, 10 months ago

Using Azure Function with blob storage trigger will not satisfy the requirement to process image in less than 1 minute. Instead use event grid.

upvoted 3 times

PierrGo 2 years, 4 months ago

I had this question this morning. I answered "B" (no), I think this is the right answer

upvoted 3 times

Question #9 *Topic 2*

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:

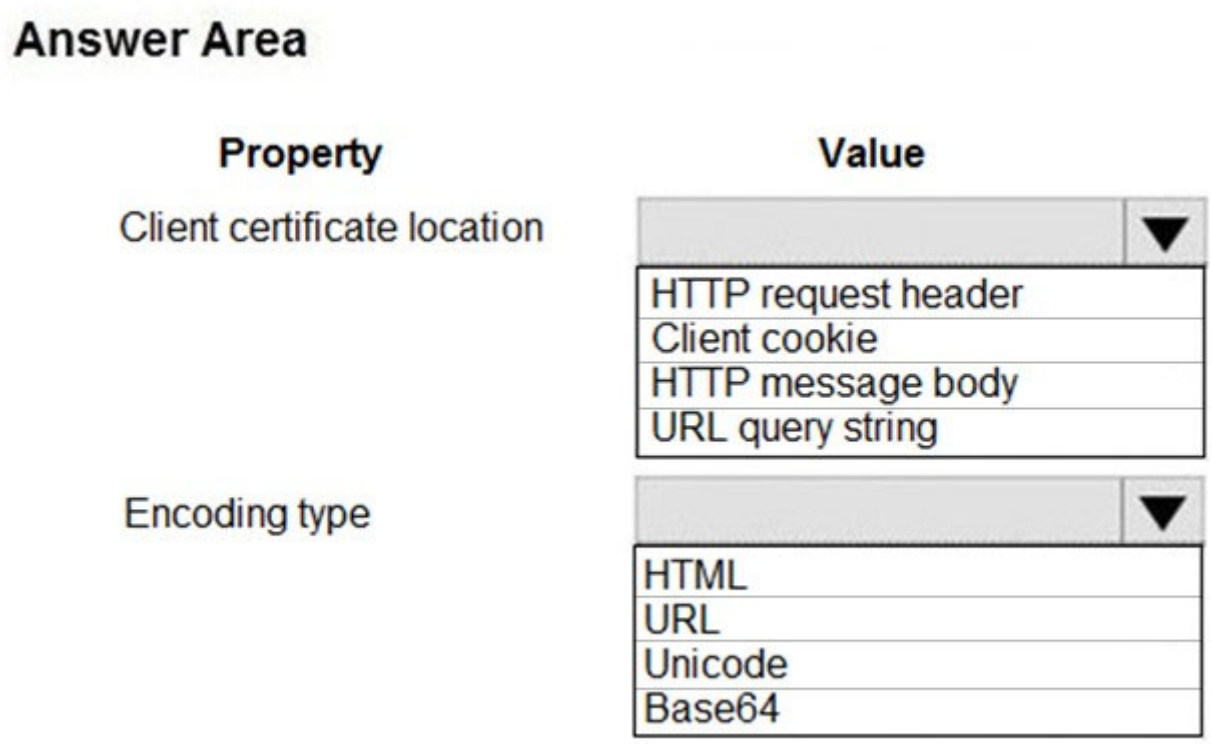

Correct Answer:

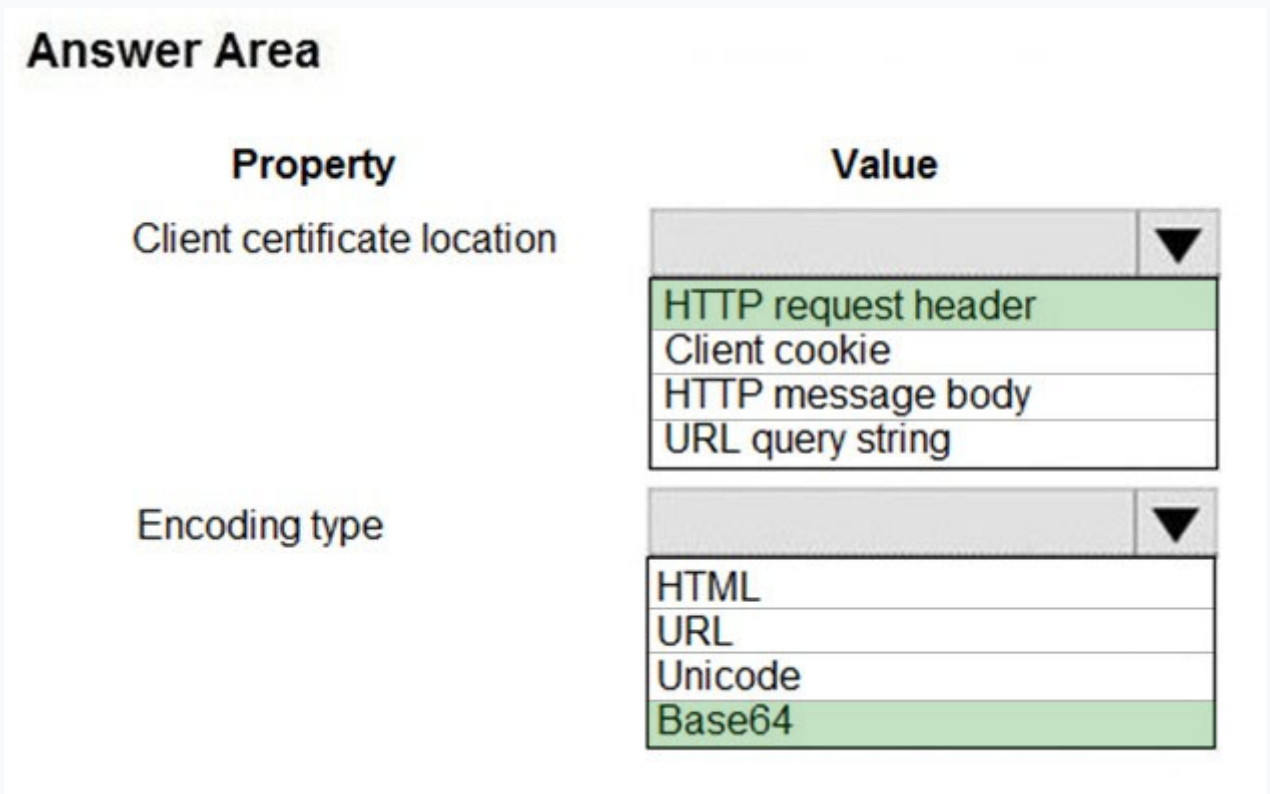

Accessing the client certificate from App Service.

If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be available through the HttpRequest.ClientCertificate property. For other application stacks, the client cert will be available in your app through a base64 encoded value in the "X-ARR-ClientCert" request header. Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.

Reference:

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

mlantonis Highly Voted 2 years, 5 months ago

Box 1: HTTP request header
If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be available through the HttpRequest.ClientCertificate property.

Box 2: Base64
For other application stacks, the client cert will be available in your app through a base64 encoded value in the "X-ARR-ClientCert" request header. Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.

Reference:

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

upvoted 51 times

27close Highly Voted 2 years, 11 months ago
With client certificates enabled, App Service injects an X-ARR-ClientCert request header with the client certificate.
upvoted 16 times

NightshadeRC Most Recent 3 months, 1 week ago
Had this question today: 2023-07-26
upvoted 3 times

Hogwarts 6 months, 4 weeks ago
Got this in exam today (5 April 2023)
upvoted 2 times

ucskips 7 months, 3 weeks ago
why is it base64?
upvoted 3 times

CODE_STS 8 months ago
Got this in the exam today! Feb 28, 2023
upvoted 1 times

uffuchsi 8 months, 1 week ago
Received this in my exam today (22/02/2023). Selected HTTP request header, and Base64. Score 927.
upvoted 3 times

Jyo221 8 months ago
Are the questions coming from dump
upvoted 1 times

Esward 8 months, 2 weeks ago
It is there in 13 Feb 2023 exam
upvoted 2 times

zb1234 10 months ago
Got this Dec 28 2022
upvoted 4 times

nvtienanh 11 months ago
Got this question on 2-Dec-2022 exam.
Answer is correct. Passed with 857 score.
upvoted 6 times

ms_master 12 months ago
Got this question on 30-Oct-2022 exam.
Answer is correct. Passed with 875 score
upvoted 2 times

AbdulMannan 1 year, 1 month ago
Got this question on 30-Sep-2022 exam.
Answer is correct. Passed with 870 score.
upvoted 2 times

serpevi 1 year, 1 month ago
Got this in 09/22 , went with the most voted answers, score 927.
upvoted 2 times

N9 1 year, 2 months ago
Box 1: HTTP request header
Box 2: Base64
upvoted 2 times

HafizSalmanMalik 1 year, 4 months ago
Statement 1: HTTP request header
If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be available through the HttpRequest.ClientCertificate property.

Statement 2: Base64
For other application stacks, the client cert will be available in your app through a base64 encoded value in the "X-ARR-ClientCert" request header. Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.

Reference:

https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth
upvoted 1 times

Eltooth 1 year, 4 months ago

HTTP request header
Base64

upvoted 1 times

AZ204Cert 1 year, 6 months ago

Got it in my exam 04/05/22 (selected HTTP request header, Base64)

upvoted 3 times

Question #10 *Topic 2*

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:

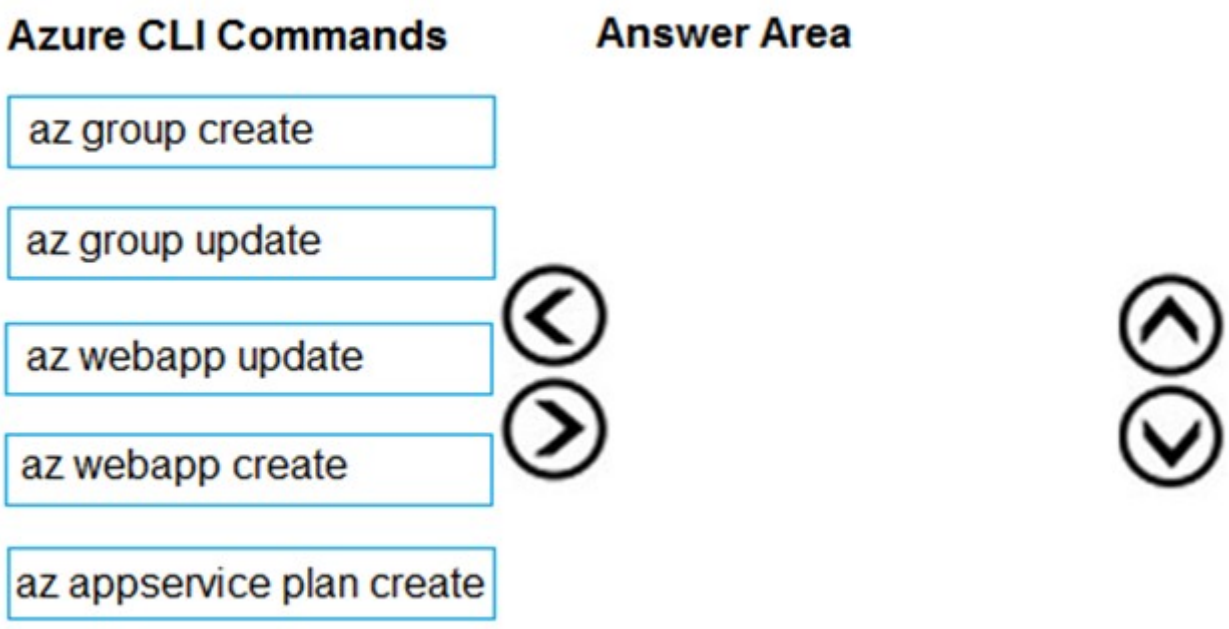

Correct Answer:

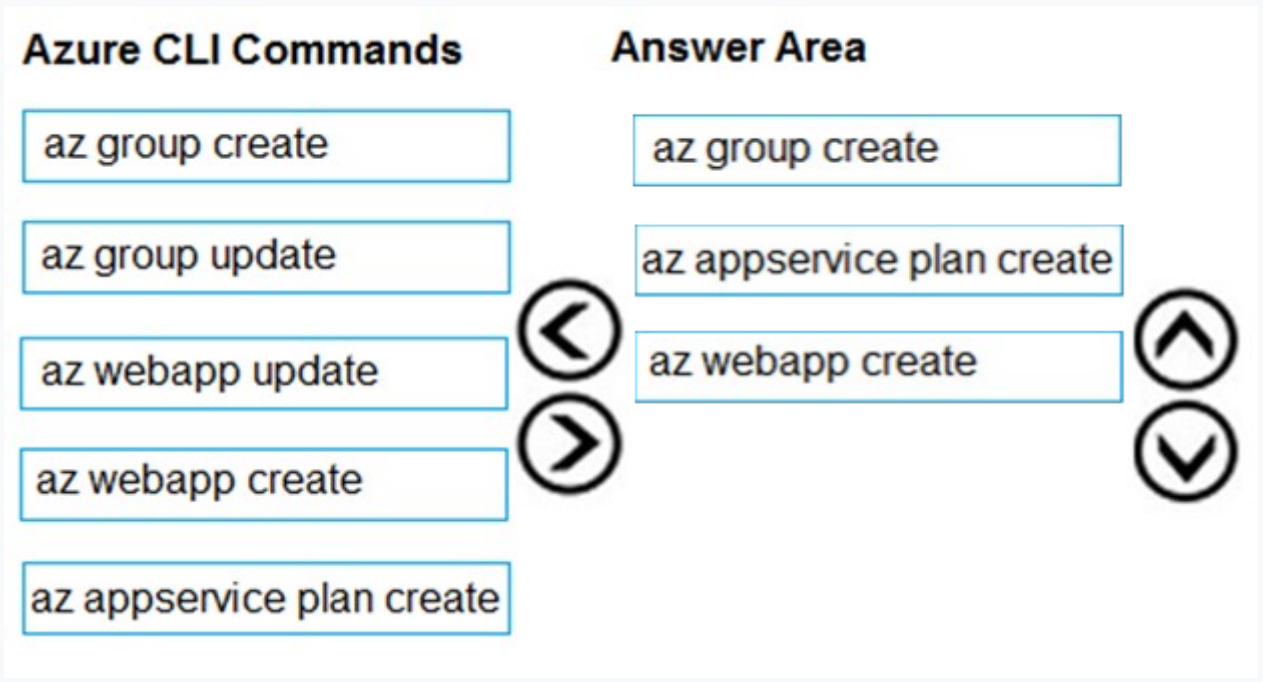

You can host native Linux applications in the cloud by using Azure Web Apps. To create a Web App for Containers, you must run Azure CLI commands that create a group, then a service plan, and finally the web app itself.

Step 1: az group create -

In the Cloud Shell, create a resource group with the az group create command.

Step 2: az appservice plan create

In the Cloud Shell, create an App Service plan in the resource group with the az appservice plan create command.

Step 3: az webapp create -

In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command. Don't forget to replace with a unique app name, and <docker-ID> with your Docker ID.

Reference:

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

Secure01 Highly Voted 2 years, 11 months ago

Within the same resource group, you can't mix Windows and Linux apps in the same region.
https://docs.microsoft.com/en-us/azure/app-service/overview#app-service-on-linux
"None of your current resource groups reside in a location that supports Linux"
So you have to create new resource group. Answer is correct.

upvoted 66 times

rgullini 2 years, 1 month ago

This explanation is correct and it is applicable BEFORE January 21, 2021. Now the article says:

"Linux and Windows App Service plans can now share resource groups. This limitation has been lifted from the platform and existing resource groups have been updated to support this"

upvoted 25 times

kags 2 years ago

Agreed. Just created both linux and windows, core3.1 and .net 5 respectively in same resource group.

upvoted 4 times

Netspud 1 year, 8 months ago

Many are missing the point here, it says "None of your current resource groups reside in a location that supports Linux". So you need to create a group in an area that does support Linux. Doesn't matter if you mix them or not, there is currently nothing that supports Linux, so something new is required.

For me the answer is correct.

upvoted 33 times

bubbleGumStroganoff 1 year, 7 months ago

The resources inside of a resource group can be in a region that differs from the region of the resource group they reside in though no?

upvoted 5 times

rohitpatil113 1 year, 2 months ago

Limitations
Note

Linux and Windows App Service plans can now share resource groups. This limitation has been lifted from the platform and existing resource groups have been updated to support this.

upvoted 13 times

JulienYork 2 years, 9 months ago

https://docs.microsoft.com/en-us/azure/app-service/overview#app-service-on-linux

upvoted 2 times

Edwardenis 2 years, 7 months ago

I don't agree with you.
1. the link you posted doesn't say anything about restriction on mixing App service plan in the same region.
2. I just created 2 app service plan in the same RG. One app service plan on windows in the region of the RG East US and the other app service plan in Linux in Central US

upvoted 1 times

azurelearner666 2 years, 4 months ago

It does if you read it... see under limitations, it says "You can't mix Windows and Linux apps in the same App Service plan."

upvoted 6 times

dookebroyle 1 year, 4 months ago

It also states your current resource group locations do not support Linux. So even though you can in your situation, the question specifies that that is not an option

upvoted 4 times

lde Highly Voted 2 years, 8 months ago

https://docs.microsoft.com/en-us/azure/app-service/overview#app-service-on-linux
Historically, you can't mix Windows and Linux apps in the same resource group. However, all resource groups created on or after January 21, 2021 do support this scenario. For resource groups created before January 21, 2021, the ability to add mixed platform deployments will be rolled out across Azure regions (including National cloud regions) soon.

upvoted 20 times

azurelearner666 2 years, 4 months ago

This should be updated on the exam to reflect reality... so we have to respond wrong now to pass... LOL

upvoted 5 times

Edwardenis 2 years, 7 months ago

Ohhh Thanks for the notice, thats why I could create both app service plans.

upvoted 2 times

pandugadu009 Most Recent 1 month, 1 week ago

the given answer is correct. the az group update command can not update the location. so you need to create a new group.

upvoted 2 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 8 times

surprise0011 6 months, 2 weeks ago

you can't update the resource group, you will get an error.
So it has to be: az group create

az group update --name acrrg --set location=centralus

(InvalidResourceGroupLocation) Invalid resource group location 'centralus'. The Resource group already exists in location 'westeurope'.
Code: InvalidResourceGroupLocation
Message: Invalid resource group location 'centralus'. The Resource group already exists in location 'westeurope'.

upvoted 3 times

ucskips 7 months, 3 weeks ago

Once the region of a resource group is set it cannot be updated. I read this before but I don't where

upvoted 2 times

Esward 8 months, 2 weeks ago

It was there in 13 Feb 2023 exam

upvoted 1 times

obeesubz5 9 months ago

All you need here is to create an app service plan and the web app.
Secondly, no one said the existing rg contain Windows vm's.
Thirdly, it does not say that the existing rg's don't support linux, it says "the locations" of the rg don't support linux. So here you have to bare in mind that a resource need not be in the same region as the rg. You can google, in fact try it out yourself. This supports the "Limit number of resource groups" requirement.

upvoted 3 times

Esward 9 months, 1 week ago

Given answers are correct!!

upvoted 1 times

ghassen204 10 months, 2 weeks ago

I think:
1: az group update
2: az appservice plan create
3: az webapp create
Note: Linux and Windows App Service plans can now share resource groups. This limitation has been lifted from the platform and existing resource groups have been updated to support this.

upvoted 3 times

josecouva 10 months ago

But in the queston it says "None of your current resource groups reside in a location that supports Linux", so you can't execute az group update.
I think the answer given is correct

upvoted 3 times

[Removed] 1 year, 1 month ago

It says the Resource groups are not in a location that supports Linux, but you can create a resource in a location other than the resource group location. So, I do not understand why there is a need to create a new RG. But I do agree that given the options the answer is correct

upvoted 3 times

Enigma___ 1 year, 1 month ago

Linux and Windows App Service plans can now share resource groups. This limitation has been lifted from the platform and existing resource groups have been updated to support this.
https://learn.microsoft.com/en-us/azure/app-service/overview#app-service-on-linux

upvoted 1 times

[Removed] 1 year, 2 months ago

* The question says ResourceGroup is residing in a location that does not support Linux. But where the resource-group resides has no relation to where the Resources (App service, for example) is residing.
* Reduce the number of resource group. The totally useless comment like that is just given to create confusion for the benefit of the certification industry

The whole discussion in the comments about Windows and Linux App service plan cannot co locate in an RG is totally unrelated. Nobody said about a Windows app service plan. Either the question is wrong or Microsoft just cannot get anything right

upvoted 1 times

xRiot007 1 year, 4 months ago

1. Create group (it very states that your current resources do *not* support Linux)
2. Create service plan
3. Create the web app

upvoted 5 times

silvatyrant 1 year, 8 months ago

1) Linux and Windows apps can now be deployed in the same RG.
2) They mention doing an initial deployment as well. "az webapp up" is best suited to do this. Not "az webapp create". The "create" action will not do an initial deployment.

upvoted 3 times

dookebroyle 1 year, 4 months ago

I agree, the "right" answer doesn't deploy the app. Webapp up will create a new resource group and app service plan though, so the answer is the same. Your current resource group locations don't support Linux, so your only option was to move them (not supported) or create a new one.

upvoted 1 times

Mev4953 1 year, 9 months ago

Got this in the exam 01/22

upvoted 3 times

Mev4953 1 year, 9 months ago

No,sorry i didnt. I read wrong :)

upvoted 2 times

resonant 1 year, 10 months ago

A couple of things to note in this question:
1. It says "Which three Azure CLI commands should you use to develop the solution?" which means that you HAVE to choose 3 options. Therefore, even if it looks like you don't have to create or update a resource group, that should be enough hint that at least you should choose one of them.
2. You can't change the location of a resource group once it is created. Not even with the "az group update" command.
With 1. and 2. it should be clear enough that you have to create the resource group anyways, even if you didn't pay attention to the fact that the question says none of the locations of the resource groups already created can have Linux (independently of the fact that Linux and Windows can NOW coexist in the same resource group).

upvoted 13 times

AmineD 11 months, 3 weeks ago

I totaly agree with the fact that resource group location cannot be changed after creation and since "None of your current resource groups reside in a location that supports Linux" so we have to use "az group create" instead of "az group update"

For me the answer is correct

Reference: https://learn.microsoft.com/en-us/rest/api/resources/resource-groups/create-or-update

upvoted 1 times

Question #11 *Topic 2*

DRAG DROP -

Fourth Coffee has an ASP.NET Core web app that runs in Docker. The app is mapped to the www.fourthcoffee.com domain.

Fourth Coffee is migrating this application to Azure.

You need to provision an App Service Web App to host this docker image and map the custom domain to the App Service web app.

A resource group named FourthCoffeePublicWebResourceGroup has been created in the WestUS region that contains an App Service Plan named AppServiceLinuxDockerPlan.

Which order should the CLI commands be used to develop the solution? To answer, move all of the Azure CLI commands from the list of commands to the answer area and arrange them in the correct order.

Select and Place:

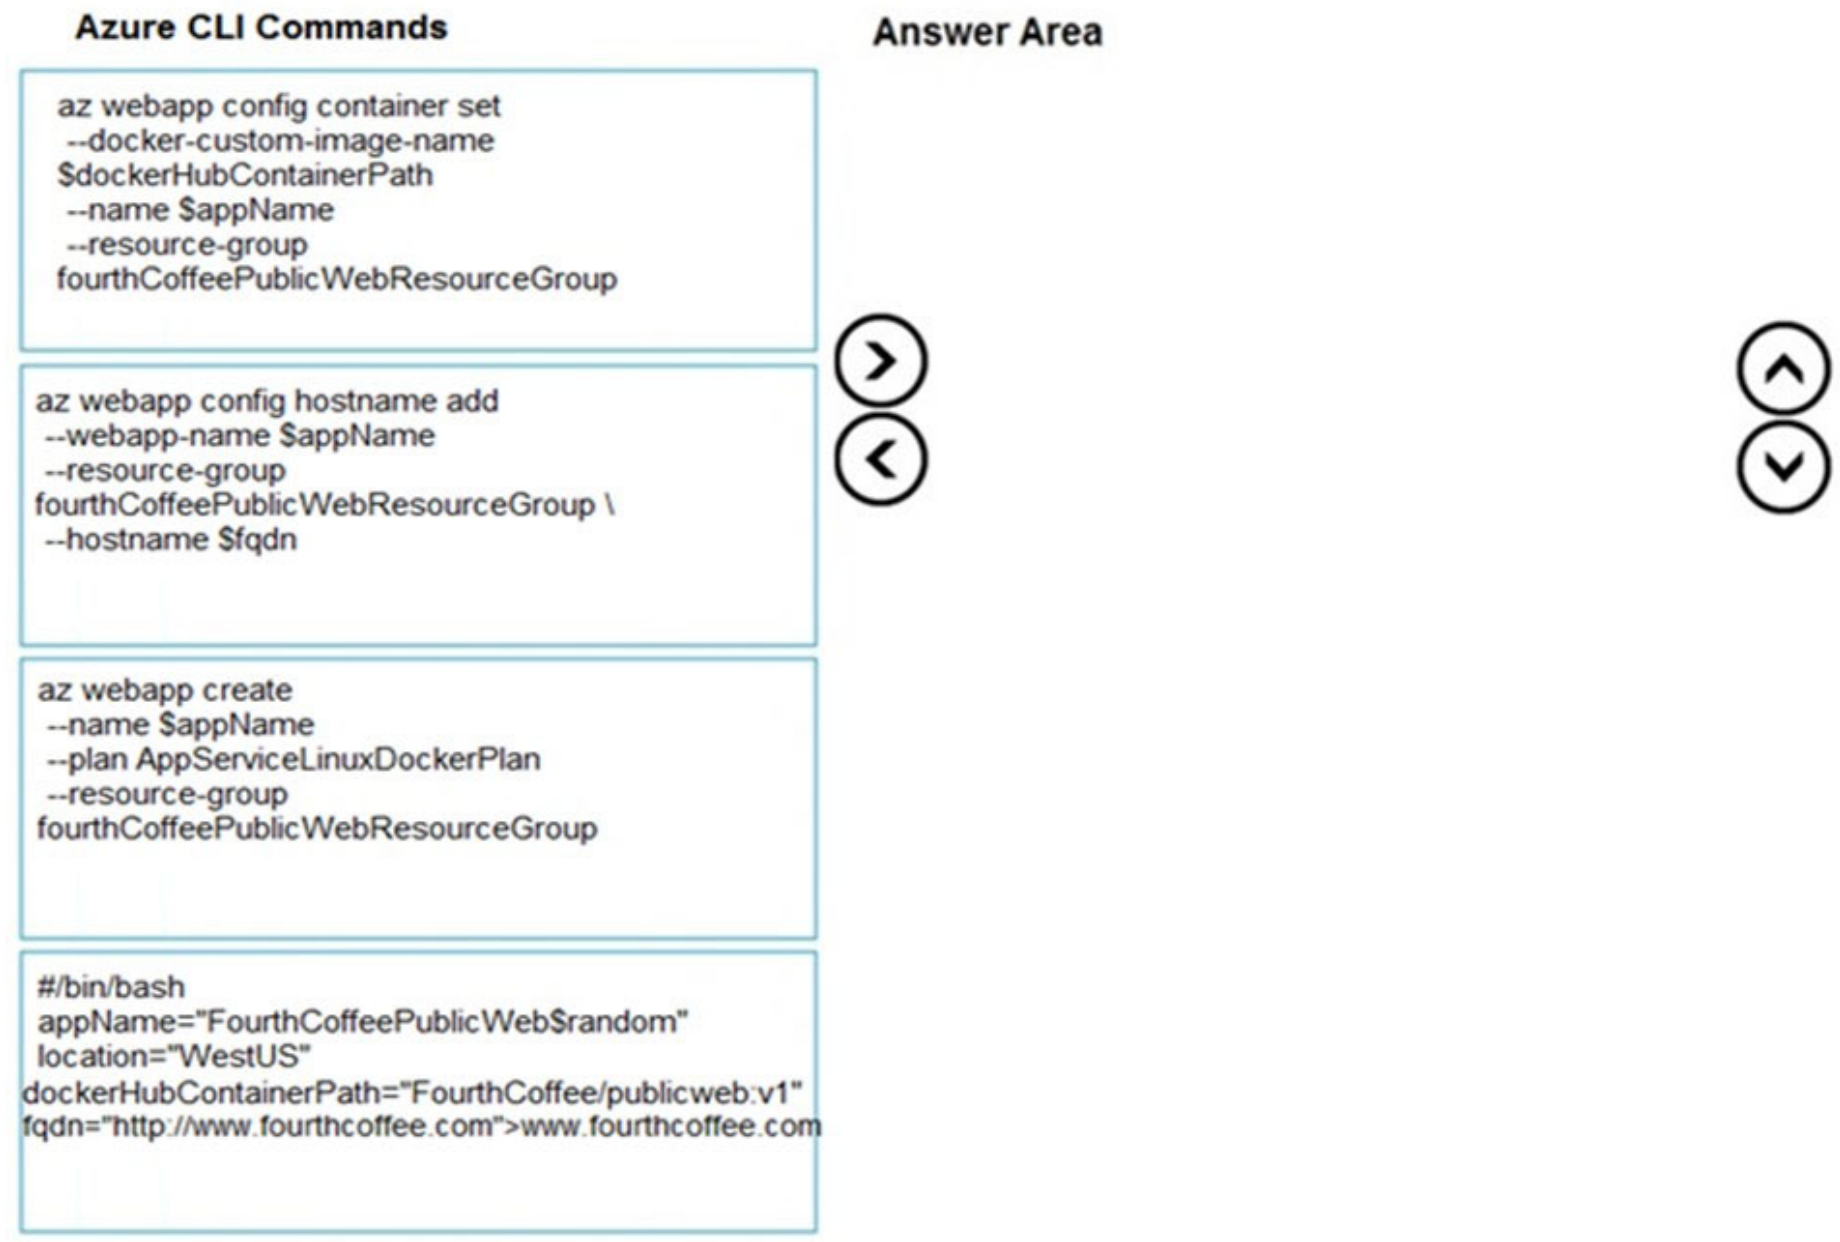

Correct Answer:

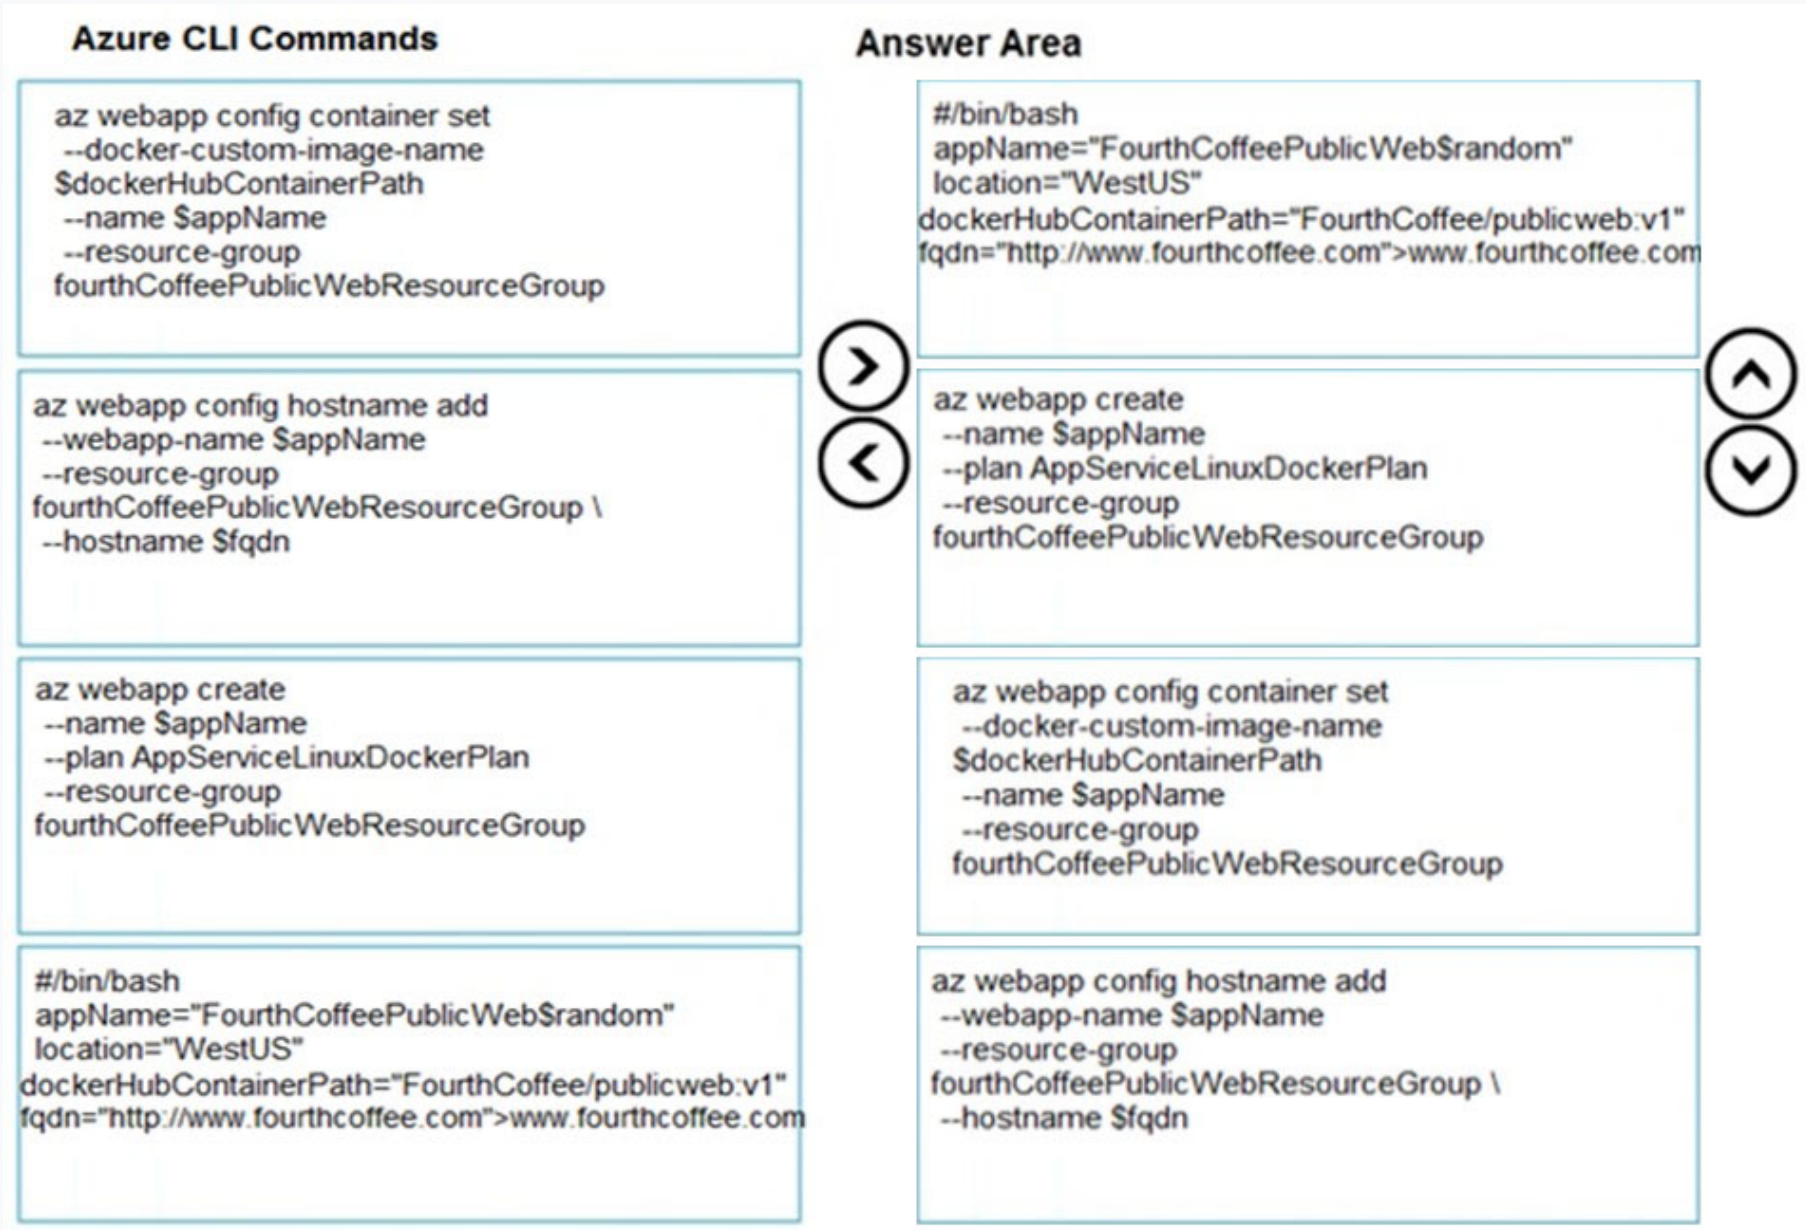

Step 1: #bin/bash -

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

Step 2: az webapp create -

Create a web app. In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command.

Step 3: az webapp config container set

In Create a web app, you specified an image on Docker Hub in the az webapp create command. This is good enough for a public image. To use a private image, you need to configure your Docker account ID and password in your Azure web app.

Step 4: az webapp config hostname add

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

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

Reference:

https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image https://docs.microsoft.com/en-us/azure/app-service/tutorial-custom-container?pivots=container-linux https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-configure-custom-domain

msuser11 Highly Voted 1 year, 5 months ago

1. /bin/bash
2. az webpp create
3. ~ config container set
4. ~ config hostname add

upvoted 91 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went above answer, score 926

upvoted 8 times

ayush66203 1 week, 2 days ago

were all the questions in the exam from exam topic? did you prepare from paid version of exam topic or the free one?

upvoted 1 times

nmatnoh 1 year, 4 months ago

I would agree with the propose solution

1. /bin/bash
2. az webpp create
3. ~ config container set

URL: https://docs.microsoft.com/en-us/azure/app-service/tutorial-custom-container?pivots=container-linux

At the bottom of tutorial, link to next tutorial
In the next tutorial, you learn how to map a custom DNS name to your app.

4. ~ config hostname add
https://docs.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-custom-domain

upvoted 5 times

Mousavi 1 year, 4 months ago

I still have doubts about the order of 3 and 4 ?

upvoted 4 times

GMartinez Highly Voted 1 year, 5 months ago

Shouldn' it be az webapp create as 2nd, az webapp config container set as 3rd and az webapp config hostname as last ?

upvoted 11 times

nightflyer 1 year, 5 months ago

Yes this is correct, i saw this question in the Microsoft Learn Labs

upvoted 4 times

JH81 Most Recent 4 months ago

Got this on 6/28/2023 and passed with 850. Went with highly voted answer.

upvoted 3 times

Maximuskh 4 months, 2 weeks ago

I think the reason container needs to be set and then hostname is because if we configure hostname first while the app does not have docker image configuration, the end users experience downtime.

upvoted 2 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

FYI:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

Saluk_DE 7 months ago

Question was in Exam 2023-03-30

upvoted 2 times

sarmaria 7 months, 2 weeks ago

Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 2 times

Esward 9 months, 1 week ago

Given answers are correct!

https://docs.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-custom-domain

upvoted 1 times

HafizSalmanMalik 10 months ago

Correct sequence is already mentioned.
1. /bin/bash
2. az webpp create
3. ~ config container set
4. ~ config hostname add

upvoted 1 times

kociuba 10 months ago

got it on my exam 30-12-2022
score: 818

upvoted 2 times

youagrar 10 months, 1 week ago

Correct. Got this in the exam today dec 23, 2022

upvoted 2 times

SoftSol 10 months, 3 weeks ago

Correct. Got this in the exam today dec 13, 2022

upvoted 2 times

jeanfmc 1 year, 2 months ago

order 3421

upvoted 1 times

ManoDevops 1 year, 2 months ago

The correct order is :
1. /bin/bash
2. az webpp create
3. ~ config container set
4. ~ config hostname add

If web-app config is executed before creating a web app it would result with the below error:
az webapp config hostname add --webapp-name mano588new -g 1-f2dc32a3-playground-sandbox --hostname www.mano.com
(ResourceNotFound) The Resource 'Microsoft.Web/sites/mano588new' under resource group '1-f2dc32a3-playground-sandbox' was not found. For more details please go to https://aka.ms/ARMResourceNotFoundFix
Code: ResourceNotFound
Message: The Resource 'Microsoft.Web/sites/mano588new' under resource group '1-f2dc32a3-playground-sandbox' was not found. For more details please go to https://aka.ms/ARMResourceNotFoundFix

upvoted 2 times

Pize 1 year, 4 months ago

Does anyone audit these? How can you add config to a webapp that isn't created?

upvoted 2 times

dookebroyle 1 year, 4 months ago

Does the order of the webapp configs matter?

upvoted 2 times

nightflyer 1 year, 5 months ago

The answer is wrong, you can't configure a web app you haven't created yet

upvoted 5 times

Question #12 *Topic 2*

DRAG DROP -

You are developing a serverless Java application on Azure. You create a new Azure Key Vault to work with secrets from a new Azure Functions application.

The application must meet the following requirements:

- Reference the Azure Key Vault without requiring any changes to the Java code.
- Dynamically add and remove instances of the Azure Functions host based on the number of incoming application events.
- Ensure that instances are perpetually warm to avoid any cold starts.
- Connect to a VNet.
- Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.

You need to grant the Azure Functions application access to the Azure Key Vault.

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

Select and Place:

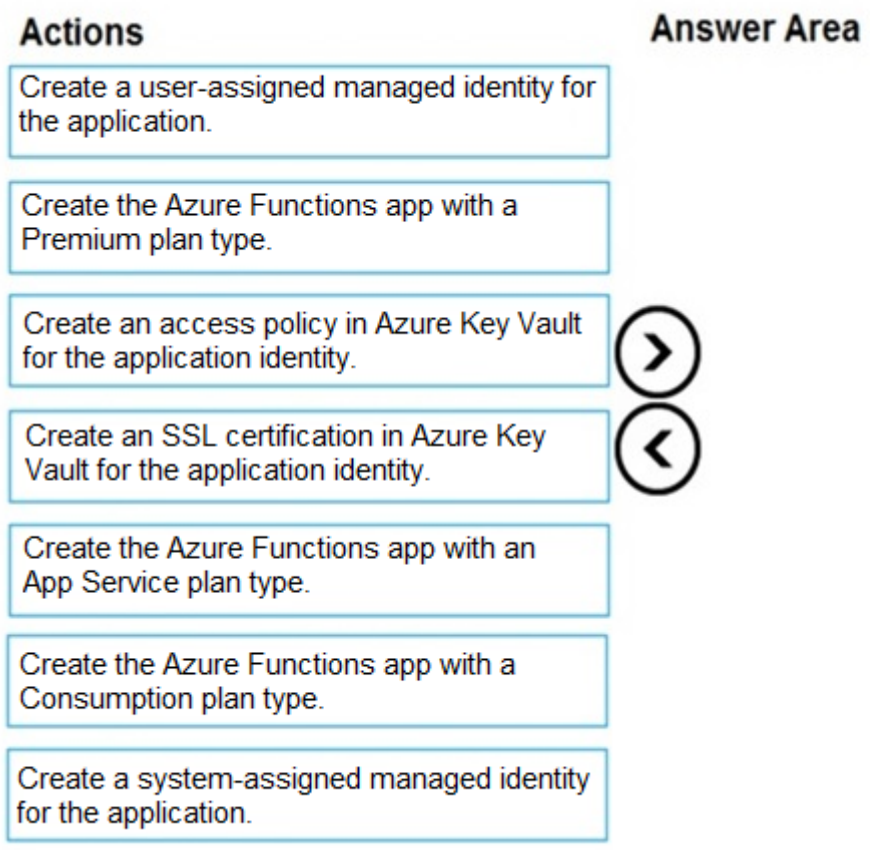

Correct Answer:

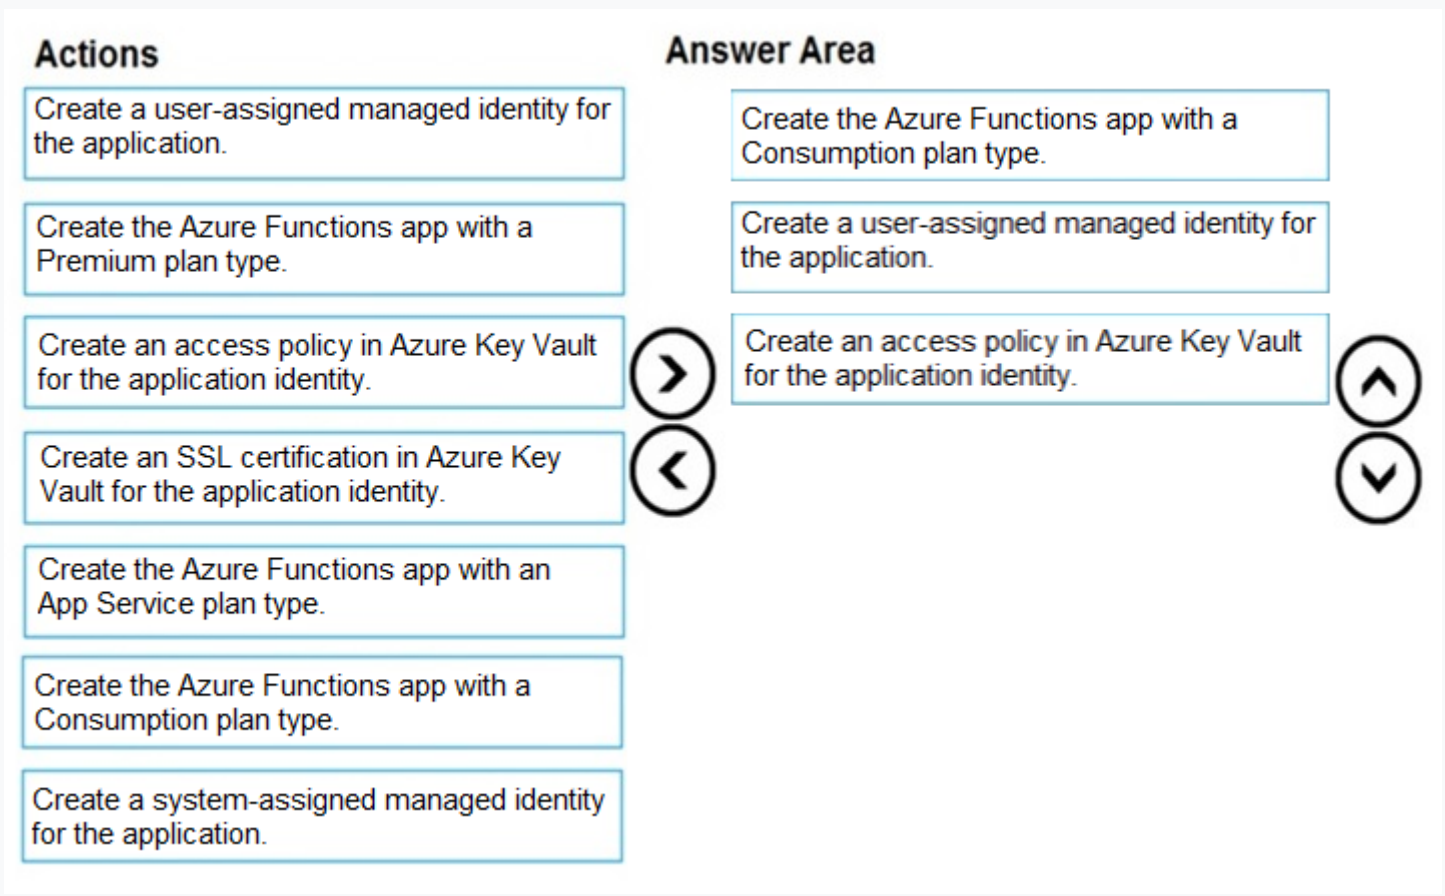

Step 1: Create the Azure Functions app with a Consumption plan type.

Use the Consumption plan for serverless.

Step 2: Create a system-assigned managed identity for the application.

Create a system-assigned managed identity for your application.

Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.

Step 3: Create an access policy in Key Vault for the application identity.

Create an access policy in Key Vault for the application identity you created earlier. Enable the "Get" secret permission on this policy. Do not

configure the
"authorized application" or applicationId settings, as this is not compatible with a managed identity.
Reference:
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

msuser11 Highly Voted 1 year, 5 months ago
1. create ~Premium plan Type (Consumption X)
2. create system-assigned ~ (user-assigned X)
3. create an access policy in Azure Key Vault~
upvoted 142 times

NaSit 1 year, 4 months ago
I agree with you.
1. Premium plan (avoid any cold starts and connect to a VNet)
Overview of plans here: https://docs.microsoft.com/th-th/azure/azure-functions/functions-scale
2. create system-assigned => "A system-assigned identity is tied to your application and is deleted if your app is deleted."
3. create an access policy
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references?toc=%2Fazure%2Fazure-functions%2Ftoc.json&tabs=azure-cli
upvoted 24 times

Xardas 3 months ago
Why not App Service plan?
upvoted 1 times

macobuzi 2 months, 2 weeks ago
App Service plan can connect to Vnet but it won't "Ensure that instances are perpetually warm to avoid any cold starts".
upvoted 2 times

imheretolearn 1 month ago
Wouldn't Always On setting on App Service plan work here?
upvoted 2 times

jakobste 2 months, 1 week ago
I agreee. I would go with system assigned identity unless the question talks about "multiple apps/vms". In that case you would go for user assigned if "administration has to be minimized" etc.
upvoted 1 times

fkaracan 8 months, 1 week ago
correct
upvoted 1 times

Esward 9 months, 1 week ago
Agreed
upvoted 1 times

uffuchsi Highly Voted 8 months, 1 week ago
Received this in my exam today (22/02/2023). Selected 'Created the Azure Function app with Premium plan type', 'Create a system-assigned managed identity for the application', and 'Create an access policy in Azure Key Vault for the application identity'. Score 927.
upvoted 14 times

Samuel316 Most Recent 3 weeks ago
Answer shown in the image does not agree with what's written in the answer description. Image says user-assigned managed identity. Description says system-assigned managed identity, which would be more correct
upvoted 2 times

Tarajee 1 month, 2 weeks ago
Got this on 2023sept
upvoted 2 times

kvtborad 2 months, 3 weeks ago
I have this question on 6th August. it was passed with 904. chose this highly voted:create ~Premium plan Type (Consumption X)
2. create system-assigned ~ (user-assigned X)
3. create an access policy in Azure Key Vault~
upvoted 2 times

tom112 5 months, 3 weeks ago
system-assigned Managed Identity's life cycle:
Shared life cycle with the Azure resource that the managed identity is created with.
When the parent resource is deleted, the managed identity is deleted as well.
upvoted 1 times

ucskips 7 months, 3 weeks ago

I agree with the answers here 1. Create function app with premium plan (premium plan allows for virtual network connection and pre-warmed starts) 2. System assigned managed identity (when the function app is deleted so is the system identity they are "tied" together https://devblogs.microsoft.com/devops/demystifying-service-principals-managed-identities/) 3. create an access policy (here, the system identity can be set as the method of accessing the key vault**)

upvoted 1 times

ucskips 7 months, 3 weeks ago

I agree with the answers here

1. Create function app with premium plan (premium plan allows for virtual network connection and pre-warmed starts)
2. System assigned managed identity (when the function app is deleted so is the system identity they are "tied" together https://devblogs.microsoft.com/devops/demystifying-service-principals-managed-identities/)
3. create an access policy (here, the system identity can be set as the method of accessing the key value)

upvoted 1 times

CODE_STS 8 months ago

Got this in the exam today! Feb 28, 2023

upvoted 1 times

Harish86 8 months, 1 week ago

Can you tell me, why this website is showing incorrect answers? And most importantly from where it is this answers. When it is saying about Warm tiers , how can it be Consumption plan

upvoted 2 times

adilkhan 9 months, 1 week ago

consumption plan does not support VNET https://learn.microsoft.com/en-us/azure/azure-functions/functions-networking-options?tabs=azure-cli

upvoted 2 times

micro9000 11 months, 3 weeks ago

The Azure Function App must be in Premium plan because we need to ensure that instances are perpetually warm to avoid any cold starts.

upvoted 1 times

vruizrob 12 months ago

The following requirements are the key:

☞ Ensure that instances are perpetually warm to avoid any cold starts.
In this link, https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#cold-start-behavior, you can read this same phrase, so, the correct answer is Premium Plan

☞ Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.
In this link, https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview#managed-identity-types, you can see that is the System-assigned the correct answer

☞ And access policy in Azure Key Vault

upvoted 4 times

vruizrob 1 year, 1 month ago

1.- Consumption -> https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale
"On the Consumption plan, instances of the Functions host are dynamically added and removed based on the number of incoming events."
2.- System Assigned, because it says "Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted"
3.- Access Policy

upvoted 2 times

[Removed] 1 year, 1 month ago

You are focusing on that one requirement. But it has to be warm and it has to connect to a vnet. Which I think consumption plan cannot handle. But it is silly that Microsoft expects developers to remember all the unfortunate plans and random options they added to them.

upvoted 7 times

sb8498 12 months ago

I don't think a Function hosted in a Consumption plan connect to a VNET, though.
https://learn.microsoft.com/en-us/azure/azure-functions/functions-networking-options?tabs=azure-cli

upvoted 1 times

DivyaRajkumar 1 year, 2 months ago

What is the correct answer for this question? Very confusing.

upvoted 1 times

franklinfocus 1 year, 1 month ago

Premium plan
system-assigned
access policy in Azure Key Vault

upvoted 5 times

rohitpatil113 1 year, 2 months ago

1. create ~Premium plan Type (Consumption X)
2. create system-assigned ~ (user-assigned X)
3. create an access policy in Azure Key Vault~

upvoted 1 times

d_vaibhav 1 year, 2 months ago

Create Azure Functions App with Premium Plan type - (Perpetually warm instances to avoid any cold start)
Create System assigned managed identity for the application - (enables vault to be automatically removed if the Azure Function application is deleted)
Create an access policy in azure key vault for the application identity.

upvoted 1 times

Question #13 *Topic 2*

You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published.
You must ensure that the website remains available and responsive while minimizing cost.
You need to deploy the website.
What should you do?

A. Deploy the website to a virtual machine. Configure the virtual machine to automatically scale when the CPU load is high.

B. Deploy the website to an App Service that uses the Shared service tier. Configure the App Service plan to automatically scale when the CPU load is high.

C. Deploy the website to a virtual machine. Configure a Scale Set to increase the virtual machine instance count when the CPU load is high.

D. Deploy the website to an App Service that uses the Standard service tier. Configure the App Service plan to automatically scale when the CPU load is high.

Correct Answer: *D*
Windows Azure Web Sites (WAWS) offers 3 modes: Standard, Free, and Shared.
Standard mode carries an enterprise-grade SLA (Service Level Agreement) of 99.9% monthly, even for sites with just one instance.
Standard mode runs on dedicated instances, making it different from the other ways to buy Windows Azure Web Sites.
Incorrect Answers:
B: Shared and Free modes do not offer the scaling flexibility of Standard, and they have some important limits.
Shared mode, just as the name states, also uses shared Compute resources, and also has a CPU limit. So, while neither Free nor Shared is likely to be the best choice for your production environment due to these limits.

Community vote distribution
D (100%)

profesorklaus Highly Voted 2 years, 10 months ago
Answer is correct. No doubts!
upvoted 93 times

27close Highly Voted 2 years, 11 months ago
answer is d- this is production and requires autoscaling when cpu is peak. the rest is constraints by limit without auto scaling
upvoted 21 times

RuffBoii Most Recent 1 month, 1 week ago
Had this on my exam today.
upvoted 2 times

juanckar 3 months, 3 weeks ago
This was on the exam (July 2023). Went with proposed. Scored 917
upvoted 2 times

alexein74 9 months, 1 week ago
Selected Answer: D
D - To ensure that the website remains available and responsive while minimizing cost, you should deploy the website to an App Service that uses the Standard service tier and configure the App Service plan to automatically scale when the CPU load is high. This way, the website can handle high traffic volumes by automatically scaling the number of instances of the website, reducing the risk of the website becoming unavailable due to high traffic.
upvoted 4 times

ReyPirata 2 months, 1 week ago
correct, in 08/20/2023, score: 925/1000
upvoted 1 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000
upvoted 2 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: D
Answer is correct.
upvoted 3 times

nhannn 1 year, 1 month ago

Selected Answer: D

The answer is D, it is very obvious both C and D are correct but we are aiming for the minimizing cost, then it should be D

upvoted 3 times

kyawtlearning 1 year, 3 months ago

The answer is correct to save the cost.

upvoted 1 times

jmirza 1 year, 3 months ago

I think the scale for the Wwe Apps should be based on Network and RAM mostly.

upvoted 1 times

AZAdam22 1 year, 3 months ago

Selected Answer: D

D - Because app service standard tier is the cheapest tier that supports auto scale.

upvoted 2 times

dookebroyle 1 year, 4 months ago

I think this would depend on whether you want to have to manage the VMs. If you choose VMs, you are responsible for OS and updates. I would think D would be the right answer

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: D

D is correct answer.

upvoted 2 times

drgn7676 1 year, 6 months ago

Selected Answer: D

Uses the Standard pricing tier

upvoted 2 times

DarkyShadoW 1 year, 6 months ago

Selected Answer: D

D. AppService Standar with AutoScale

upvoted 2 times

bubbleGumStroganoff 1 year, 7 months ago

I think C.
All other options could potentially be vertically scaling in which case outage or degradation of responsiveness will occur.

upvoted 1 times

Bartimaeus 1 year, 8 months ago

Selected Answer: D

I think the question C vs D is not really about price, because the same VMs you would run with App Service would cost you much less with a scale set. Even with Load Balancer and storage it's cheaper (at least 30$/month per instance).
However the answer says:
1. Deploy the website to a virtual machine
2. Configure a Scale Set to increase the virtual machine instance count when the CPU load is high.
Which is in fact incorrect, because you deploy the app to the scale set - you can't add an existing VM with data to a scale set or configure it to add VM instances.
You deploy a "VM scale set" as a entire entity. You can attach a custom VHD there, but it's not stated in the answer.

upvoted 6 times

ashishlakhera 1 year, 8 months ago

Selected Answer: D

selected D

upvoted 1 times

Question #14 *Topic 2*

HOTSPOT -

A company is developing a Java web app. The web app code is hosted in a GitHub repository located at https://github.com/Contoso/webapp.

The web app must be evaluated before it is moved to production. You must deploy the initial code release to a deployment slot named staging.

You need to create the web app and deploy the code.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup
```

az [▼] create --location centralus --name $resourcegroupname

az [▼] create --name $webappname --resource-group $resourcegroupname --sku S3

az [▼] create --name $webappname --resource-group $resourcegroupname --plan $webappname

az [▼] create --name $webappname --resource-group $resourcegroupname --slot staging

az [▼] config --name $webappname --resource-group $resourcegroupname \
--slot staging --repo-url $gitrepo --branch master --manual-integration

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

Correct Answer:

Answer Area

gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] create --location centralus --name $resourcegroupname

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] create --name $webappname --resource-group $resourcegroupname --sku S3

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] create --name $webappname --resource-group $resourcegroupname --plan $webappname

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] create --name $webappname --resource-group $resourcegroupname --slot staging

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] config --name $webappname --resource-group $resourcegroupname \
--slot staging --repo-url $gitrepo --branch master --manual-integration

Box 1: group -

Create a resource group.

az group create --location westeurope --name myResourceGroup

Box 2: appservice plan -

Create an App Service plan in STANDARD tier (minimum required by deployment slots). az appservice plan create --name $webappname --resource-group myResourceGroup --sku S1

Box 3: webapp -

Create a web app.

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

Box 4: webapp deployment slot -

#Create a deployment slot with the name "staging".

az webapp deployment slot create --name $webappname --resource-group myResourceGroup \
--slot staging

Box 5: webapp deployment source -

Deploy sample code to "staging" slot from GitHub.

az webapp deployment source config --name $webappname --resource-group myResourceGroup \
--slot staging --repo-url $gitrepo --branch master --manual-integration

Reference:

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

sskPL Highly Voted 2 years, 12 months ago

Given answer is correct.

upvoted 188 times

Harshad_Mehta Highly Voted 2 years, 11 months ago

why do you guys have to comment if the answer is correct.

upvoted 34 times

thomas204 2 years, 11 months ago

Because some answers are wrong and commenting when correct is just comforting for others.

upvoted 230 times

azurelearner666 2 years, 4 months ago

Agree! love this site because of the community and the discussions. they also help to understand the reasoning!!

upvoted 39 times

macobuzi 2 months, 2 weeks ago

Yes, many dump test websites have these questions but I still had to move to this site for the community!

upvoted 1 times

lasyan3 2 years, 1 month ago

Totally agreed, thanks to all the community members :)

upvoted 11 times

thebarber87 7 months ago

So vote. It is annoying when a question has 80 comments on it but they're just all answer is correct.

upvoted 1 times

azurelearner666 2 years, 4 months ago

you commented too as well, right? ;)

upvoted 12 times

DonH 4 months, 3 weeks ago

No, he asked a valid question. @thomas204 gave a valid answer. I also agree on that. Some answers are wrong so it's satisfying to read the comments when an answer is correct, just to be on the safe side.

upvoted 1 times

nlk0 Most Recent 2 months ago

Got this today 26/08/2023, went with the given. 880
Case study was contoso -> 9 questions
+ 35 question, two series.
3 of them were not from examptopics, all others were

upvoted 2 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam at 10-June-2023

upvoted 4 times

Videira 8 months ago

On my exam 2023-02-25

upvoted 4 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23
score: 817

upvoted 4 times

saravanasanthosh 10 months ago

Got this on 12/30/2022 exam. The given solution is correct

upvoted 2 times

kociuba 10 months ago

got it on my exam 30-12-2022
score: 818

upvoted 3 times

Suwani 10 months, 1 week ago

Given answer is correct

upvoted 1 times

mdg3501 10 months, 3 weeks ago

got this on 2022-12-7

upvoted 1 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 2 times

rol204 1 year ago

Box 1: group
Create a resource group.

Box 2: appservice plan
Create an App Service plan in STANDARD tier (minimum required by deployment slots)

Box 3: webapp
Create a web app

Box 4: webapp deployment slot
#Create a deployment slot with the name "staging"

Box 5: webapp deployment source
Deploy sample code to "staging" slot from GitHub

upvoted 2 times

Anitay 1 year, 1 month ago

Answers are correct...

upvoted 1 times

Elsheimy 1 year, 5 months ago

On exam, 05/22

upvoted 2 times

Evo_Morales 1 year, 6 months ago

Saw this question on a recent test

upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 3 times

Chiboy 1 year, 9 months ago

I don't know if I am the only one. But I find that even when the answers are correct, the comments from us give fresh and refreshing incite into the answers. I have severally been reminded of some factors I did not take into consideration when arriving at an answer, even though correct. I cannot thank this forum enough. Please let us continue commenting on every answer(correct and incorrect). We are greatly helping each other. Thanks.

upvoted 10 times

Question #15 *Topic 2*

HOTSPOT -

You have a web service that is used to pay for food deliveries. The web service uses Azure Cosmos DB as the data store.

You plan to add a new feature that allows users to set a tip amount. The new feature requires that a property named tip on the document in Cosmos DB must be present and contain a numeric value.

There are many existing websites and mobile apps that use the web service that will not be updated to set the tip property for some time.

How should you complete the trigger?

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
function ensureTip() {
    var r = [ __.value(); | __.readDocument('item'); | getContext().getRequest(); | getContext().getResponse(); ]
    var i = r.getBody();
    [ if (!("tip" in i)) { | if (request.getValue("tip") = = = null) { | if (isNaN(i)["tip"] || i["tip"] = = = null) { | if (typeof __.pluck("tip") = = 'number') { ]
        i["tip"] = 0;
    }
    [ r.setBody(i); | r.setValue(i); | __.upsertDocument(i); | __.replaceDocument(i); ]
}
```

Correct Answer:

Answer Area

```
function ensureTip() {
    var r = [ __.value(); | __.readDocument('item'); | getContext().getRequest(); | getContext().getResponse(); ]
    var i = r.getBody();
    [ if (!("tip" in i)) { | if (request.getValue("tip") = = = null) { | if (isNaN(i)["tip"] || i["tip"] = = = null) { | if (typeof __.pluck("tip") = = 'number') { ]
        i["tip"] = 0;
    }
    [ r.setBody(i); | r.setValue(i); | __.upsertDocument(i); | __.replaceDocument(i); ]
}
```

dol Highly Voted 2 years, 11 months ago

Right answer in second drop down is the first one (..."tip" in i...)
Similiar example can be found on https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-write-stored-procedures-triggers-udfs

upvoted 163 times

MikeAWS 1 day, 11 hours ago

Correct! Second dropdown is: if(!("tip" in i)) {
answer can also be found here: https://www.whizlabs.com/blog/microsoft-azure-az-204-exam-questions//

upvoted 1 times

Juanlu 2 years, 11 months ago

After develop a simple Javascript test, de correct option is (..."tip" in i). In cases where "tip" is no defined, sentence "if(isNaN(i)["tip"] || i["tip"] === null)", return "undefined", so it doesn't work properly !
Here this simple test:
---
var i = { tip: 5, name: "just a test"}

if ('tip' in i){
console.log("YES");
}
else {
console.log("NO");
}

console.log("Option 2");

if(isNaN(i)["tip"] || i["tip"] === null){ // Undefined when "tip: N" not in array !
console.log("NO");
}else {
console.log("YES");
}

upvoted 10 times

nychollas 2 years, 4 months ago

because the sintax in the proposed option is NaN(i)["tip"]... it should have been Nan(i["tip"]). In this case, this option might be taken into account as a possible correct answer. This is my opinion.

upvoted 2 times

Satheesh88 2 years, 11 months ago

question says tip value should be numeric. What if is gave var i = { tip: 'junk', name: "just a test"}
Result:
"YES"
"Option 2"
"NO"

upvoted 6 times

thomastrain 2 years, 10 months ago

Sadheesh88 has a valid point. Although the question states "there will be many apps not updating their value for some time", a good solution should always validate user input. If the value of "tip" is null or "junk", the first option will fail.

Option 3 has a typo and is missing a closing parenthesis. If we make a small change to option 3, so it reads:
if(isNaN(i["tip"]) || i["tip"] === null){
then this becomes the best answer. Check for yourself with this fiddle:
https://jsfiddle.net/d4x3fota/2/

So I think it comes down to if the answers were copied correctly.

upvoted 21 times

DonH 4 months, 3 weeks ago

thanks for the Fiddle.

upvoted 1 times

zoan 2 years, 11 months ago

I tested the isNaN in and I get 'true' when 'tip' is undefined, null or "a" . But I use isNaN(i['tip']) and not isNaN(i)['tip']. For 'tip' in i I get only if tip is defined the i object regardless of the type of the value.

upvoted 1 times

Tom87 2 years, 8 months ago

isNaN(null) is false. That's why i["tip"] === null is needed too.

upvoted 1 times

jakobste 2 months, 1 week ago

Yes. This question is on Wizlabs also with a variation, and the code there is (..."tip" in i...).

upvoted 1 times

sujitwarrier11 2 years, 4 months ago

where is the validation that the tip value is a number? so the current answer is right. isNan checks if the value is not a number. if true it sets value to 0.

upvoted 8 times

altafpatel1984 1 year, 11 months ago

But they suggest IsNaN(i), not isNaN(i["tip"]) which makes no sense as syntax is incorrect.

upvoted 6 times

MiraA 2 years ago

My opinion to NaN()...
The "old" clients will send no "tip" value at all => set "tip" to 0 to match the new expectations.
The "new" clients will send some "tip" value => pass "tip" value to a web service as it is; the web service has possibility to detect NaN() itself and to reject the request as invalid with proper HTTP status code (and log such invalid request to detect malformed clients). The client knows something went wrong and that his request wasn't processed.

I consider 0 to be a default value of "tip" which will be used for "old" clients only.
Correcting wrong (NaN) values in "tip" to 0 at this moment is bad idea as it hides a problem with the client's communication.

A reference to "a numeric value" requirement in the assignment is related to CosmosDB primarily I think.

upvoted 1 times

Skyrocket Highly Voted 2 years, 9 months ago

Admin, requst you to update answers to avoid any further confusion.
1. getRequest
2. (!"tip" in i)
3. setBody

upvoted 91 times

sujitwarrier11 2 years, 4 months ago

where is the validation that the tip value is a number? so the current answer is right. isNan checks if the value is not a number. if true it sets value to 0.

upvoted 4 times

Basu525 2 years, 8 months ago

this is the correct answer. admins please rectify it.

upvoted 3 times

sujitwarrier11 2 years, 4 months ago

where is the validation that the tip value is a number? so the current answer is right. isNan checks if the value is not a number. if true it sets value to 0.

upvoted 3 times

Poops 10 months ago

isNaN(i) checks that i is NaN, it will always evaluate to false. That whole snippet is bad, it only compiles because JavaScript has a very forgiving syntax

upvoted 1 times

btrump 2 years, 8 months ago

This is not correct. The prompt states, "must be present and contain a numeric value."

if i["tip"] == "foo" your passes but fails to meet the requirement.

upvoted 11 times

jasifu3 1 year, 8 months ago

your argument is valid, but there is no better alternative. isNaN(i)["tip"] is evaluated as false["tip"] -> undefined every time.

upvoted 2 times

Forhallf 2 years, 8 months ago

Skyrocket is correct.

```
function validateToDoItemTimestamp() {
var context = getContext();
var request = context.getRequest();

// item to be created in the current operation
var itemToCreate = request.getBody();

// validate properties
if (!("timestamp" in itemToCreate)) {
var ts = new Date();
itemToCreate["timestamp"] = ts.getTime();
}

// update the item that will be created
request.setBody(itemToCreate);
}
```

https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-write-stored-procedures-triggers-udfs

upvoted 9 times

ENGs Most Recent 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

jakobste 2 months, 2 weeks ago

This was on the exam in August 2023

upvoted 3 times

Sleepers 4 months ago

Got this on 06/27/2023 exam

upvoted 3 times

Devendra_Shukla 3 months, 3 weeks ago

which option you select for second dropdown?

upvoted 1 times

DonH 4 months, 2 weeks ago

Just for information: I just had this question on my AZ204 exam - 16-jun-2023.
I barely made it (with only 767 points) so I can't inform anyony if this answer is correct or not, just stating that this is an actual exam question.

upvoted 4 times

dirtuebus 5 months, 3 weeks ago

The right answer in the second dropdown, even though it has a little syntax error, has to be by exclusion:
if(isNaN(i)["tip"] || i["tip"] === null) {
The other options don't check if "tip" is a number, at all.

Anyway, even if there is a syntax error, I don't believe the exam is supposed to be a test of our ability to compile the code in our minds :)

upvoted 2 times

D58 5 months, 3 weeks ago

The goal is to set the tip property on the document in Cosmos DB. 3. question's answer should be __.replaceDocument(i) and not r.Body(i) since the document should be replaced (updated) and not only returned to the user (r.Body(i)). See link with examples here: https://learn.microsoft.com/en-us/training/modules/work-with-cosmos-db/4-cosmos-db-stored-procedures

upvoted 1 times

D58 5 months, 3 weeks ago

Oh never mind: Pretrigger can be used:
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-write-stored-procedures-triggers-udfs?tabs=javascript#pre-triggers

upvoted 2 times

fkaracan 8 months, 1 week ago

Question says "tip" should be numeric. Thus isNaN(tip) is correct

upvoted 2 times

dimsok 8 months, 3 weeks ago

The isNaN(i)["tip"] makes no sense as a syntax, isNaN returns true or false.
From the line i["tip"] = 0; we can assume i is an array and thus ("tip" in i) is correct

upvoted 2 times

saravanasanthosh 10 months ago

Got this on 12/30/2022 exam

upvoted 2 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 2 times

[Removed] 1 year, 1 month ago

Microsoft expect us to be polygots knowing Javascript, .Net, Powershell and Bash. Or maybe this exam is to prove that the person is full time busy doing hello-world deployments using different options in Azure. A decent developer uses Azure not try to remember only Azure. This is pointless.
The only thing to know is that if other webapps are not yet ready to use that new variable (tip) then it should be set in the trigger to be 0 if not provided (validations are not strictly required because it will fail if the tip is given wrongly as a string and assuming there are no business requirement to set boundaries)
The code is not required to check this knowledge. A decent developer knows the art of googling-copy-and-paste.

upvoted 3 times

macobuzi 2 months, 2 weeks ago

I understand your struggle, but Microsoft also doesn't expect us to study their questions bank to pass their exam. Just like a job interview, many interview questions might not relate to your actual work at all. Still, you need to pass it.

upvoted 1 times

r3verse 11 months ago

Hey man i keep seeing you post negative about every question in this exam. I do understand your points, but it doesn't add any value to the people trying to learn the questions. It actually decreseas the value to add to it. Your feedback about Microsoft in general can be posted at lots of places, but i don't think this is the best place. Please try to contribute in a way that helps people get the exam!

upvoted 7 times

NombreFalso 8 months, 2 weeks ago

I think he's pretty fun for comedic relief

upvoted 3 times

[Removed] 1 year, 1 month ago

I'm a developer and I get a headache looking at that "code". Microsoft should stop thinking these kind of code is the only way to test a "developer". Just look at all these discussions about this answer and what knowledge is it checking? (Memorize what the Microsoft guy wrote in some document as example?)

upvoted 6 times

Knightie 1 year, 2 months ago

isNaN(i)["tip"] <=== syntax error. isNaN(i) will return only True/False which is boolean, where ["tip"] is expecting the left side variable being an Array.

upvoted 2 times

UmarAbbas 1 year, 5 months ago

In my opinion, the answer would be A as Durable functions are made for such kind of long-running operations. My answer would be A.

upvoted 1 times

gursimran_s 1 year, 6 months ago

Came on April 20, 2023

upvoted 2 times

rgullini 7 months ago

Are you from the feature?

upvoted 2 times

Question #16 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data.
You need to ensure the app does not time out and processes the blob data.
Solution: Use the Durable Function async pattern to process the blob data.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:
Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the
HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices

Community vote distribution
A (82%) B (18%)

sasisang Highly Voted 2 years, 11 months ago
if someone knows the exact answer or any reference , then only comment. It is so confusing
upvoted 87 times

john4p 1 year, 10 months ago
A. Yes

"230 seconds is the maximum amount of time[...] For longer processing times, consider using the DURABLE FUNCTIONS ASYNC PATTERN[...]"
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
upvoted 45 times

Netspud 1 year, 8 months ago
See question 17, which would suggest the correct answer for this question is NO. The issues is the calling app is timing out not the function (i.e. the function in not replying as completed to the calling app in time, so use a queue, queue replies as completed in time, and work is deferred). At least after reading the two that is how I understand it. Very confusing question.
upvoted 4 times

Magis 11 months, 4 weeks ago
I agree with xRiot007. This question have two correct answers as MS saying that you can use Durable Function OR Service Bus in this scenario. So are correct.

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.
upvoted 1 times

xRiot007 1 year, 4 months ago
these types of questions can have multiple correct responses, and sometimes no correct responses at all. In this case, Durable Functions and using http triggers that can differ to a queue for long time processing are good answers.
upvoted 2 times

rgullini 6 months, 3 weeks ago
Somehow you need to investigate and understand the answers, not just memorize without understanding.
upvoted 2 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with most voted answer, score 926

upvoted 3 times

Awry Highly Voted 2 years, 12 months ago

Shouldn't this be TRUE?
https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-overview?tabs=csharp#async-http

upvoted 58 times

WillPassExam 2 years, 6 months ago

Agree :)
Durable function is one of the suggested two options for long processing tasks (the other is to defer the actual work (e.g. push to service bus queue + queue triggered function solution)
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

upvoted 28 times

sieiro 1 year, 11 months ago

I thing that maybe durable functions exceed de timeout of other functions, but if the functions fails or interrumpt don't cover the "The app must process the blob data".

upvoted 1 times

noro5 1 year, 8 months ago

True, I think the explanation mentions this
https://docs.microsoft.com/en-us/azure/architecture/patterns/async-request-reply#example
which looks exactly like the Async Pattern of Durable Function

upvoted 1 times

Secure01 2 years, 11 months ago

Seems you are right...
Durable Functions provides built-in support for this pattern(Pattern #3: Async HTTP APIs), simplifying or even removing the code you need to write to interact with long-running function executions.

upvoted 8 times

Vladimir_Gajinov Most Recent 1 month, 3 weeks ago

Selected Answer: A

The answer is YES.

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout

upvoted 1 times

Michael19842 3 months, 2 weeks ago

This site is not helping me at all... Every third awnser is wrong

upvoted 3 times

sarmaria 7 months, 2 weeks ago

Got this question in the exam on 16/03/2023. Chosen yes. Scored 846. Make sure to prepare for case studies. I got city and lights case study.

upvoted 3 times

daothanhyen 8 months, 1 week ago

Selected Answer: B

No.
Durable changes timeout settings. It can't resolve the problem.
Use queues to response immediately

upvoted 1 times

jimmyyml 8 months, 1 week ago

Yes, using the Durable Function async pattern can solve the problem of the app timing out and ensure that the blob data is processed. The Durable Function allows long-running functions to be executed in a more scalable and reliable way, by splitting the functions into smaller, asynchronous tasks that can be managed by the framework. This can help prevent timeouts by allowing the function to continue processing in the background, while also providing additional benefits such as state management and error handling.

upvoted 1 times

dimsok 8 months, 3 weeks ago

I wouldn't use durable functions, they are build to solve stateful problems, and in any case If the calling app is timing out, then durable functions will do nothing.

upvoted 3 times

Esward 9 months, 1 week ago

Yes, Is the correct answer
upvoted 1 times

Esward 9 months, 1 week ago
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
upvoted 1 times

narenazure 9 months, 3 weeks ago
Selected Answer: A
Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.
upvoted 1 times

hubekpeter 11 months, 2 weeks ago
Selected Answer: A
Answer is yes. You can find the same example here. https://learn.microsoft.com/en-us/azure/developer/javascript/how-to/with-web-app/azure-function-file-upload And here's the code change required to deal with HTTP API Async pattern. Basically the probably the problem is with a big blobs, when you are hitting the LB/HTTP timeout limits, which is 230s. With an async pattern, you get a webhook you can query instead.
upvoted 1 times

hubekpeter 11 months, 2 weeks ago
I'm curious, when we'll get an edit feature from examtopics to fix typos.
upvoted 1 times

TheExamMaster2020 11 months, 2 weeks ago
Did my exam on 15th November 2022. This question was on it.
upvoted 2 times

micro9000 11 months, 3 weeks ago
A or B can be an answer, it depends on your preference
The key here is "The app continues to time out after four minutes." and it seems it is expected for the processing to run more than four minutes, but based on this https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
For longer processing times, consider using the Durable Functions async pattern OR defer the actual work and return an immediate response.
upvoted 1 times

TheExamMaster2020 1 year ago
Selected Answer: A
There are more solutions. I quote here the description "Some question sets might have more than one correct solution, while others might not have a correct solution" so both Durable func & HTTP is right.
upvoted 3 times

[Removed] 1 year, 1 month ago
Http triggered function: okay
The app is triggered using an output binding on the blob: What? How can the app be triggered with an output binding? Or do you mean it get the blob data with an input binding after being triggered by an HttpTrigger? Assuming that.
YES: A durable function properly implemented can return immediately to the HttpTrigger with a 202 (Accepted), process the blob, and the http trigger can get the status
NO: Because it never says if the http trigger can check the status later. Even a queue in between cannot make it work.
Basically these kind of questions are typical of Microsoft. Does not make sense.
But I will vote for YES, considering Microsoft is proud of Durable functions (I do like it and used it)
upvoted 1 times

serpevi 1 year, 1 month ago
Got this in 09/22 , went with the most voted answers, score 927. This was the only I marked as YES.
upvoted 3 times

Ameet9 1 year, 1 month ago
Selected Answer: B
* Regardless of the Function App Timeout configuration, 230 seconds is the maximum amount of time that an HTTP Triggered function can take to respond to a request.
so the answer is "No"
upvoted 1 times

Question #17 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data.
You need to ensure the app does not time out and processes the blob data.
Solution: Pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*
Large, long-running functions can cause unexpected timeout issues. General best practices include:
Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the
HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices

Community vote distribution
A (64%) B (36%)

msdevanms Highly Voted 2 years, 11 months ago
Answer is correct
upvoted 61 times

piotrek1993 1 year, 7 months ago
What about sending files via Service Bus ? I guess this is not a good idea especially when maximum message size is 256 kb.
upvoted 3 times

hubekpeter 11 months, 2 weeks ago
You are not going to send the file throu message buss, only reference to a blob, 256kB is plenty enough for that purpose.
upvoted 5 times

tom112 5 months, 3 weeks ago
But it's a message, not an event. There are chances that a message would contain a blob in their body, max size 256k for standard and 100MB for premium.
upvoted 1 times

vruizrob 1 year, 1 month ago
This link https://docs.microsoft.com/en-us/azure/azure-functions/performance-reliability#avoid-long-running-functions confirms that the correct answer is A.
upvoted 5 times

surprise0011 6 months, 2 weeks ago
good finding. now it is clear.
IMO it not seems optimal but it is certainly a solution
upvoted 1 times

henry1985 1 year, 3 months ago
Async request-reply pattern https://docs.microsoft.com/en-us/azure/architecture/patterns/async-request-reply#example
upvoted 1 times

noro5 1 year, 8 months ago
Yes, the note section here confirms that
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
upvoted 1 times

Secure01 Highly Voted 2 years, 11 months ago

Answer is NO. The best solution is durable functions
Durable Functions provides built-in support for this pattern(Pattern #3: Async HTTP APIs), simplifying or even removing the code you need to write to interact with long-running function executions.
Durable Functions provides built-in support for this pattern(Pattern #3: Async HTTP APIs), simplifying or even removing the code you need to write to interact with long-running function executions.
https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-overview?tabs=csharp#async-http
upvoted 24 times

Figa 2 years, 11 months ago
Why should this Answer be false?

Can`t be both correct? They are not asking for the optimal solution.
upvoted 20 times

Cornholioz 2 years, 11 months ago
You may be right. There are questions where two answers are true. I think both Durable Timers and this solution are both true.
upvoted 10 times

Cornholioz 2 years, 11 months ago
But then again, can you really chunk the Blob data into multiples just so the refactored smaller function sets can process in parallel? What if the Blob data is just one huge chunk of a single transaction that has to be completed in shot even if takes long. That makes the Durable Functions the better/right/only answer :(
upvoted 1 times

tevivi8222 2 years, 7 months ago
It was timing out because you were doing a long task during HTTP request, which can't live forever. If you're saving a huge blob and it's not being done during the lifecycle of that request, the time doesn't really matter.
upvoted 1 times

rgullini 6 months, 3 weeks ago
Question is not asking for the BEST solution. Is asking if the provided solution will solve the issue or not, despite it is more expensive or not.
upvoted 6 times

Yah88 2 years, 7 months ago
you aren't asked for best solution, but for solution that works.
upvoted 19 times

TakumaK 2 years, 4 months ago
don't overthink. stick to what the question is asking. and why does this comment have the upvotes that many?
upvoted 19 times

Vladimir_Gajinov Most Recent 1 month, 3 weeks ago
Selected Answer: A
The answer is YES.

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
upvoted 2 times

nardk 1 month, 3 weeks ago
Selected Answer: A
Answer is A
upvoted 1 times

Kanasan 2 months ago
Selected Answer: A
https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.
upvoted 1 times

longnguyendh 3 months ago
A is correct.
upvoted 1 times

WuksaUK 3 months, 2 weeks ago
Answer is correct, the timeout time for HTTP triggers is always 230s. Since the compute time is 240s the minimum timeout for a Standard azure function that isn't from a HTTP trigger is 5 minutes. So this works
upvoted 3 times

deathRac3 5 months, 1 week ago
Selected Answer: A
For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

upvoted 2 times

tom112 5 months, 3 weeks ago

Selected Answer: B

"return an immediate HTTP success response." part makes the answer No.
Max size of a message is 100mb, we don't know how long it takes to process, and the condition is saying that the app times out after 4 min. then you cannot process through servicebus queue and give immediate response in time.

upvoted 1 times

tom112 5 months, 3 weeks ago

Soz, I was stupid. the answer is A.

upvoted 1 times

lyggwtwtczxnhznebw 6 months ago

This solution may not meet the goal of processing the blob data within the Azure Function app without timing out.

Passing the HTTP trigger payload to a Service Bus queue will remove the processing logic from the original Azure Function app, and instead, pass the message to a separate Azure Function that is triggered by the Service Bus queue.

While this approach can be useful for decoupling processing logic and scaling the processing across multiple functions, it does not directly address the timeout issue in the original function. The original function will still timeout after four minutes, and the Service Bus queue-triggered function will also be subject to a default timeout limit of five minutes.

Therefore, while passing the HTTP trigger payload to an Azure Service Bus queue and processing it with a queue trigger function may be a useful approach for certain scenarios, it may not meet the goal of ensuring that the original HTTP-triggered Azure Function app processes the blob data without timing out.

upvoted 1 times

crica007 8 months ago

Passing the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function would not necessarily prevent the Azure Function from timing out after four minutes. This solution would introduce additional complexity and may not be necessary depending on the specific requirements of the application. It would be better to investigate options such as increasing the function timeout limit, breaking up the processing into smaller chunks, or using Durable Functions to handle long-running operations. B.No

upvoted 3 times

Michael2023 9 months, 1 week ago

the wording is somehow really confusing here, because I have an HTTP triggered function app... and also the app is triggered using an output binding on the blob... so what is meant here exactly? ...
also the question describtion says, that there can be more than 1 right solutions

upvoted 1 times

rotimislaw 10 months ago

Selected Answer: A

Comparing with previous answer, I expect one of them to be TRUE and the other FALSE

upvoted 2 times

rgullini 6 months, 3 weeks ago

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.

Please read carefully to avoid wrong asumptions

upvoted 1 times

hubekpeter 11 months, 2 weeks ago

Selected Answer: B

You need to ensure the app does not time out and processes the blob data. I think this is the key. Yes it's technically doable but if you send immediate response, there's no guarantee that a function will really process the data. Yes there are retry policies etc., but relying on something like that is considered as anti pattern from architecture standpoint.

upvoted 1 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 1 times

Magis 11 months, 4 weeks ago

Selected Answer: A

MS saying that you can select Durable function OR service bus.

So in this YES/NO question series we have two correct answers not just a one as usually.

This one is correct.

upvoted 3 times

TheExamMaster2020 1 year ago

Selected Answer: A

The maximum amount of time that an HTTP triggered function can take to respond to a request is 230 seconds. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern.
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout

upvoted 1 times

Question #18 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data.
You need to ensure the app does not time out and processes the blob data.
Solution: Configure the app to use an App Service hosting plan and enable the Always On setting.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:
Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the
HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices

Community vote distribution
B (80%) A (20%)

00avatar Highly Voted 2 years, 9 months ago
Answer "No" is correct. Always On enables waking up on HTTPTrigger, but does not prevent the exceeding the max time out time of 230 seconds.
https://docs.microsoft.com/en-us/azure/azure-functions/dedicated-plan#always-on
upvoted 61 times

surprise0011 6 months, 2 weeks ago
received 2023-04-17 went given answer, score 926
upvoted 2 times

Knightie 1 year, 2 months ago
No is the answer, but the reason is the timeout being raised by HTTP layer from the Azure Load Balancer, not the App layer that at least it gives 5 minutes for the cheapest type, Consumption, so however you enhance the app layer, the http layer Azure Load Balance will not wait more than 230 second and will reply it as timeout. Use the durable function pattern to poll the status for completion will be the easiest solution, else avoid the http layer like service bus will work too.

https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout:~:text=1%20Regardless%20of,immediate%20response.
upvoted 5 times

Amrit862 2 years, 8 months ago
Answer is still 'No' but default time for dedicated host is 30 mins, it can be technically configured to infinite...but have to do that manually and that is not mentioned in solution.
ref: https://docs.microsoft.com/en-us/azure/azure-functions/functions-host-json#functiontimeout
upvoted 1 times

Basu525 2 years, 8 months ago
Yes absolutely, as per Microsoft documentation, "Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer." So App service plan with Always On is not a solution.
upvoted 4 times

clarionprogrammer 2 years, 6 months ago
Agreed. 'No' is correct.
Here is the reference cited above.
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
upvoted 5 times

Cornholioz Highly Voted 2 years, 11 months ago

Answer "B.No" seems right because this is not about App Services Best Practices.

upvoted 9 times

Kanasan Most Recent 2 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer.

upvoted 1 times

RenatoJr 6 months ago

Selected Answer: B

Answer "No" is correct

upvoted 1 times

Anitay 1 year, 1 month ago

Use Durable function for this requirement

upvoted 1 times

Perplex 1 year, 2 months ago

Selected Answer: B

Always On does not prevent a timeout so is not a solution to this problem. Answer is B (No).

upvoted 1 times

Rakeshpro 1 year, 3 months ago

Selected Answer: A

Correct!

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer (No).
Max timeout for HTTP triggers is 230 seconds < 240 seconds (4 minutes).

upvoted 1 times

TakumaK 2 years, 4 months ago

"Always On" feature of Azure App Service is to keep the host process running to allow more responsive to requests after significant idle periods.
This is UNRELATED to the timeout.
So the answer is obvious!!

upvoted 4 times

mlantonis 2 years, 5 months ago

Correct Answer: No

Always On enables waking up on HTTP trigger, but does not prevent the exceeding the max time out time of 230 seconds.

If you run on an App Service plan, you should enable the Always on setting so that your function app runs correctly. On an App Service plan, the functions runtime goes idle after a few minutes of inactivity, so only HTTP triggers will "wake up" your functions. The Always on setting is available only on an App Service plan. On a Consumption plan, the platform activates function apps automatically.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/dedicated-plan#always-on

upvoted 8 times

abdou1987 2 years, 5 months ago

Answer is YES
"Best for long-running scenarios where Durable Functions can't be used. Consider an App Service plan in the following situations:" link:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#overview-of-plans
another link https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout that mentionned the duration is 30 min and upper

upvoted 1 times

abdou1987 2 years, 5 months ago

Sorry i cant delete my previois comment.
I confirm the answer is NO. Whatever the app service plan the maximum timeout for HTTP triggers is 230 seconds.

upvoted 3 times

TakumaK 2 years, 4 months ago

you are messing up this comment section by correcting your own wrong answer. :)

upvoted 1 times

UnknowMan 2 years, 5 months ago

Answer "No" is correct

upvoted 2 times

glam 2 years, 5 months ago

B. No,,

upvoted 3 times

bugimachi 2 years, 10 months ago

The page, however, does not talk about Service Bus queues (any more).

upvoted 1 times

Question #19 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure
Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Move photo processing to an Azure Function triggered from the blob upload.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*
Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file- oriented workflow.
Events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener.
Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

Community vote distribution
B (50%) A (50%)

AnkanG Highly Voted 3 years, 3 months ago
as is correct
upvoted 56 times

fadikh Highly Voted 2 years, 8 months ago
The answer is correct, but it should have also mentioned that the function app must not be on a consumption plan. because in that case, it might take up to 10 minutes to process the event.
upvoted 27 times

dookebroyle 1 year, 4 months ago
The time limit is for the process to START only. It can take longer to process.
upvoted 4 times

bbou 2 years, 6 months ago
No 10 seconds
upvoted 1 times

bbou 2 years, 6 months ago
I was wrong - yes up to 10 minutes
upvoted 4 times

BrettusMaximus 2 years, 6 months ago
So the answer is NO. Cant wait 10 minutes
upvoted 6 times

altafpatel1984 1 year, 11 months ago
That limitation is only for Consumption plan. Since no mention of it, we can assume user can opt for premium, app service plan.
upvoted 5 times

dddddd111 Most Recent 1 day, 3 hours ago
Selected Answer: A
I think the answer is A.
upvoted 1 times

Dixavado 1 month ago

Selected Answer: A

Go with A

upvoted 1 times

Kanasan 2 months ago

> The process to produce a mobile-friendly version of the image must start in less than one minute.

This cannot be fulfilled with the proposed solution.
So, vote for B

upvoted 1 times

LxstInNxght 2 months, 3 weeks ago

Selected Answer: A

Answer is A

upvoted 1 times

EliteAllen 3 months ago

Selected Answer: A

A. Yes

The solution does meet the goal. Azure Functions can be triggered by blob storage events, such as the uploading of a new photo. When a new photo is uploaded to blob storage, the Azure Function would be triggered to process the photo and create a mobile-friendly version. The trigger is almost instantaneous, so the process would indeed start in less than a minute after a photo is uploaded. This meets the requirement mentioned in the question.

upvoted 1 times

Sleepers 4 months ago

Got this on 06/27/2023 exam

upvoted 1 times

Priyankakanna 4 months, 1 week ago

Answer A, just to vote

upvoted 1 times

deathRac3 5 months, 1 week ago

Selected Answer: A

just to vote

upvoted 1 times

surprise0011 6 months, 2 weeks ago

In conjuction with Event Grid it will work, but without it is not guaranteed to start within 1 minute.
So, since Event Grid is not mentioned I would say: No

upvoted 5 times

bdumois 7 months, 1 week ago

Correct answer is B:

There is no service level agreement around the time it takes for a message to arrive. It's not uncommon for messages to arrive anywhere from 30 minutes to two hours.
https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

upvoted 6 times

adilkhan 7 months, 2 weeks ago

YES CHATGPT

upvoted 1 times

red0400 7 months, 2 weeks ago

Selected Answer: A

A is correct.

upvoted 1 times

fkaracan 8 months, 1 week ago

Selected Answer: B

I would say "No" because the question says "must" be less than 10 minutes. The latency with the Azure blob storage trigger is up to 10 minutes.

https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=in-process&pivots=programming-language-csharp

upvoted 5 times

szumik 8 months, 1 week ago

I had this set of questions on exam on 22/02/23, it consisted of solutions:
1. Function on Consumption plan
2. Implement change feed
3. Change to BlockBlob
I selected No on all of them, not sure if correct, but scored overall 846

upvoted 5 times

AymanAkk 10 months, 1 week ago

Selected Answer: A

A is correct

upvoted 1 times

Question #20 *Topic 2*

You are developing an application that uses Azure Blob storage.
The application must read the transaction logs of all the changes that occur to the blobs and the blob metadata in the storage account for auditing purposes. The changes must be in the order in which they occurred, include only create, update, delete, and copy operations and be retained for compliance reasons.
You need to process the transaction logs asynchronously.
What should you do?

A. Process all Azure Blob storage events by using Azure Event Grid with a subscriber Azure Function app.

B. Enable the change feed on the storage account and process all changes for available events.

C. Process all Azure Storage Analytics logs for successful blob events.

D. Use the Azure Monitor HTTP Data Collector API and scan the request body for successful blob events.

Correct Answer: *B*
Change feed support in Azure Blob Storage
The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes. Client applications can read these logs at any time, either in streaming or in batch mode. The change feed enables you to build efficient and scalable solutions that process change events that occur in your Blob Storage account at a low cost.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed

Community vote distribution
B (100%)

Ummara Highly Voted 3 years, 3 months ago
B: the change feed provides transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, a durable, immutable, read-only log of these changes. You can process these logs asynchronously, incrementally or in-full.
upvoted 71 times

azurelearner666 2 years, 4 months ago
Right!
upvoted 3 times

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: B

The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes. Client applications can read these logs at any time, either in streaming or in batch mode. The change feed enables you to build efficient and scalable solutions that process change events that occur in your Blob Storage account at a low cost.

Reference:

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed
upvoted 18 times

deathRac3 Most Recent 5 months, 1 week ago
Selected Answer: B
as is correct
upvoted 2 times

Saluk_DE 7 months ago
Question was in Exam 2023-03-30
upvoted 4 times

NombreFalso 8 months, 2 weeks ago
Selected Answer: B
Exactly what Change Feed offers
upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. Enable the change feed on the storage account and process all changes for available events.
Azure Blob storage change feed provides a log of all the create, update, delete, and copy operations that occur on blobs and blob metadata in the storage account. It allows you to track the changes to the blobs in the order in which they occurred, which is what you need for auditing purposes. By enabling the change feed, your application can asynchronously process the changes and retain them for compliance reasons. You can also use Azure Event Grid to route the change feed events to a subscriber Azure Function app for further processing, this way you can keep the logs for the compliance reasons.

upvoted 1 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 2 times

AZAdam22 1 year, 3 months ago

Selected Answer: B

B - Because change feed contains an ordered list of operations that have been made in a storage account.

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.

upvoted 1 times

hamzabts 1 year, 6 months ago

B : The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account.

upvoted 1 times

Freidrich 1 year, 8 months ago

Selected Answer: B

Correct.

upvoted 1 times

HimanshuNankani 2 years ago

What about the delete operations that are not logged by the change feed?

upvoted 6 times

santoshsidnal 1 year, 4 months ago

That's a valid question!! I think we should use event grid approach i.e A

upvoted 1 times

[Removed] 1 year, 1 month ago

BlobDeleted is supported eventType and even if it is not supported still the requirement that it should be ordered cannot be satisfier with the blob-storage-events through event-grid. Change feed supports ordering

upvoted 1 times

ozbonny 2 years, 3 months ago

enabled change feed to get information and for auditing or compliance purpose

upvoted 1 times

UnknowMan 2 years, 5 months ago

Correct B

upvoted 3 times

glam 2 years, 5 months ago

B. Enable the change feed on the storage account and process all changes for available events.

upvoted 1 times

siddharth 2 years, 6 months ago

B is correct

upvoted 1 times

jokergester 2 years, 7 months ago

For those who are confused with Azure Storage Analytics logging, read below. It also contains information about "Change Feed"

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal#what-is-the-difference-between-change-feed-and-storage-analytics-logging

upvoted 8 times

Question #21 *Topic 2*

DRAG DROP -

You plan to create a Docker image that runs an ASP.NET Core application named ContosoApp. You have a setup script named setupScript.ps1 and a series of application files including ContosoApp.dll.

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

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

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

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

Select and Place:

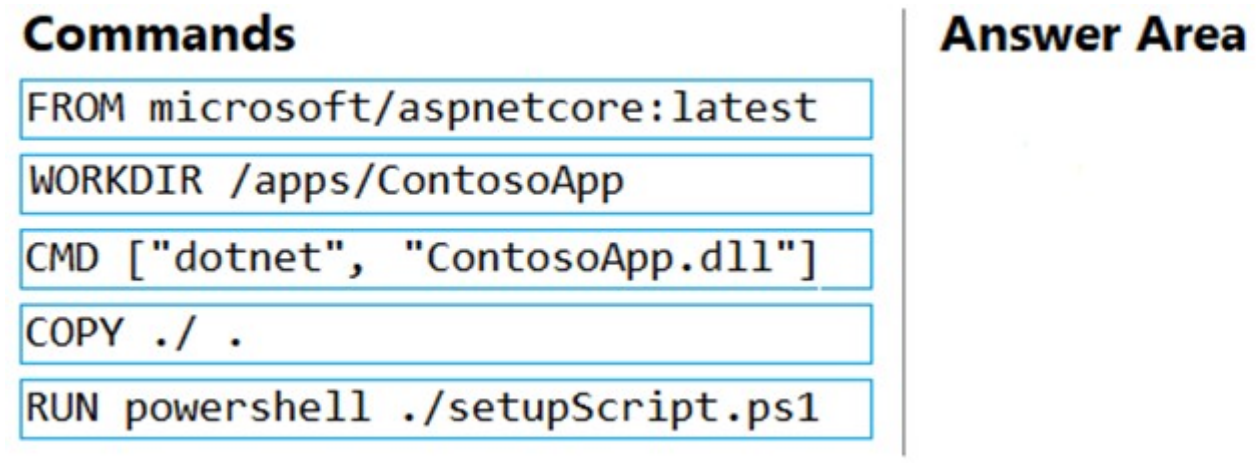

Correct Answer:

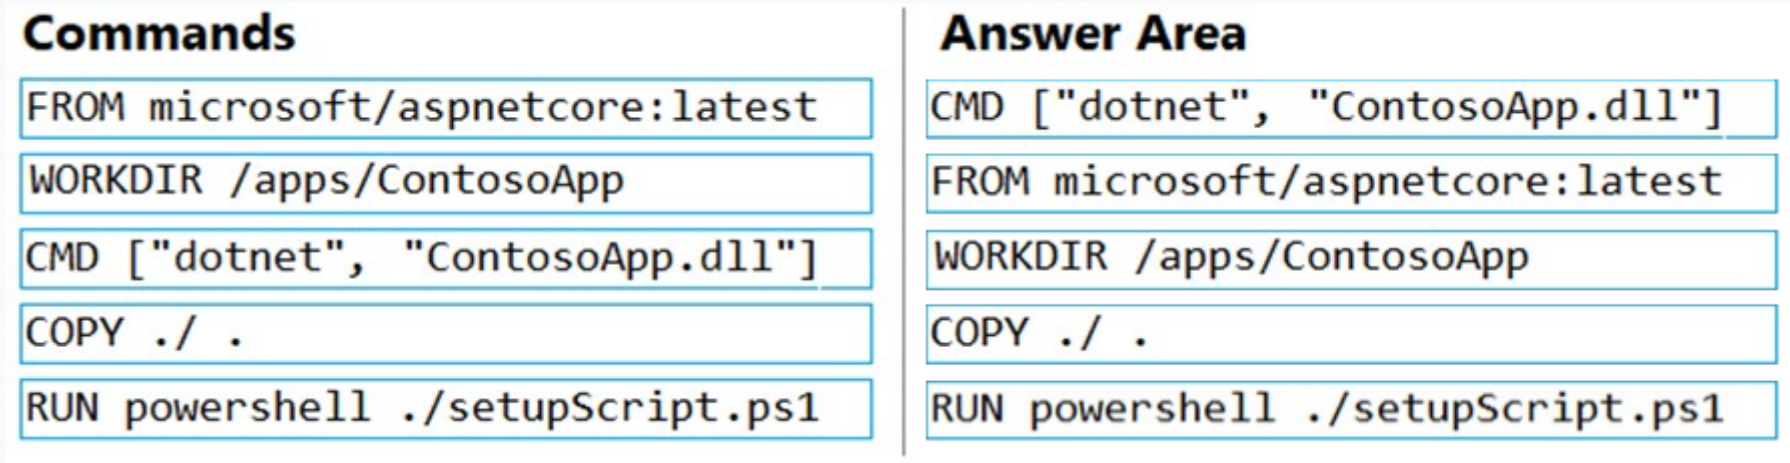

Box 1: CMD [..]

Cmd starts a new instance of the command interpreter, Cmd.exe.

Syntax: CMD <string>

Specifies the command you want to carry out.

Box 2: FROM microsoft/aspnetcore-build:latest

Box 3: WORKDIR /apps/ContosoApp -

Bxo 4: COPY ./ .

Box 5: RUN powershell ./setupScript.ps1

agueda Highly Voted 2 years, 7 months ago

It should be:
- FROM
- WORKDIR
- COPY
- RUN
- CMD

Same question on:
https://www.examtopics.com/discussions/microsoft/view/13131-exam-az-300-topic-3-question-4-discussion/

And:
https://www.examtopics.com/discussions/microsoft/view/11045-exam-az-203-topic-1-question-7-discussion/

upvoted 268 times

ReyPirata 2 months, 1 week ago

This was on the exam (08/20/2023). Scored 925

- FROM
- WORKDIR

- COPY
- RUN
- CMD

upvoted 2 times

balis 1 year, 7 months ago

This is correct answer
It should be:
- FROM
- WORKDIR
- COPY
- RUN
- CMD

because WORKDIR will create directory if it doesn't exist https://docs.docker.com/engine/reference/builder/#workdir

upvoted 12 times

Dinima 2 years, 7 months ago

You are correct. This has been discussed in Udemy course as well as follows,
The first statement in the Dockefile must be the FROM statement to specify the image to use as the base image.

Then specify the Image working directory

Then copy all of the application contents using the COPY command

And then use the CMD command to run the PowerShell command and the ENTRYPOINT statement to run the dotnet application.

upvoted 29 times

Basu525 2 years, 7 months ago

as per Udemy, the last steps would be CMD powershell ./script.ps1 and then ENTRYPOINT (dotnet, xx.dll) which I believe is the correct answer. But unfortunately the options are not there in Examtopic

upvoted 3 times

ranjitklive 2 years, 3 months ago

ENTRYPOINT instruction works very similarly to CMD in that it is used to specify the command executed when the container is started.

upvoted 3 times

solidrock 2 years ago

which udemy course you guys are talking about?

upvoted 2 times

Den1354 1 year, 9 months ago

- FROM
- COPY
- WORKDIR
- RUN
- CMD

Otherwise we are going to set for work directory path which doesn't exist yet

upvoted 4 times

TakumaK Highly Voted 2 years, 4 months ago

Just wondering who put the answers for the questions in this site? most of them are not correct.

upvoted 22 times

SlavMar 2 years, 4 months ago

VCE exams have same issues
Probably this are some braindumps of people who were not well prepared to take test or they too test just to scrap exam questions but they have no matter knowledge

upvoted 3 times

p2006 Most Recent 1 month ago

Got on 9/25/2023
From
Workdir
Copy
Run powershell
Cmd ContosoApp.dll

upvoted 1 times

RuffBoii 1 month, 1 week ago

Had this on my exam today.

upvoted 1 times

mihailos 2 months ago

Got it in exam 28/08/23. Scored 912. Went with the following:
- FROM
- WORKDIR
- COPY

- RUN
- CMD

upvoted 1 times

TechyNetty 1 month, 3 weeks ago

Im yet to appear for exam soon. Since you attempted recently, were most of the questions from this site? Please respond, it will be helpful. thanks

upvoted 1 times

longnguyendh 3 months ago

■ FROM ...
■ WORKDIR ...
■ COPY ...
■ RUN ...
■ CMD ...

upvoted 1 times

deathRac3 5 months ago

Question was there for me on 29th May 2023

upvoted 1 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23
score: 817

upvoted 1 times

Priya0703 8 months, 1 week ago

Got this question today on 20-02-2023 exam.

upvoted 1 times

HafizSalmanMalik 9 months, 2 weeks ago

It should be:
- FROM
- WORKDIR
- COPY
- RUN
- CMD

upvoted 3 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 3 times

ghuna 1 year, 2 months ago

Answer seems to be wrongly documented, it should be like this.
- FROM
- WORKDIR
- COPY
- RUN
- CMD

upvoted 3 times

samraw83 1 year, 3 months ago

For any doubt between RUN and CMD:
The requirement is something like this as per the question:
☞ Call setupScripts.ps1 when the container is built. (so definitely RUN comes first)
☞ Run ContosoApp.dll when the container starts. (Therefore CMD comes next)
Rest of the above are quite clear for most of the users:
FROM
WORKDIR
COPY
RUN
CMD

upvoted 3 times

angrybird2007 1 year, 3 months ago

https://docs.docker.com/samples/dotnetcore/

```
# syntax=docker/dockerfile:1
FROM mcr.microsoft.com/dotnet/sdk:6.0 AS build-env
WORKDIR /app

# Copy csproj and restore as distinct layers
COPY *.csproj ./
RUN dotnet restore

# Copy everything else and build
COPY ../engine/examples ./
```

```
RUN dotnet publish -c Release -o out

# Build runtime image
FROM mcr.microsoft.com/dotnet/aspnet:6.0
WORKDIR /app
COPY --from=build-env /app/out .
ENTRYPOINT ["dotnet", "aspnetapp.dll"]
```

upvoted 3 times

Pize 1 year, 4 months ago

do they deliberately give you the wrong answers at this website to prompt discussion?

upvoted 3 times

prachi190892 1 year, 4 months ago

The Dockerfile will have below series of commands:
FROM
WORKDIR
COPY
RUN: command that runs while the images builds
CMD: command that runs when the container is created

upvoted 1 times

Eltooth 1 year, 4 months ago

- FROM
- WORKDIR
- COPY
- RUN
- CMD

upvoted 1 times

Question #22 *Topic 2*

You are developing an Azure Function App that processes images that are uploaded to an Azure Blob container.
Images must be processed as quickly as possible after they are uploaded, and the solution must minimize latency. You create code to process images when the
Function App is triggered.
You need to configure the Function App.
What should you do?

A. Use an App Service plan. Configure the Function App to use an Azure Blob Storage input trigger.

B. Use a Consumption plan. Configure the Function App to use an Azure Blob Storage trigger.

C. Use a Consumption plan. Configure the Function App to use a Timer trigger.

D. Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger.

E. Use a Consumption plan. Configure the Function App to use an Azure Blob Storage input trigger.

Correct Answer: *B*
The Blob storage trigger starts a function when a new or updated blob is detected. The blob contents are provided as input to the function.
The Consumption plan limits a function app on one virtual machine (VM) to 1.5 GB of memory.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger

Community vote distribution
D (100%)

Kitkit Highly Voted 2 years, 8 months ago
The answer is D. Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger.
Consumption plan can cause a 10-min delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled.
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp
upvoted 188 times

OscarL 4 months ago
totally agree
upvoted 1 times

Dinima 2 years, 7 months ago
True, D, We have to use an Azure Blob storage trigger. In order to ensure the function is invoked immediately , m
upvoted 8 times

SlavMar 2 years, 4 months ago
Why not B then.
Using input for function apps seems to make more sense
upvoted 2 times

zzt 2 years, 3 months ago
Seemingly there is no such thing as "input trigger" it is "input binding"
See: https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-input?tabs=csharp
upvoted 31 times

SlavMar 2 years, 4 months ago
I ment A
upvoted 1 times

samraw83 1 year, 4 months ago
App service plan is best for the long running scenarios where you cannot use durable functions.
upvoted 2 times

danielcr Highly Voted 2 years, 8 months ago
A or D, never B.
With Consumption plan, you could have cold start, and the question say "must be processed as quickly as possible" so you need an App Service Plan.
Between A and D ... All triggers are Input, so i don´t undertand the difference.
upvoted 24 times

ewertonews 2 years, 1 month ago

It cannot be A. There is no such thing as a input trigger. it's either a (Blob Storage) trigger or an input binding. They are different things.
upvoted 9 times

surprise0011 6 months, 2 weeks ago
exactly my friend. I was trying to find info about it and as I was thinking, there is no such thing
upvoted 1 times

ranjitklive 2 years, 3 months ago
Agree with your point, all triggers are input..!!
upvoted 3 times

ewertonews 2 years, 1 month ago
But that nomenclature is not used by MS.
A function can have a trigger and a input binding (HTTP trigger with a Storage Queue as input binding for example).
upvoted 3 times

RuffBoii Most Recent 1 month, 1 week ago
Had this on my exam today.
upvoted 1 times

Guurr40__ 1 month, 1 week ago
What is the difference between "storage trigger" and "storage input trigger"? In the end, isn't the second one just a reinforcement to indicate that the moment a file is placed, precisely as input, in the hypothetical folder, then it is picked up and processed by the trigger. Thank you in advance for your answers
upvoted 1 times

sachinrikhe 7 months, 2 weeks ago
Answer is B.
When using a Consumption plan, Azure automatically scales out instances of the function as needed to handle incoming requests, and you only pay for the actual number of executions of the function. This ensures that the function is always available to process images without any delays caused by having to scale up or down manually.
upvoted 3 times

macobuzi 2 months, 2 weeks ago
You can't use Consumption Plan. Consumption Plan can take up to 10 mins to warm up which doesn't fit the requirement.
upvoted 1 times

nedlo 7 months, 2 weeks ago
Selected Answer: D
D is correct "For this scenario where images must be processed as quickly as possible and latency must be minimized, it is recommended to use the App Service plan instead of the Consumption plan.

The App Service plan provides dedicated resources to the function app, which can help improve performance and minimize the risk of cold start delays. Additionally, the App Service plan allows you to scale out the function app to multiple instances, which can help handle high traffic loads and further improve performance.

On the other hand, the Consumption plan is a serverless hosting option that scales automatically based on demand, but may result in cold start delays and limited resources during high traffic periods.

Therefore, it is recommended to use the App Service plan for this scenario."
upvoted 3 times

vguimars 8 months, 1 week ago
Selected Answer: D
Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger.
Consumption plan can cause a 10-min delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled.
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp
upvoted 1 times

OPT_001122 10 months, 4 weeks ago
Selected Answer: D
D. Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger.
upvoted 1 times

madhubanti0007 11 months, 4 weeks ago
Selected Answer: D
Consumption plan can cause a 10-min delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled.
upvoted 2 times

ms_master 12 months ago
Got this question on 30-Oct-2022 exam.
Answer is D. Passed with 870 score
upvoted 4 times

OPT_001122 10 months, 4 weeks ago

thanks for mentioning the date

upvoted 2 times

Ameet9 1 year, 1 month ago

Selected Answer: D

As quickly as possible then ofcourse it is App service plan

upvoted 2 times

ghuna 1 year, 2 months ago

Selected Answer: D

if you want to be fast, then you need to avoid cold starts, which are part of consumption plan

upvoted 1 times

rohitpatil113 1 year, 2 months ago

Selected Answer: D

D is correct as consumption plan will not process the data immediately.

upvoted 1 times

Pize 1 year, 4 months ago

Another wrong answer on this site

upvoted 2 times

xRiot007 1 year, 4 months ago

Solution is D : App service plan with Blob storage.
Answers with Consumption plan do not satisfy the requirement to minimize latency.
Also, there is no such thing as an input trigger.

upvoted 2 times

Hornwood509 1 year, 6 months ago

Selected Answer: D

From the URL: Minimizing latency: If your function app is on the Consumption plan, there can be up to a 10-minute delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled

upvoted 2 times

noobpinkfloyd 1 year, 7 months ago

Selected Answer: D

Consumption plan can have up to 10 minutes of cold start, and the 'input trigger' does not exist, there are triggers and input bindings.

upvoted 1 times

Question #23 *Topic 2*

HOTSPOT -

You are configuring a new development environment for a Java application.

The environment requires a Virtual Machine Scale Set (VMSS), several storage accounts, and networking components.

The VMSS must not be created until the storage accounts have been successfully created and an associated load balancer and virtual network is configured.

How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
{
. . .
"resources": [
  {
    "apiVersion": "2016-01-01",
    "type": "Microsoft.Storage/storageAccounts",
    "name": "[concat(                    (), 'storage', uniqueString(resourceGroup().id))]",
  "location": "[resourceGroup().location]",
  . . .
    "sku": {
    "name": "Standard_LRS"
  },
  "kind": "Storage",
  "properties": {},
  "                    ": {
    "name": "storagesetup",
    "count": 3
  }
},
{
  "apiVersion": "2015-06-15",
  "type": "Microsoft.Compute/virtualMachines",
  "name": "[concat('VM', uniqueString(resourceGroup().id))]",
  "                    ": [
      "[variables('loadBalancerName')]",
      "[variables('virtualNetworkName')]",
      "storagesetup",
      ],
    . . .
  }
],
"outputs": {}
}
```

- copy
- copyIndex
- priority
- dependsOn

Correct Answer:

Answer Area

```
{
. . .
"resources": [
  {
    "apiVersion": "2016-01-01",
    "type": "Microsoft.Storage/storageAccounts",
    "name": "[concat(                    (), 'storage', uniqueString(resourceGroup().id))]",
     "location": "[resourceGroup().location]",
     . . .
       "sku": {
       "name": "Standard_LRS"
     },
     "kind": "Storage",
     "properties": {},
     "                    ": {
       "name": "storagesetup",
       "count": 3
     }
  },
  {
     "apiVersion": "2015-06-15",
     "type": "Microsoft.Compute/virtualMachines",
     "name": "[concat('VM', uniqueString(resourceGroup().id))]",
     "                    ": [
       "[variables('loadBalancerName')]",
       "[variables('virtualNetworkName')]",
       "storagesetup",
       ],
     . . .
   }
],
"outputs": {}
}
```

copy / copyIndex / priority / dependsOn

Box 1: copyIndex -

Notice that the name of each resource includes the copyIndex() function, which returns the current iteration in the loop. copyIndex() is zero-based.

Box 2: copy -

By adding the copy element to the resources section of your template, you can dynamically set the number of resources to deploy.

Box 3: dependsOn -

Example:

"type": "Microsoft.Compute/virtualMachineScaleSets",
"apiVersion": "2020-06-01",
"name": "[variables('namingInfix')]",
"location": "[parameters('location')]",
"sku": {
"name": "[parameters('vmSku')]",
"tier": "Standard",
"capacity": "[parameters('instanceCount')]"
},
"dependsOn": [
"[resourceId('Microsoft.Network/loadBalancers', variables('loadBalancerName'))]",

"[resourceId('Microsoft.Network/virtualNetworks', variables('virtualNetworkName'))]"
],
Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/copy-resources https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/quick-create-template-windows

cloudlabadm Highly Voted 2 years, 7 months ago

Answer: copyIndex, copy, Dependson. Check this link https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/define-resource-dependency.

upvoted 65 times

Esward 9 months, 1 week ago

Agreed

upvoted 3 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: copyIndex
Notice that the name of each resource includes the copyIndex() function, which returns the current iteration in the loop. copyIndex() is zero-based.

Box 2: copy
By adding copy loop to the resources section of your template, you can dynamically set the number of resources to deploy. You also avoid having to repeat template syntax.

Box 3: dependsOn
Within your Azure Resource Manager template (ARM template), the dependsOn element enables you to define one resource as a dependent on one or more resources.

Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/copy-resources
https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/quick-create-template-windows
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/define-resource-dependency

upvoted 52 times

MrAzz Most Recent 4 days, 20 hours ago

For the second part: "Copy"
What is the point of adding the same property with the same name 3 times?

upvoted 1 times

deathRac3 5 months ago

Question was there for me on 29th May 2023

upvoted 3 times

CODE_STS 8 months ago

Got this in the exam today! Feb 28, 2023

upvoted 3 times

uffuchsi 8 months, 1 week ago

Received this in my exam today (22/02/2023). Selected copyIndex, copy, and dependsOn. Score 927.

upvoted 8 times

saravanasanthosh 10 months ago

Got this on exam 12/30/2022

upvoted 2 times

OPT_001122 10 months, 4 weeks ago

Thanks to all who have mentioned the exam date

upvoted 1 times

[Removed] 1 year, 1 month ago

Do I need to remember these functions? The options are made in a way to confuse you and make a mistake. These are things developers just lookup in documentation. Not writing an ARM template every day

upvoted 7 times

Rini100 1 year, 6 months ago

Got this on 20 Apr 2022.. (copyIndex, copy, dependsOn)

upvoted 5 times

AZ204Cert 1 year, 6 months ago

Got this on 04/05/22 (selected copyIndex, copy, dependsOn)

upvoted 4 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 3 times

AidenYoukhana 1 year, 10 months ago

CORRECT ANSWER.

upvoted 2 times

Rev1201 2 years, 1 month ago

Answer is Correct!

upvoted 3 times

glam 2 years, 5 months ago

Answer: copyIndex, copy, Dependson

upvoted 3 times

[Removed] 2 years, 7 months ago

https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/template-tutorial-create-multiple-instances?tabs=CLI%2Cazure-cli

upvoted 5 times

agueda 2 years, 7 months ago

Agree with the answer

upvoted 9 times

Question #24 *Topic 2*

HOTSPOT -

You are developing an Azure Function App by using Visual Studio. The app will process orders input by an Azure Web App. The web app places the order information into Azure Queue Storage.

You need to review the Azure Function App code shown below.

```
public static class OrderProcessor
{
  [FunctionName("ProcessOrders")]
  public static void ProcessOrders([QueueTrigger("incoming-orders")]CloudQueueMessage myQueueItem, [Table("Orders")]ICollector<Order> tableBindings, TraceWriter log)
  {
    log.Info($"Processing Order: {myQueueItem.Id}");
    log.Info($"Queue Insertion Time: {myQueueItem.InsertionTime}");
    log.Info($"Queue Expiration Time: {myQueueItem.ExpirationTime}");
    tableBindings.Add(JsonConvert.DeserializeObject<Order>(myQueueItem.AsString));
  }
  [FunctionName("ProcessOrders-Poison")]
  public static void ProcessFailedOrders([QueueTrigger("incoming-orders-poison")]CloudQueueMessage myQueueItem, TraceWriter log)
  {
    log.Error($"Failed to process order: {myQueueItem.AsString}");
    . . .
  }
}
```

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

	Yes	No
The code will log the time that the order was processed from the queue.	○	○
When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.	○	○
When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders.	○	○
The ProcessOrders function will output the order to an Orders table in Azure Table Storage.	○	○

Correct Answer:

Answer Area

	Yes	No
The code will log the time that the order was processed from the queue.	○	○
When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.	○	○
When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders.	○	○
The ProcessOrders function will output the order to an Orders table in Azure Table Storage.	○	○

Box 1: No -

ExpirationTime - The time that the message expires.

InsertionTime - The time that the message was added to the queue.

Box 2: Yes -

maxDequeueCount - The number of times to try processing a message before moving it to the poison queue. Default value is 5.

Box 3: Yes -

When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them.

By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.

Box 4: Yes -

Reference:

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

mlantonis Highly Voted 2 years, 5 months ago

Box 1: No
It logs the following:
- ExpirationTime - The time that the message expires.
- InsertionTime - The time that the message was added to the queue.

Box 2: Yes
maxDequeueCount: The number of times to try processing a message before moving it to the poison queue. Default value is 5.

Box 3: Yes
When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them. By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.

Box 4: Yes
[Table("Orders")]ICollector<Order> table bindings
And in the code it adds the order:
tableBindings.Add(JsonConvert.DeserializeObject<Object>(myQueueItem.AsString));

upvoted 122 times

NombreFalso 8 months, 2 weeks ago

Egg salad I mean excellent

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

very good and helpful explanation

upvoted 2 times

[Removed] Highly Voted 1 year, 1 month ago

Microsoft thinks that a good developer should remember the default value for the dequeCount (and not forget that in a stress exam situation that there is that property, which is not shown in the code, very convenient).

upvoted 29 times

Anvsoc Most Recent 1 year, 10 months ago

correct

upvoted 1 times

tis_truth 1 year, 10 months ago

Would love to answer this correctly but the full code isn't showing. The maxDequeueCount value isn't showing in the question codeset provided. Does anyone have the full codeset?

upvoted 8 times

[Removed] 1 year, 1 month ago

It is not in the function but in the host.json. Microsoft is testing whether you forgot about that. I really do not know why they try to trick people but then it is Microsoft

upvoted 4 times

coffecold 1 year, 1 month ago

maxDequeueCount; deafult =5
https://docs.microsoft.com/fr-fr/azure/azure-functions/functions-bindings-storage-queue#hostjson-settings

upvoted 2 times

still6dark 2 years, 5 months ago

Image isn't showing a full code? Please send me full code

upvoted 7 times

j888 2 years, 3 months ago

Same.. obviously something is missing

upvoted 2 times

glam 2 years, 5 months ago

Correct.

upvoted 1 times

kapetan 2 years, 6 months ago

The last statement is true: take a look at the input parameters:
...[Table("Orders")]ICollector<Order> table bindings...
and in the code it adds the order:
tableBindings.Add(JsonConvert.DeserializeObject<Object>(myQueueItem.AsString));

upvoted 2 times

idrisfl 2 years, 7 months ago

Seems correct
https://docs.microsoft.com/fr-fr/azure/azure-functions/functions-bindings-storage-queue#hostjson-settings

upvoted 22 times

Marusyk 2 years, 7 months ago

Answer is correct

upvoted 3 times

Kuna_Lambo 2 years, 7 months ago

AZ-203 Topic 1 Q#9

upvoted 2 times

Question #25 *Topic 2*

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:

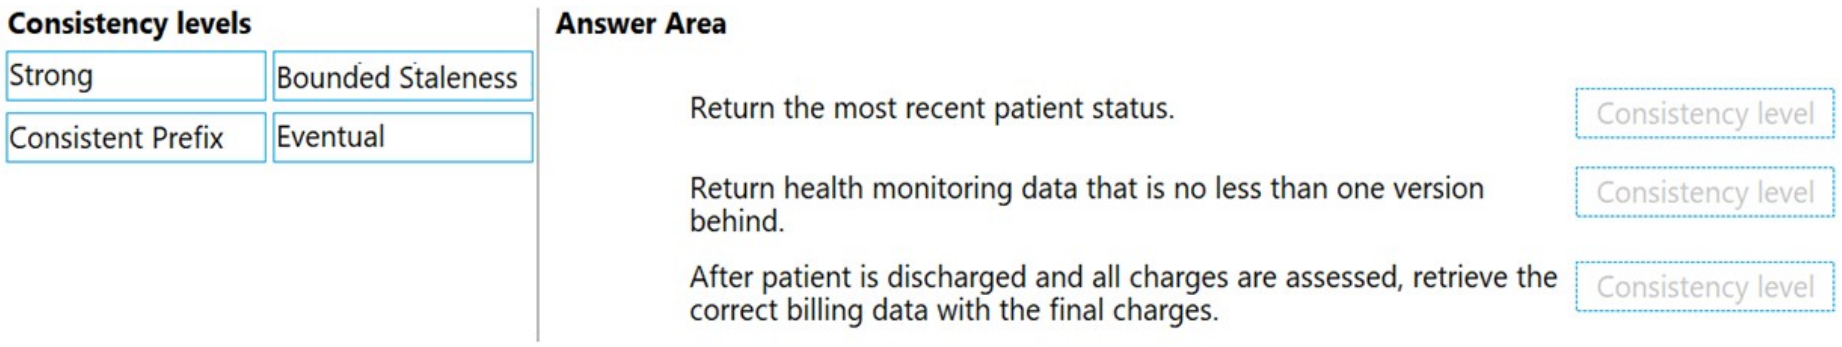

Correct Answer:

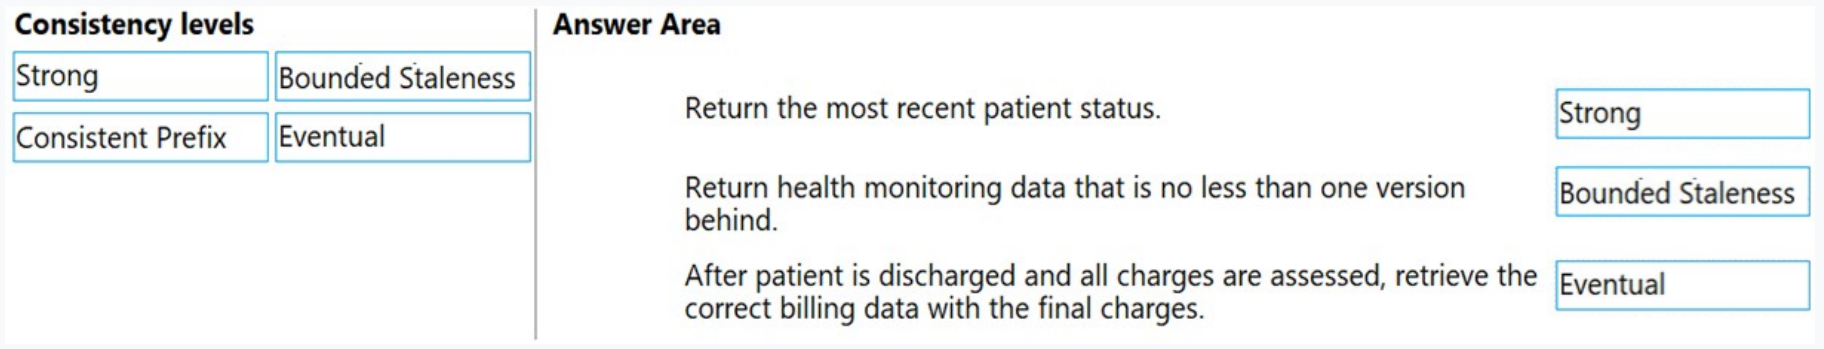

Box 1: Strong -

Strong: Strong consistency offers a linearizability guarantee. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Box 2: Bounded staleness -

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is

"updates") of an item or by "t" time interval. When you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (t) by which the reads might lag behind the writes

Box 3: Eventual -

Eventual: There's no ordering guarantee for reads. In the absence of any further writes, the replicas eventually converge.

Incorrect Answers:

Consistent prefix: Updates that are returned contain some prefix of all the updates, with no gaps. Consistent prefix guarantees that reads never see out-of-order writes.

Reference:

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

AndresMza Highly Voted 2 years, 7 months ago

Answers are correct

upvoted 75 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: Strong

Box 2: Bounded staleness

Box 3: Eventual

Note: Consistent prefix: Updates that are returned contain some prefix of all the updates, with no gaps. Consistent prefix guarantees that reads never see out-of-order writes.

Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels
upvoted 30 times

TakumaK 2 years, 4 months ago
what is the Note for in your comment? you just copied it from the correct answer description which doesn't make sense to me.
upvoted 22 times

MrAzz 4 days, 20 hours ago
Consistent Prefix: This guarantees that a read operation will return the most recent version of the data written (the order is correct), but it could be outdated. It does not guarantee the most recent value of all writes anywhere.

For example, if data is written in A, B, C order, the user may get A, B or A, B, C. It will not get out of order writes such as A, C or A, C, B.
https://www.mssqltips.com/sqlservertip/7158/azure-cosmos-db-consistency-levels-strong-bounded-staleness-session-consistent-prefix-eventual/
upvoted 2 times

dddddd111 Most Recent 1 day, 2 hours ago
This is in my exam last Sept 20, 2023. I chose same answers.
upvoted 1 times

agiler 1 month ago
for box 3, why should i risk to have a wrong bill because the database was not aligned yet?
upvoted 1 times

shekhar11 1 month, 1 week ago
In my exam on 22sep 2023
upvoted 1 times

kingAzure 2 months ago
Got this Q on my exam 25/08/2023
upvoted 2 times

yawxir 1 month, 3 weeks ago
Hi, How was your exam, considering the latest Open Book development ?
upvoted 1 times

jakobste 2 months, 1 week ago
In a real world scenario, if the last charge is applied and then the final bill is immediately generated from a query against the database you might be in trouble with eventual consistency. However, if you can read the data as events in order asynchronous it would not be a problem.

I don't think i would go through the effort though, i would just use strong consistency to generate the final bill.
upvoted 2 times

yawxir 1 month, 3 weeks ago
Agreed. I think the answer in last should be Strong consistency due to financially critical scenario.
upvoted 1 times

JH81 4 months ago
Got this on 6/28/2023 and passed with 850. Went with answer.
upvoted 2 times

Sleepers 4 months ago
Got this question on 27/06/2023
upvoted 1 times

LanGo 6 months, 2 weeks ago
Got this question today (17/04/2023). Answered strong, bounded staleness, eventual. Result: 846
upvoted 4 times

Papit 6 months, 3 weeks ago
rcvd this question 04/2023
upvoted 1 times

Una20 7 months, 3 weeks ago
received the qn in exam 02/2022
upvoted 2 times

[Removed] 1 year, 1 month ago

Box 1: Strong
Box 2: Argument 1: You cannot set the staleness less than 10 versions, so Box-2 is not Bounded staleness. It is not Consistent prefix because there is no control at all on how old the version you are getting. So, it has to be Strong.
Argument 2: Because default consistency is Strong the writes will be synchronous in all regions. And the fact that this is Strong consistency shows that this is not multi-region writes because Strong consistency is not supported in Multi-region writes. So, a bounded-staleness will work (but it could be consistent prefix as well or even be eventual because all the replicas are kept in sync while writing. That part I'm still confused). But still making it Strong or Bounded-staleness does not matter for RUs considering both use local-minority for read.
In conclusion, I plan to bring a lawyer into the exam. Absolutely silly question with full vagueness from Microsoft
Box 3: This could be eventual considering the good application developer has made it as events and there is no ordering required. But it is better to choose Consistent prefix instead of Eventual.

upvoted 6 times

[Removed] 1 year, 1 month ago

The amount of discussion in this question shows how silly the question from Microsoft is. Microsoft should consider not asking scenario questions because it all depends on a lot of assumptions.

upvoted 4 times

dookebroyle 1 year, 3 months ago

I believe it is Strong, Strong, and possibly Strong for the last answer as well.
The first Strong is self explanatory. The second must be strong because bounded staleness' minimum value is 10 versions for single region and even more for multiple regions. The last Strong I believe is true because you need to be sure that all transactions are recorded before you access the final bill as a read transaction. However, it doesn't specify the maximum latency required to access the bill, and many hospitals send bills out after the patient has checked out and the data isn't needed immediately. A difficult question...the answer could be eventual as well. If anyone knows what the answer was on the exam, it would be helpful to know :)

upvoted 2 times

[Removed] 1 year, 1 month ago

I understand how you feel. A very silly question from Microsoft. It all depends on how it is handled. Impossible to guess what the question creator employee of Microsoft knows and thinks is the right answer. If his linkedin profile is given I can research what he has done in his life and guess the answer

upvoted 5 times

NombreFalso 8 months, 2 weeks ago

I love reading your reactions

upvoted 1 times

Leo1328 1 year, 6 months ago

https://docs.microsoft.com/en-us/learn/modules/explore-azure-cosmos-db/5-choose-cosmos-db-consistency-level
When the consistency level is set to bounded staleness, Cosmos DB guarantees that the clients always read the value of a previous write, with a lag bounded by the staleness window.

upvoted 1 times

xRiot007 1 year, 7 months ago

1 - Strong - strong consistency will guarantee the read of the most recent data.
2 - Bounded staleness
3. - Eventual - billing info is like a shopping list. It is irrelevant in what order billing info is retrieved as final charges will be summarized into a total anyway.

upvoted 7 times

Question #26 *Topic 2*

HOTSPOT -

You are configuring a development environment for your team. You deploy the latest Visual Studio image from the Azure Marketplace to your Azure subscription.

The development environment requires several software development kits (SDKs) and third-party components to support application development across the organization. You install and customize the deployed virtual machine (VM) for your development team. The customized VM must be saved to allow provisioning of a new team member development environment.

You need to save the customized VM for future provisioning.

Which tools or services should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Action	Tool or service
Generalize the VM.	Azure PowerShell Visual Studio command prompt Azure Migrate Azure Backup
Store images.	Azure Blob Storage Azure Data Lake Storage Azure File Storage Azure Table Storage

Correct Answer:

Answer Area

Action	Tool or service
Generalize the VM.	Azure PowerShell Visual Studio command prompt Azure Migrate Azure Backup
Store images.	Azure Blob Storage Azure Data Lake Storage Azure File Storage Azure Table Storage

Box 1: Azure Powershell -

Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks.

Before you begin, make sure that you have the latest version of the Azure PowerShell module.

You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image.

Box 2: Azure Blob Storage -

You can store images in Azure Blob Storage.

Reference:

https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource#create-an-image-of-a-vm-using-powershell

st003 Highly Voted 2 years, 6 months ago

The answer is right, it is show in AZ-900 as well.

upvoted 55 times

clarionprogrammer 2 years, 6 months ago

Powershell is correct.
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource

upvoted 9 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: Azure Powershell
Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks. Before you begin, make sure that you have the latest version of the Azure PowerShell module. You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image.

Box 2: Azure Blob Storage
A VM Image is a collection of metadata and pointers to a set of VHDs (one VHD per disk) stored as page blobs in Azure Storage.

Reference:

https://azure.microsoft.com/en-us/blog/vm-image-blog-post

https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource

upvoted 41 times

nlk0 Most Recent 3 months, 2 weeks ago

Is it still relevant for az-204 since they removed VMs from scope in 28th of april?

upvoted 3 times

macobuzi 1 month, 3 weeks ago

I would say better safe than sorry because Microsoft tends to give a sh*t about scope.

upvoted 3 times

tom112 5 months, 3 weeks ago

It's Visual Studio Command Prompt. You can't generalize VM to make an image in azure.
<Windows>
Sysprep removes all your personal account and security information, and then prepares the machine to be used as an image. For information about Sysprep, see Sysprep overview.

Once Sysprep has finished, set the status of the virtual machine to Generalized.
=> So, you generalize VM in the OS side(windows/linux/etc), and just set the status in Azure powershell.
https://learn.microsoft.com/en-us/azure/virtual-machines/generalize

upvoted 1 times

IrshadF 7 months, 1 week ago

It was asked in exam on 26/03/23 as well.

upvoted 1 times

sarmaria 7 months, 2 weeks ago

Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 2 times

Esward 9 months ago

Given answers are correct!
https://learn.microsoft.com/en-us/azure/virtual-machines/generalize

upvoted 2 times

Magis 11 months, 4 weeks ago

I think that answer is incorrect.

Correct answer:
-Visual Studio Command Prompt as first of all you need to Generalize VM. Generalize means that you need to run SysPrep and you can do it from OS side only. So Azure PowerShell cannot be used at this point but you can do it from Visual Studio Command Prompt as it has integration with OS CMD.
https://learn.microsoft.com/en-us/azure/virtual-machines/generalize

- Second option is correct. Azure Blob Storage.

upvoted 4 times

AllNickNamesTaken 8 months, 1 week ago

I agree. Everyone else appears to be answering the question "What tool should you use to create an image?" which would be PowerShell. But, that is not the question being asked. The question is "What tool should you use to generalize the image?" and you are correct with VC Cmd Prompt.

upvoted 1 times

rol204 1 year ago

why blob image I can not understand

upvoted 1 times

Dani_ac7 1 year, 2 months ago

Answer is correct, Powershell because we need arm Blob Storage templates to store virtual machine images, specifically Page Blobs

upvoted 1 times

Eltooth 1 year, 4 months ago

Powershell
Blob Storage

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose as Generalize the VM = Azure PowerShell ; Store images = Azure Blob Storage

upvoted 4 times

elequiel 11 months, 2 weeks ago

thanks for coment date

upvoted 1 times

BobCui 1 year, 10 months ago

Why not Azure Backup for the first box?

upvoted 5 times

RochaG2 1 year, 2 months ago

I've got no ideia.. and no one seems to explain it

upvoted 2 times

ozbonny 2 years, 3 months ago

powershell and blob storage

upvoted 2 times

glam 2 years, 5 months ago

correct.

upvoted 3 times

titombo 2 years, 7 months ago

What is the problem to store the VM images on Azure File Storage?

upvoted 4 times

DanVe 2 years, 7 months ago

Page blobs are the standard storage area for VM

upvoted 5 times

andsol 2 years, 7 months ago

Looks like the answer is correct.

upvoted 5 times

Question #27 *Topic 2*

You are preparing to deploy a website to an Azure Web App from a GitHub repository. The website includes static content generated by a script.
You plan to use the Azure Web App continuous deployment feature.
You need to run the static generation script before the website starts serving traffic.
What are two possible ways to achieve this goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Add the path to the static content generation tool to WEBSITE_RUN_FROM_PACKAGE setting in the host.json file.

B. Add a PreBuild target in the websites csproj project file that runs the static content generation script.

C. Create a file named run.cmd in the folder /run that calls a script which generates the static content and deploys the website.

D. Create a file named .deployment in the root of the repository that calls a script which generates the static content and deploys the website.

Correct Answer: *AD*
A: In Azure, you can run your functions directly from a deployment package file in your function app. The other option is to deploy your files in the d:\home\site
\wwwroot directory of your function app (see A above).
To enable your function app to run from a package, you just add a WEBSITE_RUN_FROM_PACKAGE setting to your function app settings.
Note: The host.json metadata file contains global configuration options that affect all functions for a function app.
D: To customize your deployment, include a .deployment file in the repository root.
You just need to add a file to the root of your repository with the name .deployment and the content:
[config]
command = YOUR COMMAND TO RUN FOR DEPLOYMENT
this command can be just running a script (batch file) that has all that is required for your deployment, like copying files from the repository to the web root directory for example.
Reference:
https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package

Community vote distribution
BD (85%) Other

minsma Highly Voted 2 years, 8 months ago
I think it is B and D
upvoted 67 times

rdemontis 2 years, 7 months ago
I think you're correct: in fact
- run.cmd is used only to start a project as a dll file https://www.sohaibtariq.com/console-webjob/%C3%ACndex/

- WEBSITE_RUN_FROM_PACKAGE doesn't consent the execution of any script. You can only run your web project from a package (.zip file tipically)
https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package
https://github.com/Azure/app-service-announcements/issues/84

- Instead in .csproj file PreBuild Event Target you can specify any command to execute before the compilation and the application execution.
https://docs.microsoft.com/en-us/visualstudio/msbuild/msbuild-targets?view=vs-2019
https://stackoverflow.com/questions/44818730/is-there-a-net-core-cli-pre-before-build-task
https://stackoverflow.com/questions/28916414/visual-studio-add-pre-build-event-that-always-runs-c-project
upvoted 10 times

[Removed] 1 year, 1 month ago
Me: Do I need to know .Net shit to pass AZ-204?
Microsoft: No, it is only concept you need to know. It does not check your knowledge of .Net (or whatever it is)
Questions: .csproj file PreBuild Event Target
Me: Microsoft oh Microsoft.
upvoted 29 times

dimsok 8 months, 3 weeks ago
This is not .NET shit
upvoted 5 times

Arrrqqq 2 years, 3 months ago

More direct info https://docs.microsoft.com/en-us/azure/app-service/deploy-best-practices#net
"By default, Kudu executes the build steps for your .NET application (dotnet build)" - so pre-build step should work there. There is similar note for node.js

upvoted 3 times

Vegetta95 1 year, 7 months ago

If the solution is built on the local machine, you can add target before build and do operations like copy files, run .bat(which can generate script content). However, if the repository is hosted on GitHub, any target will not work, because GitHub does not build apps

upvoted 2 times

ZodiaC 2 years, 3 months ago

1000000% CORRECT !!!!!

upvoted 2 times

rolandcha Highly Voted 2 years, 8 months ago

the answer is :
C,D

upvoted 30 times

Drgn 2 years, 8 months ago

I agree C and D. The explanation of the Answers point to a Azure Function App, either the question was entered wrong (as they meant Azure Function App) into ExamTopics or the answers they chose were wrong.

upvoted 7 times

ACCP1 2 years, 8 months ago

It could be A. It is used here for a web app. https://docs.microsoft.com/en-us/azure/app-service/deploy-run-package

upvoted 5 times

Robert12345Robert 2 years, 4 months ago

Why is this upvoted? When you check the link you can read yourself that this will not run any scripts for you.

upvoted 10 times

[Removed] 1 year, 1 month ago

But the repository could have a script and that will be run automatically by kudu ... Sorry, I do not know it myself whether kudu will be applied when it is from github. There are a lot of permutation combination and ways in which something can be done. To expect someone to remember all the unfortunate choices Microsoft made in these choices is ridiculous. But then it is my mistake I chose to do this certification from Microsoft

upvoted 1 times

Kitkit 2 years, 8 months ago

Can you explain why you think c is the answer?

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

'C' makes no sense. Nothing exists on the Azure website in regard to a "run.cmd".

upvoted 6 times

dreamcoder 1 year, 10 months ago

C, D is correct.
To customize your deployment, include a .deployment file in the repository root. For more information, see Customize deployments and Custom deployment script.
https://docs.microsoft.com/en-us/azure/app-service/deploy-continuous-deployment?tabs=github
https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script

upvoted 3 times

aragones Most Recent 5 months, 3 weeks ago

Got this 2023-05-12.

Had cases:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 3 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

FYI:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 5 times

notedo 6 months, 4 weeks ago

Thanks for commenting

upvoted 1 times

Saluk_DE 7 months ago

Question was in Exam 2023-03-30

upvoted 2 times

PP2015 7 months ago

Did my exam on 3/29/2023. This question was on it. I Selected B and D. Score 850

upvoted 2 times

adilkhan 7 months, 2 weeks ago

C D chatGPT

upvoted 1 times

macobuzi 2 months, 2 weeks ago

ChatGPT is not always right, actually

upvoted 1 times

DonH 4 months, 3 weeks ago

without any additional comment this answer is useless. Just typing it in ChatGPT doesn't mean it's helpfull.

upvoted 4 times

uffuchsi 8 months, 1 week ago

Selected Answer: CD

It has to be C, or DIf you want to customize the deployment process, for example you want to run your tests before deploying (or after) and cancel the deployment if they fail?

Any of these commands will generate the files required to deploy your site, mainly:

.deployment - Contains the command to run for deploying your site.
deploy.cmd - Contains the deployment script (or deploy.sh if running on Mac/Linux)

upvoted 2 times

NombreFalso 8 months, 2 weeks ago

Selected Answer: BD

A is wrong, it's never mentioned to be an Azure Function. There is nothing wrong with B and D.

upvoted 1 times

JamieS 9 months, 1 week ago

Selected Answer: BD

Without research I thought A and D, but static files (from what I can tell) are housed in wwwroot, but https://learn.microsoft.com/en-us/azure/azure-functions/run-functions-from-deployment-package states that "When you run from a package, the wwwroot folder becomes read-only and you'll receive an error when writing files to this directory. Files are also read-only in the Azure portal."

so unless it has been reconfigured not to use wwwroot for content, this can't be right.

I'd then lean towards B and D like the majority, but then this is assuming msbuild for the website, but what if it's another technology like Ruby? It wouldn't have a .csproj. Unfortunately we're not given a big enough world view, so we'll have to make assumptions, if we're making assumptions I think it must be B & D, B would certainly work if was an MSBuild project.

upvoted 2 times

rotimislaw 10 months ago

Selected Answer: BD

As stated by others here, B&D

upvoted 1 times

hubekpeter 11 months, 2 weeks ago

Selected Answer: BD

A - wrong, WEBSITE_RUN_FROM_PACKAGE this useful to run a function from a package, B - correct, that's a standard way, C - wrong, run.cmd does nothing, D - correct- custom deployment script https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script

upvoted 3 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: BD

B, D is correct.

upvoted 2 times

rol204 1 year ago

should I choose the most voted or the answer selected by the site in the exam?

upvoted 3 times

mmdex 8 months, 2 weeks ago

You should consider both, do some research and thinking, and then decide for yourself.

upvoted 2 times

[Removed] 1 year, 1 month ago

It is not clear if the answer should be about option to customize the deployment then both the .csproj and .deployment options are possible. But if there are two things to do be done, then it is to first deploy the stuff from github.
https://github.com/projectkudu/kudu/wiki/Customizing-deployments

upvoted 1 times

dookebroyle 1 year, 3 months ago

Selected Answer: BD

D.) https://docs.microsoft.com/en-us/azure/app-service/deploy-local-git?tabs=cli
B.) https://docs.microsoft.com/en-us/visualstudio/msbuild/msbuild-targets?view=vs-2022 (for build pipeline in Azure DevOps)

upvoted 5 times

Question #28 *Topic 2*

DRAG DROP -

You are developing an application to use Azure Blob storage. You have configured Azure Blob storage to include change feeds.

A copy of your storage account must be created in another region. Data must be copied from the current storage account to the new storage account directly between the storage servers.

You need to create a copy of the storage account in another region and copy the data.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

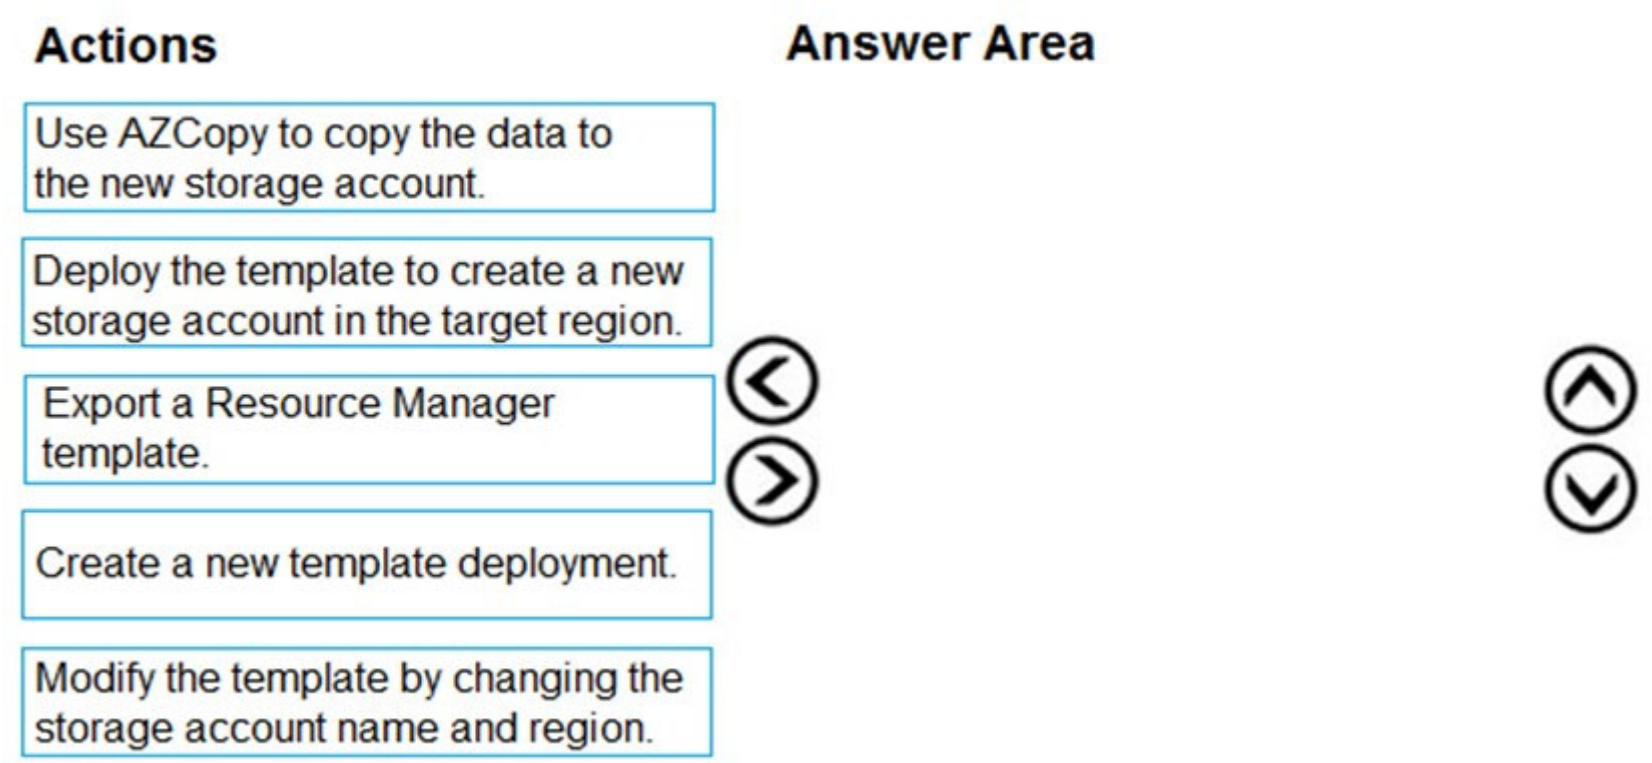

Correct Answer:

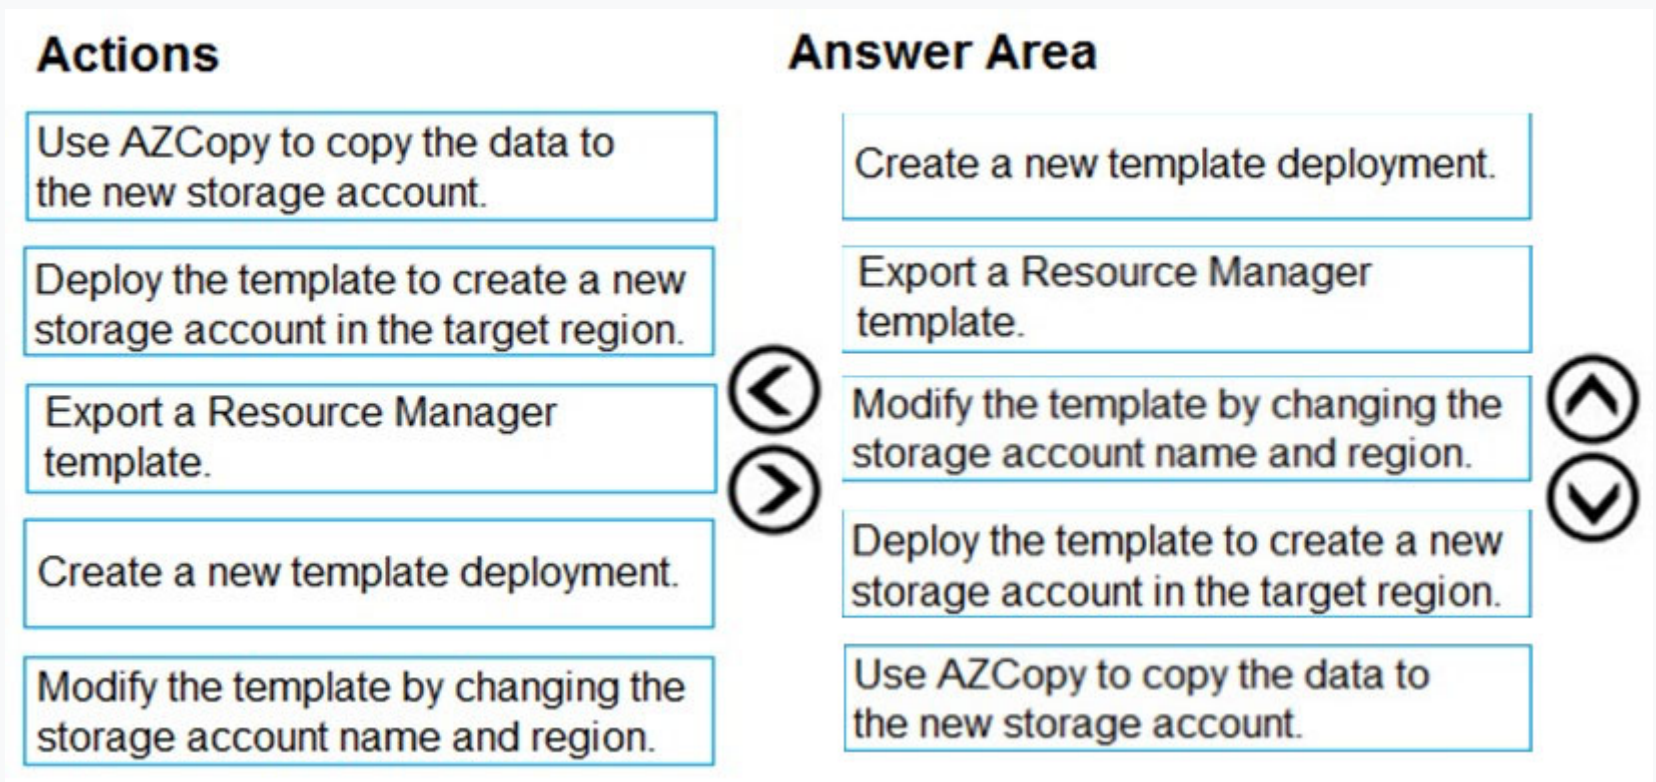

To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice.

The steps are:

☞ Export a template.

☞ Modify the template by adding the target region and storage account name.

☞ Deploy the template to create the new storage account.

☞ Configure the new storage account.

☞ Move data to the new storage account.

☞ Delete the resources in the source region.

Note: You must enable the change feed on your storage account to begin capturing and recording changes. You can enable and disable changes by using Azure

Resource Manager templates on Portal or Powershell.

Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed

MariusN Highly Voted 2 years, 11 months ago

I think you first have to export the Resource Manager template before you can create a new template deployment. So, swap options 1 and 2 from the solution. In the first link of the solution's text, exporting is also considere as the first step.

upvoted 143 times

uffuchsi 8 months, 1 week ago

100% agree, but I think as long as you create the deployment befpore you deploy that is okay. So the order could also be:
* Export a template.
* Modify the template by adding the target region and storage account name.
* Create a new Template deployment
* Deploy the template to create the new storage account.
* Use AzCopy

Please correct me and explain if this is not true.

upvoted 9 times

shadownin 7 months, 1 week ago

Do you speak English? You've said 100% agree and changed what he said
Guys its
Export -
Create New Template Deployment
Modify
Deploy
AzCopy

As per the guide, Yes the header is confusing but look at the instructions. Create Template deployment is before modifying

upvoted 4 times

LucreziaB 3 months, 2 weeks ago

No need to be this aggressive tho

upvoted 12 times

b_sidra 2 years, 11 months ago

Agreed with this statement.

reference: https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal

upvoted 5 times

KarstenWolff 1 year, 2 months ago

I believe it is only important that the template is exported before modifying it.
So, first "Create new template deployment", then "Export Resource Manager template", or first "Export Resource Manager template" and then "Create new template deployment". Then "Modify", "Deploy" and "AZCopy".
So, the suggested solution and solution proposed by Marius should be both correct.

upvoted 3 times

Kobee 2 years, 11 months ago

Solution is right.
The option "Create a new Template deployment" is not present on the first link.
And you can first create the template if you want to export it. So Create first and then export the new template

upvoted 1 times

victor Highly Voted 2 years, 9 months ago

Answer is Wrong. correct Sequence is

Export
Create
Modify
Deploy
AZ copy

upvoted 105 times

noro5 1 year, 8 months ago

Yes, confirmation: https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal

upvoted 5 times

Robert12345Robert 2 years, 3 months ago

I agree, see:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/quickstart-create-templates-use-the-portal

upvoted 4 times

macobuzi Most Recent 2 months, 2 weeks ago

According to Microsoft document:https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal
- Export
- Modify
- Deploy
- Create
-AzCopy

upvoted 1 times

macobuzi 1 month, 3 weeks ago

Sorry, correct order is:
- Export
- Modify
- Create
- Deploy
- AzCopy

Need to create the template before deploying it

upvoted 2 times

maqsoodshah 7 months ago

1. Export a template.
2. Modify the template by adding the target region and storage account name.
3. Deploy the template to create the new storage account.
4. Create a new template deployment
5. Use AZCopy to Copy data to the new storage account.

https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal

upvoted 3 times

macobuzi 2 months, 2 weeks ago

You messed up at 3 and 4. You need to create the template before deploying it.

upvoted 1 times

Defa 7 months, 3 weeks ago

Modify before create.
https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal

upvoted 2 times

Strawberry_123 8 months, 3 weeks ago

As per my understanding, correct answer is-Export/Create/Modify/Deploy/Az Copy
1. Export the template and download/save it.
2.Create a new template and load the previously exported template's json file.
3.Modify the required details.
4.Deploy it.
4.Use Az copy
https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal

upvoted 3 times

narenazure 9 months, 3 weeks ago

https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal

Export
Create
Modify
Deploy
AZ copy

upvoted 2 times

Knightie 1 year, 2 months ago

Creating a Template vs Creating a Template Deployment are two different things.
Export Template. Modify the Template, Create the Deployment, Deployment the Template, Do Copy.

upvoted 4 times

kbsivaiyer 1 year, 3 months ago

https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal
it is export, modify, create, deploy, azcopy

upvoted 4 times

Wengatz 1 year, 2 months ago

Thanks for the link, but according to that, the order is export > create > modify > deploy > azcopy

upvoted 3 times

BasitJamil 1 year, 3 months ago

Export
New Template Manager

Modify
Deploy
AzCopy
upvoted 3 times

Eltooth 1 year, 4 months ago

Export
Create
Modify
Deploy
AZ copy
upvoted 2 times

UmarAbbas 1 year, 5 months ago

We need to export resource before creating new template so the answer would be
Export
Create
Modify
Deploy
AZ copy
upvoted 2 times

rick1010 1 year, 9 months ago

Generate a template using the portal
Modify (Edit template) (Changes Value like region)
Create the template deployment
Deploy the template
Az COpy
upvoted 8 times

trickerk 2 years ago

The first step described in the solution is Export a template, so MaruisN is right!
upvoted 1 times

debanjan10 2 years ago

The order must be: Export > Create > Modify > Deploy> Use Az Copy
upvoted 3 times

debanjan10 2 years ago

Order:
Export -> Create -> Modify -> Deploy -> Use AzCopy
upvoted 2 times

cool_tool 2 years, 2 months ago

First, we need to export the exiting resource manager template that was used to create the original storage account.
==> Export a Resource Manager template

Then you can modify the template and change the storage account name and region for the new storage account.
==> Modify the template by changing the storage account and region

Then create a new template in Azure with the help of the MODIFIED template
==> Create a new template deployment.

Then
==> Deploy the template to create a new storage account in the target region

And then
==> Use the AzCopy tool to transfer the data from the source to destinationstorage account.
upvoted 19 times

Question #29 *Topic 2*

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:

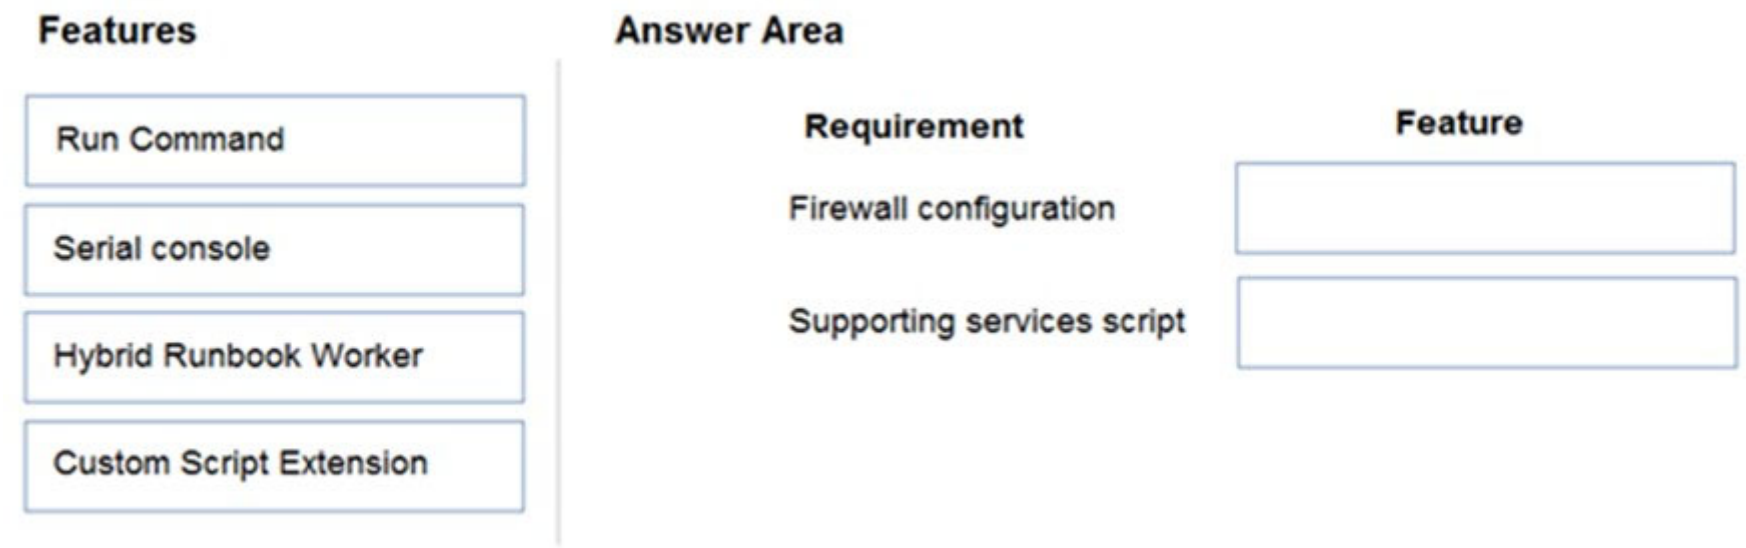

Correct Answer:

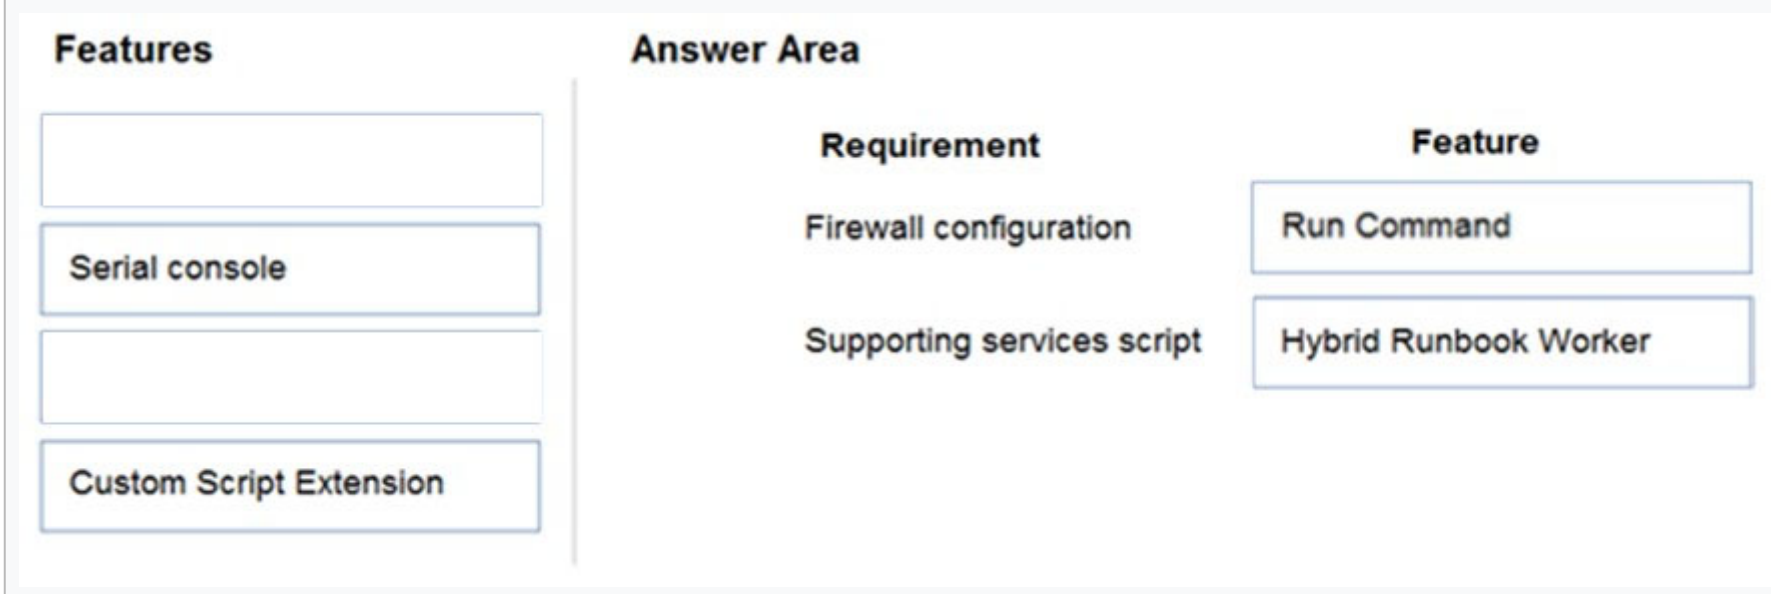

Reference:
https://docs.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-command

agueda Highly Voted 2 years, 7 months ago

1. Run Command
2. Customer Script Extension

Some question: https://www.examtopics.com/discussions/microsoft/view/12062-exam-az-203-topic-1-question-11-discussion/

upvoted 94 times

if54uran 2 years, 4 months ago

but it says the firewall cofiguration has to happen automatically. How does "Run command" do that?

upvoted 8 times

Magis 11 months, 4 weeks ago

You can run Run Command by using Azure CLI and automate it in this way.

https://learn.microsoft.com/en-us/azure/virtual-machines/windows/run-command

upvoted 2 times

titombo 2 years, 7 months ago

On this link is showing the same as we have now, 1. Run Command and 2. Hybrid Runbook Worker

upvoted 3 times

BrettusMaximus 2 years, 6 months ago

Cant use run to configure firewall. The run command uses the VM OS but the VM does not have access at that point.
upvoted 2 times

fesioche 2 years, 2 months ago
The Run Command feature enables virtual machine and application management and troubleshooting using scripts, and is available even when the machine is not reachable, for example if the guest firewall doesn't have the RDP or SSH port open.

Ref: https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-scripts-in-vm#run-command
upvoted 8 times

pmsiva 2 years, 7 months ago
The Custom Script Extension downloads and executes scripts on Azure virtual machines. This extension is useful for post deployment configuration, software installation, or any other configuration or management tasks. Scripts can be downloaded from Azure storage or GitHub, or provided to the Azure portal at extension run time.
upvoted 5 times

mlantonis Highly Voted 2 years, 5 months ago
Box 1: Run Command
This capability is useful in all scenarios where you want to run a script within a VM. It's one of the only ways to troubleshoot and remediate a VM that doesn't have the RDP or SSH port open, because of improper network or administrative user configuration.

Box 2: Customer Script Extension
The Custom Script Extension downloads and executes scripts on Azure virtual machines. This extension is useful for post deployment configuration, software installation, or any other configuration or management tasks. Scripts can be downloaded from Azure storage or GitHub, or provided to the Azure portal at extension run time. The Custom Script Extension integrates with Azure Resource Manager templates, and can be run using the Azure CLI, PowerShell, Azure portal, or the Azure Virtual Machine REST API.
upvoted 69 times

mlantonis 2 years, 5 months ago
Reference:

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

https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-scripts-in-vm
upvoted 10 times

nlk0 Most Recent 3 months, 2 weeks ago
Is it still relevant for az-204 since they removed VMs from scope from 28th april 2023
upvoted 1 times

macobuzi 2 months, 2 weeks ago
As someone has stated, the question appeared on 29th May 2023, which means it is still within the scope. But seriously, I don't care whether it is out of scope or not, I just study everything from Examtopic, better safe than sorry!
upvoted 5 times

nlk0 3 months, 2 weeks ago
is it still scope of az 204, since VMs were removed from scope in 28th april 2023?
upvoted 1 times

LuchoGonzalez 4 months, 1 week ago
1) Run command
2) Customer script extension(https://learn.microsoft.com/en-us/azure/virtual-machines/extensions/custom-script-windows)
upvoted 1 times

deathRac3 5 months ago
Question was there for me on 29th May 2023
upvoted 6 times

tom112 5 months, 3 weeks ago
Azure Automation Agent-based User Hybrid Runbook Worker (Windows and Linux) will retire on 31 August 2024
https://learn.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker
upvoted 1 times

KarthikCS 6 months, 1 week ago
Why are we tested on ridiculously low level details like this in these azure exams? It makes no sense. These AZ exams are the most senseless exams I've ever seen.
upvoted 3 times

PP2015 7 months ago
Did my exam on 3/29/2023. This question was on it. Selected Run Command, Customer Script Extension. Score 850
upvoted 5 times

Esward 9 months ago
Ru coomand, Custom Script Extensions are correct!

https://learn.microsoft.com/en-us/azure/virtual-machines/windows/run-command
upvoted 1 times

narenazure 9 months, 3 weeks ago

1. Run
2. CSE (The Custom Script Extension downloads and runs scripts on Azure virtual machines (VMs). This extension is useful for post-deployment configuration, software installation, or any other configuration or management task. You can download scripts from Azure Storage or GitHub, or provide them to the Azure portal at extension runtime.)

upvoted 2 times

micro9000 11 months, 3 weeks ago

It should be
1. Run Command
2. Custom Script Extension
because there is no hybrid setup mentioned on the question

upvoted 1 times

vcfvct 1 year ago

Got this on 10/21/2022. went with Run/CSE

upvoted 3 times

coffecold 1 year, 1 month ago

firewall configuration : hybrid runbook worker
suporting services : Custom Script Extension

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

firewall configuration : must be executed on the VM and eventually on the host machine and must be done automatically.
suporting services : once access to the VM, it can be run using an ARM

upvoted 1 times

warchoon 8 months, 1 week ago

No
Runbook worker
- Requires manual installation
- the firewall must be automatically configured

upvoted 1 times

Eltooth 1 year, 4 months ago

Run
CSE

upvoted 2 times

meoukg 1 year, 7 months ago

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

upvoted 6 times

Bartimaeus 1 year, 8 months ago

Why not Serial console for Firewall? It also doesn't require SSH access and it's easier to use that Run Command.

upvoted 1 times

Bogdan75 1 year, 7 months ago

"the firewall must be AUTOMATICALLY configured"

upvoted 1 times

Question #30 *Topic 2*

HOTSPOT -

A company is developing a Node.js web app. The web app code is hosted in a GitHub repository located at https://github.com/TailSpinToys/webapp.

The web app must be reviewed before it is moved to production. You must deploy the initial code release to a deployment slot named review.

You need to create the web app and deploy the code.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
[ ▼ ] -Name myResourceGroup -Location $location
[ ▼ ] -Name $webappname -Location $location -ResourceGroupName myResourceGroup -Tier Standard
[ ▼ ] -Name $webappname -Location $location -AppServicePlan $webappname -ResourceGroupName myResourceGroup
[ ▼ ] -Name $webappname -ResourceGroupName myResourceGroup -Slot review
$PropertiesObject = @{repoUrl = "$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
```

New-AzWebAppSlot / New-AzWebApp / New-AzAppServicePlan / New-AzResourceGroup

Correct Answer:

Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
[ ▼ ] -Name myResourceGroup -Location $location          → New-AzResourceGroup
[ ▼ ] -Name $webappname -Location $location -ResourceGroupName myResourceGroup -Tier Standard   → New-AzAppServicePlan
[ ▼ ] -Name $webappname -Location $location -AppServicePlan $webappname -ResourceGroupName myResourceGroup   → New-AzWebApp
[ ▼ ] -Name $webappname -ResourceGroupName myResourceGroup -Slot review   → New-AzWebAppSlot
$PropertiesObject = @{repoUrl = "$gitrepo";branch = "master";}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup
-SourceSlotName review -DestinationSlotName production
```

Box 1: New-AzResourceGroup -

The New-AzResourceGroup cmdlet creates an Azure resource group.

Box 2: New-AzAppServicePlan -
The New-AzAppServicePlan cmdlet creates an Azure App Service plan in a given location

Box 3: New-AzWebApp -
The New-AzWebApp cmdlet creates an Azure Web App in a given a resource group

Box 4: New-AzWebAppSlot -
The New-AzWebAppSlot cmdlet creates an Azure Web App slot.
Reference:
https://docs.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroup?view=azps-2.3.2 https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azappserviceplan?view=azps-2.3.2 https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebapp?view=azps-2.3.2 https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebappslot?view=azps-2.3.2

AndresMza Highly Voted 2 years, 7 months ago
Answer is correct
upvoted 80 times

mlantonis Highly Voted 2 years, 5 months ago
Box 1: New-AzResourceGroup
The New-AzResourceGroup cmdlet creates an Azure resource group.

Box 2: New-AzAppServicePlan
The New-AzAppServicePlan cmdlet creates an Azure App Service plan in a given location

Box 3: New-AzWebApp
The New-AzWebApp cmdlet creates an Azure Web App in a given a resource group

Box 4: New-AzWebAppSlot
The New-AzWebAppSlot cmdlet creates an Azure Web App slot.
upvoted 33 times

mlantonis 2 years, 5 months ago
Reference:
https://docs.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroup?view=azps-2.3.2
https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azappserviceplan?view=azps-2.3.2
https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebapp?view=azps-2.3.2
https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebappslot?view=azps-2.3.2
upvoted 7 times

shekhar11 Most Recent 1 month, 1 week ago
Question in my exam on 22sept 2023
upvoted 2 times

Saluk_DE 7 months ago
Question was in Exam 2023-03-30
upvoted 3 times

sarmaria 7 months, 2 weeks ago
Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.
upvoted 1 times

Ayman99 8 months ago
Was on the exam 27.02.2023
upvoted 2 times

70PineApple 8 months, 1 week ago
Got this in exam today..20/02/23
score: 817
upvoted 1 times

NombreFalso 8 months, 2 weeks ago
As vanilla as it gets
upvoted 1 times

menlapwa 9 months, 3 weeks ago
The answer make sense to me
upvoted 1 times

saravanasanthosh 10 months ago
Got this on exam 12/30/2022

upvoted 1 times

SoftSol 10 months, 3 weeks ago

Correct, got this in the exam today dec 13, 2022

upvoted 1 times

uditrm12 10 months, 1 week ago

examtopic questions are enough to pass?

upvoted 1 times

coffecold 1 year, 1 month ago

Finally an easy to answer question, without bad wordings

upvoted 1 times

Eltooth 1 year, 4 months ago

ResourceGroup
ServicePlan
WebApp
WebAppSlot

upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 4 times

debanjan10 2 years ago

Always:
Resource Group / Group -> App Service Plan -> Web App -> Web App Slot -> Web App Source

upvoted 11 times

[Removed] 2 years, 4 months ago

There is another question similar to this one but with AzureCLI. But share the same structure.

upvoted 3 times

UnknowMan 2 years, 5 months ago

Answer is correct

upvoted 2 times

Question #31 *Topic 2*

HOTSPOT -

You are developing an application that needs access to an Azure virtual machine (VM).

The access lifecycle for the application must be associated with the VM service instance.

You need to enable managed identity for the VM.

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

NOTE: Each correct selection is worth one point.

Hot Area:

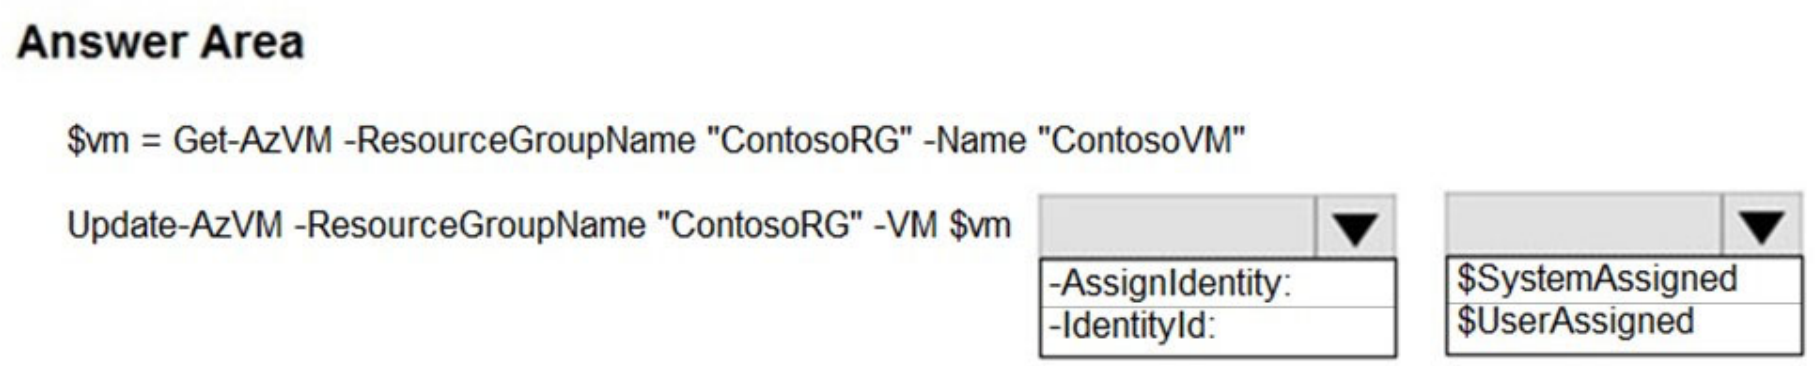

Correct Answer:

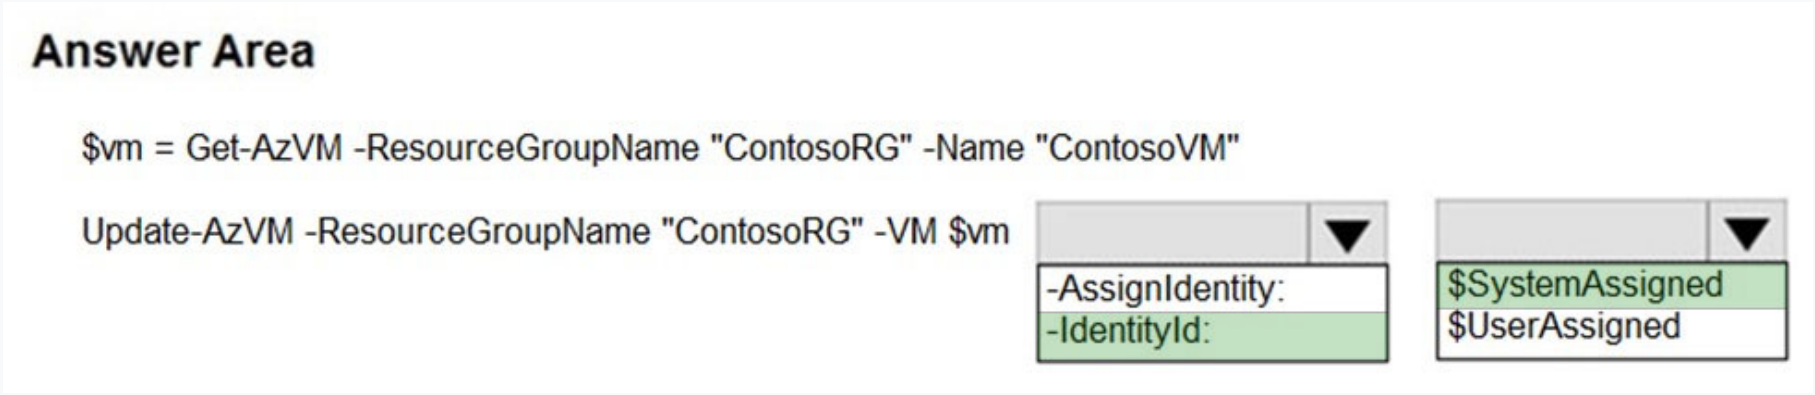

Box 1: -IdentityType -

Enable system-assigned managed identity on an existing Azure VM:

To enable a system-assigned managed identity, use the -IdentityType switch on the Update-AzVM cmdlet (see below).

Box 2: $SystemAssigned -

$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM

Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -IdentityType SystemAssigned

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm

agueda Highly Voted 2 years, 7 months ago

The parameter should be "IdentityType", not "IdentityId" as it's stated in the reference link.

"SystemAssigned" is correct

upvoted 73 times

shocks 2 years, 6 months ago

The $ before "SystemAssigned" is superflous.

upvoted 5 times

Cholo981 2 years, 4 months ago

1- You can think there are two distinct errors in the question.
2- ...or then you can try to answer thinking both IdentityId and the "$" in the second ddl are ok.
The $ means that's a variable.
So it must contains a specific User Managed Identity path for the paramenter IdentityId. Can't be System Assigned.
I'll go for the second one...

upvoted 7 times

surprise0011 6 months, 2 weeks ago

I thought like you also but there is no -IdentityType UserAssigned parameter along with -IdentityId $UserAssigned

I will assume that on exam will occur parameter IdentityType, so then the choice should be SystemAssigned either when it is variable (with $) or not.
If on exam will show -IdentityType and $SystemAssigned i will assume that variable just contains "SystemAssigned"

here some samples:
https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm#user-assigned-managed-identity

upvoted 1 times

TakumaK 2 years, 4 months ago

then where is the variable defined in the question? and how are you sure the variable are set to what you guess?

upvoted 1 times

Cholo981 2 years, 4 months ago

It is in the question:
"$" INDICATES a poweshell variable. Is right there.

I'm not "guessing" anything. You guys are guessing BOTH the possible answers are wrong written...
My answer is the only possible if the question is correct.

Someone could point that IdentityId require also the IdentityType parameter but:
1- If you don't specify it, the shell wil request it to you (I tried the script, you should do that too, don't count to much on the comments and the upvotes here, many are wrong);
2- Is it really needed if the VM already have a UAMI? this is an update command, you may want to add a UAMI or edit one.

BTW, I just passed the Exam (and I got the question too). So is not my problem anymore. Good luck. :)

upvoted 4 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000

upvoted 7 times

Esward 9 months ago

Agreed

upvoted 1 times

clarionprogrammer 2 years, 6 months ago

Correct. It should be -IdentityType.
-IdentityId is another valid, optional parameter.

upvoted 5 times

clarionprogrammer 2 years, 6 months ago

References:
https://docs.microsoft.com/en-us/powershell/module/az.compute/update-azvm?view=azps-5.8.0&viewFallbackFrom=azps-5.6.0
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.management.compute.models.resourceidentitytype?view=azure-dotnet

upvoted 3 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: -IdentityType
-IdentityType: The type of identity used for the virtual machine. Valid values are SystemAssigned, UserAssigned, SystemAssignedUserAssigned, and None.

-IdentityId: Specifies the list of user identities associated with the virtual machine. The user identity references will be ARM resource IDs in the form:

Box 2: $SystemAssigned
There are two types of managed identities:
- System-assigned: Some Azure services allow you to enable a managed identity directly on a service instance. When you enable a system-assigned managed identity an identity is created in Azure AD that is tied to the lifecycle of that service instance. So when the resource is deleted, Azure automatically deletes the identity for you. By design, only that Azure resource can use this identity to request tokens from Azure AD.
- User-assigned: You may also create a managed identity as a standalone Azure resource. You can create a user-assigned managed identity and assign it to one or more instances of an Azure service. In the case of user-assigned managed identities, the identity is managed separately from the resources that use it.

upvoted 63 times

mlantonis 2 years, 5 months ago

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview#managed-identity-types

upvoted 6 times

edengoforit 1 year, 9 months ago

This answer deserves a top comment

upvoted 1 times

ENGs Most Recent 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

p2006 1 month, 2 weeks ago

https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm#enable-system-assigned-managed-identity-on-an-existing-azure-vm

upvoted 2 times

DonH 4 months, 2 weeks ago

Just for information: I just had this question on my AZ204 exam - 16-jun-2023.
I barely made it (with only 767 points) so I can't inform anyony if this answer is correct or not, just stating that this is an actual exam question.

upvoted 3 times

tom112 5 months, 3 weeks ago

<IdentityType>
The type of identity used for the virtual machine. Valid values are SystemAssigned, UserAssigned, SystemAssignedUserAssigned, and None.

Type: Nullable<T>[ResourceIdentityType]
Accepted values: SystemAssigned, UserAssigned, SystemAssignedUserAssigned, None

upvoted 1 times

fkaracan 8 months, 1 week ago

since $ states they both are variables. given answer is correct.

upvoted 1 times

AllNickNamesTaken 8 months, 1 week ago

Please correct me if I am wrong, but the "The access lifecycle for the application must be associated with the VM service instance" part looks like a red herring. Since it says that the *access* lifecycle must be associated with the service instance and not the *managed identity lifecycle*, then a user assigned identity would work here. Delete the VM and the application has no more access to it. If true, that would make the "IdentityID" parameter correct in this case.

upvoted 1 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23
score: 817

upvoted 2 times

Priya0703 8 months, 1 week ago

Got this question in exam today on 20-02-2023

upvoted 1 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

micro9000 11 months, 3 weeks ago

There's no correct option, it should be "IdentityType"
https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm#enable-system-assigned-managed-identity-on-an-existing-azure-vm
We should update this question and the answer

upvoted 1 times

warchoon 8 months, 1 week ago

I think it's a trick. $SystemAssigned is a variable for ID

upvoted 1 times

[Removed] 1 year, 1 month ago

Another of the evil questions. Microsoft says no you need not know in-depth powershell or azure cli to be able to answer. And look at the options. In azure cli it is : az vm identity assign (see the similarity with assignidentity and that given as an option to punish the person who knows azure cli? That's why "don't be evil")
and in power shell it is -identitytype
The unfortunate choice made by Microsoft "developers" to not be consistent in their naming conventions or lack of communication across their organizational islands (I guess powershell side and bash) is weighing high on the unfortunate people who are trying to complete this certification.
If this kind of questions come in the real exam I'm done with Microsoft related anything (except for a laptop)

upvoted 3 times

NombreFalso 8 months, 2 weeks ago

Man these sure make studying less boring

upvoted 1 times

Dani_ac7 1 year, 2 months ago

Given answer is correct because "The access lifecycle for the application must be associated with the VM service instance"

upvoted 1 times

dookebroyle 1 year, 3 months ago

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-ua-arm#assign-identity

According to that article, the -IdentityType parameter is missing, nonetheless, the answer should be IdentityId and $SystemAssigned

upvoted 1 times

Eltooth 1 year, 4 months ago

IdentityID
SystemAssigned

upvoted 3 times

suku555 1 year, 5 months ago

I am not sure why no one in the discussions tell this.
Hint in the Q - "The access lifecycle for the application must be associated with the VM service instance."
Answer is -AssignIdentity:$SystemAssigned
https://infra.engineer/azure/60-azure-stop-and-deallocate-a-windows-vm-using-a-managed-identity

upvoted 2 times

Question #32 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure
Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Create an Azure Function app that uses the Consumption hosting model and that is triggered from the blob upload.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*

In the Consumption hosting plan, resources are added dynamically as required by your functions.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function

Community vote distribution

B (100%)

AndresMza Highly Voted 2 years, 7 months ago

Answer should be "No". Consumption plan can take up to several minutes to trigger the function. See note from https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function.

"When your function app runs in the default Consumption plan, there may be a delay of up to several minutes between the blob being added or updated and the function being triggered. If you need low latency in your blob triggered functions, consider running your function app in an App Service plan."

upvoted 140 times

forgetfulalligator 10 months, 3 weeks ago

Yeap , you can also use a premium plan to use pre-warmed instances so no cold start

upvoted 3 times

Esward 9 months ago

Use the app service plan and set always on property to true to start processing immediately

upvoted 2 times

rdemontis 2 years, 7 months ago

In my opinion you're right. Often in a consumption plan the function starts immediately but it's not guaranteed and in production environment minimizing latency is very important as required in the question. For further info about cold starts i found an interesting article: https://azure.microsoft.com/en-us/blog/understanding-serverless-cold-start/#:~:text=In%20the%20context%20of%20Azure,haven't%20been%20called%20recently.

upvoted 7 times

DirkK Highly Voted 2 years, 7 months ago

Hi all.
I would say, that this has to be "No"...
Quote: "The process to produce a mobile-friendly version of the image must start in less than one minute."
And this is only a feature from the "Premium plan" (Perpetually warm instances to avoid any cold start)....
Consumption plan = up to several minutes...

upvoted 18 times

Tarajee Most Recent 1 month, 2 weeks ago

Got this on2023sept

upvoted 2 times

Chris17 6 months, 1 week ago

Selected Answer: B

The answer should be "No" because with consuption plan take more than 1 minuts, in the before question. same to the answer, it is with app service plan.

upvoted 1 times

NombreFalso 8 months, 2 weeks ago

The explanation is bull. Don't buy it. It's B.

upvoted 1 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

dookebroyle 1 year, 3 months ago

Selected Answer: B

Answer should be "no" after reviewing this article. The processing needs to begin sooner and Microsoft specifies the difference between an Event Grid trigger and a storage blob trigger.

https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=in-process%2Cextensionv5&pivots=programming-language-csharp
"Minimizing latency: If your function app is on the Consumption plan, there can be up to a 10-minute delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled. You can also use an Event Grid trigger with your Blob storage account. For an example, see the Event Grid tutorial."

upvoted 4 times

Satish_Babu 1 year, 4 months ago

Selected Answer: B

Correction answer is B (No)
Because - When your function app runs in the default Consumption plan, there may be a delay of up to several minutes between the blob being added or updated and the function being triggered. If you need low latency in your blob triggered functions, consider running your function app in an App Service plan.

upvoted 2 times

azure900practice 1 year, 4 months ago

Selected Answer: B

Functin app should be in App serice plan with 'Always On' setting as true. Consumption plan have cold starts and latency.

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.
Consumption plan needs to warm up (up to 10 minutes) before processing.

upvoted 2 times

Freidrich 1 year, 8 months ago

Selected Answer: B

The correct answer is B.

upvoted 2 times

sozturk88 1 year, 8 months ago

Selected Answer: B

10 min delay for Consuming

upvoted 2 times

philsboies 1 year, 8 months ago

Selected Answer: B

I'd say no

upvoted 2 times

Zamaletto 1 year, 8 months ago

Selected Answer: B

I would go for B

upvoted 2 times

MFahd 1 year, 9 months ago

The answer is NO, as the consumption plan have low latency and It can take some time to add files to blob and trigger function

upvoted 1 times

DiegoManinetti 1 year, 9 months ago

Selected Answer: B

Highly Voted

upvoted 2 times

mabusalma 1 year, 9 months ago

Selected Answer: B

Answer: B
Consumption plan not guaranteed

upvoted 1 times

Question #33 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.
You need to ensure that scripts run and resources are available before a swap operation occurs.
Solution: Update the app with a method named statuscheck to run the scripts. Update the app settings for the app. Set the WEBSITE_SWAP_WARMUP_PING_PATH and WEBSITE_SWAP_WARMUP_PING_STATUSES with a path to the new method and appropriate response codes.
Does the solution meet the goal?

A. No

B. Yes

Correct Answer: *A*
These are valid warm-up behavior options, but are not helpful in fixing swap problems.
Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.
Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot.
Here's a sample web.config fragment.
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostName="[app hostname]" />
<add initializationPage="/Home/About" hostName="[app hostname]" />
</applicationInitialization>
</system.webServer>
Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps

Community vote distribution
B (76%) A (24%)

Carlous Highly Voted 2 years, 7 months ago
Should be YES?
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots
You can also customize the warm-up behavior with one or both of the following app settings:

WEBSITE_SWAP_WARMUP_PING_PATH: The path to ping to warm up your site. Add this app setting by specifying a custom path that begins with a slash as the value. An example is /statuscheck. The default value is /.
WEBSITE_SWAP_WARMUP_PING_STATUSES: Valid HTTP response codes for the warm-up operation. Add this app setting with a comma-separated list of HTTP codes. An example is 200,202 . If the returned status code isn't in the list, the warmup and swap operations are stopped. By default, all response codes are valid.
WEBSITE_WARMUP_PATH: A relative path on the site that should be pinged whenever the site restarts (not only during slot swaps). Example values include /statuscheck or the root path, /.
upvoted 78 times

warchoon 8 months, 1 week ago
NO. Auto swap should be on the source
upvoted 4 times

Tom87 2 years, 6 months ago
I agree.
For ASP.NET and ASP.NET Core developers, setting app settings in App Service are like setting them in <appSettings> in Web.config or appsettings.json, but the values in App Service override the ones in Web.config or appsettings.json.
https://docs.microsoft.com/en-us/azure/app-service/configure-common#configure-app-settings
upvoted 5 times

altafpatel1984 1 year, 11 months ago
Question is not only for warm-up but also to execute custom script. So Answer No is correct.

upvoted 10 times

mcanic 1 year, 10 months ago

it states that a new method called statuscheck is added to the application that executes these scripts. if you add the path to /statuscheck in WEBSITE_SWAP_WARMUP_PING_PATH the application executes the scripts

upvoted 9 times

Santileo 2 years, 6 months ago

This guy is right. Better url: https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up

upvoted 6 times

Percy2112 2 years, 4 months ago

I Agree.

upvoted 3 times

simonverma Highly Voted 2 years, 6 months ago

The answer is right to be marked as 'No' as the warm up is done correctly but the web config part is not mentioned i.e, AppInitialization.

upvoted 18 times

lugospod 1 year, 9 months ago

you dont need appinit tif you set WEBSITE_SWAP_WARMUP_PING_PATH and response... it is an alternate method to achieve the same thing.

upvoted 6 times

Tarajee Most Recent 1 month, 2 weeks ago

Got this on 2023sept

upvoted 1 times

mihailos 2 months ago

Got it in exam 28/08/23. Went with YES for this. Scored 912

upvoted 2 times

yawxir 1 month, 3 weeks ago

Hey,
Glad you pass. Can you guide more on exam pattern as this should be open book now ?

upvoted 2 times

JH81 4 months, 1 week ago

Selected Answer: A

Answer is NO. While this solution may work put on your MS thinking hat because modifications to the app should not be necessary when there are other ways to get this to work with no app modifications.

upvoted 1 times

red0400 7 months, 2 weeks ago

Selected Answer: B

Answer is YES

upvoted 1 times

uffuchsi 8 months, 1 week ago

Surely it's yes?
You can customize the warm-up behavior with one or both of the following app settings:

WEBSITE_SWAP_WARMUP_PING_PATH: The path to ping to warm up your site. Add this app setting by specifying a custom path that begins with a slash as the value. An example is /statuscheck. The default value is /.

WEBSITE_SWAP_WARMUP_PING_STATUSES: Valid HTTP response codes for the warm-up operation. Add this app setting with a comma-separated list of HTTP codes. An example is 200,202 . If the returned status code isn't in the list, the warmup and swap operations are stopped. By default, all response codes are valid.

WEBSITE_WARMUP_PATH: A relative path on the site that should be pinged whenever the site restarts (not only during slot swaps). Example values include /statuscheck or the root path, /.

https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up

upvoted 1 times

rgullini 10 months ago

Selected Answer: B

The solution doesn´t include the scripts. Answer is NO.

upvoted 1 times

Lucifer14 10 months ago

In 28-12-2022 exam

upvoted 1 times

memooo 10 months, 2 weeks ago

Selected Answer: B

Correct answer is B:
Reference:
https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up
upvoted 1 times

mdg3501 10 months, 3 weeks ago
got this on 2022-12-7
upvoted 1 times

nvtienanh 11 months ago
Got this question on 2-Dec-2022 exam. Went with most voted
upvoted 3 times

uditrm12 10 months, 1 week ago
examtopics questions are enough to pass?
upvoted 2 times

hubekpeter 11 months, 2 weeks ago
Selected Answer: B
You can bind virtually whatever you want under /status path as far as it'll return a correct status code when it's ready. Yes, you can implement this differently, as many of you stated, but that's not the point. It's perfectly valid solution.
upvoted 1 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: B
B. Yes
upvoted 1 times

[Removed] 1 year, 1 month ago
Those two settings are enough for custom warm-up. From Microsoft documentation:
"The <applicationInitialization> configuration element is part of each app start-up, whereas the two warm-up behavior app settings apply only to slot swaps."
So that means that the appInitialization is not required. The requirement is only about warmup during swap and not with every startup. So, it satisfies. I hope to not get questions like this for which I need a lawyer.
upvoted 6 times

aliyevmirzakhan 1 year, 1 month ago
Selected Answer: A
I think the answer is NO,
Because ,IMO, the /statuscheck is just for ping, although you can run scripts inside the method invoked, but that is not its purpose. So it may work, but it is not "elegant" solution.
upvoted 1 times

rqFamily 1 year, 3 months ago
the answer is no
first such scripts are needed prior to application startup.
second http timeout is 230 seconds, and scripts might take longer.
so scripts needs to be part of applicationInitialization
upvoted 1 times

mhall1 1 year, 1 month ago
Warm-up can include scripts run after the actual web app has initialized. The instance won't be ready for swapping until the endpoint specified in the variable responds with the specified status. Whether that happens before or after the app has loaded doesn't matter.
upvoted 1 times

Question #34 *Topic 2*

HOTSPOT -

You create the following PowerShell script:

```
$source = New-AzScheduledQueryRuleSource -Query 'Heartbeat | where TimeGenerated > ago(1h)' -DataSourceId "contoso"
$schedule = New-AzScheduledQueryRuleSchedule -FrequencyInMinutes 60 -TimeWindowInMinutes 60
$triggerCondition = New-AzScheduledQueryRuleTriggerCondition -ThresholdOperator "LessThan" -Threshold 5
$aznsActionGroup = New-AzScheduledQueryRuleAznsActionGroup -ActionGroup "contoso" -EmailSubject "Custom email subject"
 -CustomWebhookPayload "{ '"alert'":'"#alertrulename'", '"IncludeSearchResults'":true }"
$alertingAction = New-AzScheduledQueryRuleAlertingAction -AznsAction $aznsActionGroup -Severity "3" -Trigger $triggerCondition
New-AzScheduledQueryRule -ResourceGroupName "contoso" -Location "eastus" -Action $alertingAction -Enabled $true
 -Description "Alert description" -Schedule $schedule -Source $source -Name "Alert Name"
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	○	○
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	○	○
The log alert is scheduled to run every two hours.	○	○

Correct Answer:

Answer Area

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	○	●
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	●	○
The log alert is scheduled to run every two hours.	○	●

Box 1: No -

The AzScheduledQueryRuleSource is Heartbeat, not CPU.

Box 2: Yes -

The AzScheduledQueryRuleSource is Heartbeat!

Note: New-AzScheduledQueryRuleTriggerCondition creates an object of type Trigger Condition. This object is to be passed to the command that creates Alerting

Action object.

Box 3: No -

The schedule is 60 minutes, not two hours.

-FrequencyInMinutes: The alert frequency.

-TimeWindowInMinutes: The alert time window

The New-AzAscheduledQueryRuleSchedule command creates an object of type Schedule. This object is to be passed to the command that creates Log Alert

Rule.

Reference:
https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryrule https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryruletriggercondition

mlantonis Highly Voted 2 years, 5 months ago

I agree.
No
Yes
No

upvoted 80 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000

upvoted 7 times

wolf_lu Highly Voted 2 years, 4 months ago

I agree.
NO
YES
NO

upvoted 8 times

barretowski 1 year, 7 months ago

don't confuse the staff, the answer is: no, yes, no

upvoted 1 times

JH81 Most Recent 4 months ago

Got this on 6/28/2023 and passed with 850. Went with No, Yes, No.

upvoted 1 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

FYI:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

ucskips 7 months, 2 weeks ago

1. No
2. Yes
3. No

I used the last command to get the following understanding

My logic:
Commands will send an email
Where threshold is < 5
Every 60 mins
Where heartbeat is > 1h ago

upvoted 1 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23
score: 817

upvoted 2 times

Priya0703 8 months, 1 week ago

Got this question on exam today 20-02-2023.

upvoted 2 times

18Marks 10 months, 3 weeks ago

Got it in exam. Use highly voted answer

upvoted 2 times

uditrm12 10 months, 1 week ago

examtopics questions are enough to pass?

upvoted 1 times

70PineApple 8 months, 1 week ago

yes they are

upvoted 2 times

70PineApple 8 months, 1 week ago

yes they are

upvoted 1 times

Nokaido 1 year, 1 month ago

What about the 60 min trigger frequency + the 60 min recording window? As far as I know it only triggers after the previous trigger is finished. It should be NO, YES, YES?

upvoted 2 times

JustHereToLearn 10 months, 3 weeks ago

No, the answer should be No, Yes, No as per this link https://learn.microsoft.com/en-us/javascript/api/@azure/arm-monitor/schedule?view=azure-node-latest

upvoted 3 times

nuiowmokka 1 year, 2 months ago

Did anyone get this question in real exam?

upvoted 3 times

ImPavan 10 months ago

I got today 28-12-2022

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

Answer is correct: NO, YES, NO

upvoted 1 times

yaroo1 2 years ago

Answer is correct

upvoted 3 times

Percy2112 2 years, 4 months ago

The Answers are Correct:
No
Yes
No

upvoted 3 times

Arul4 2 years, 4 months ago

No
Yes
No
are the correct answers.

upvoted 4 times

Question #35 *Topic 2*

DRAG DROP -

You are developing an Azure Function app.

The app must meet the following requirements:

- Enable developers to write the functions by using the Rust language.
- Declaratively connect to an Azure Blob Storage account.

You need to implement the app.

Which Azure Function app features should you use? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

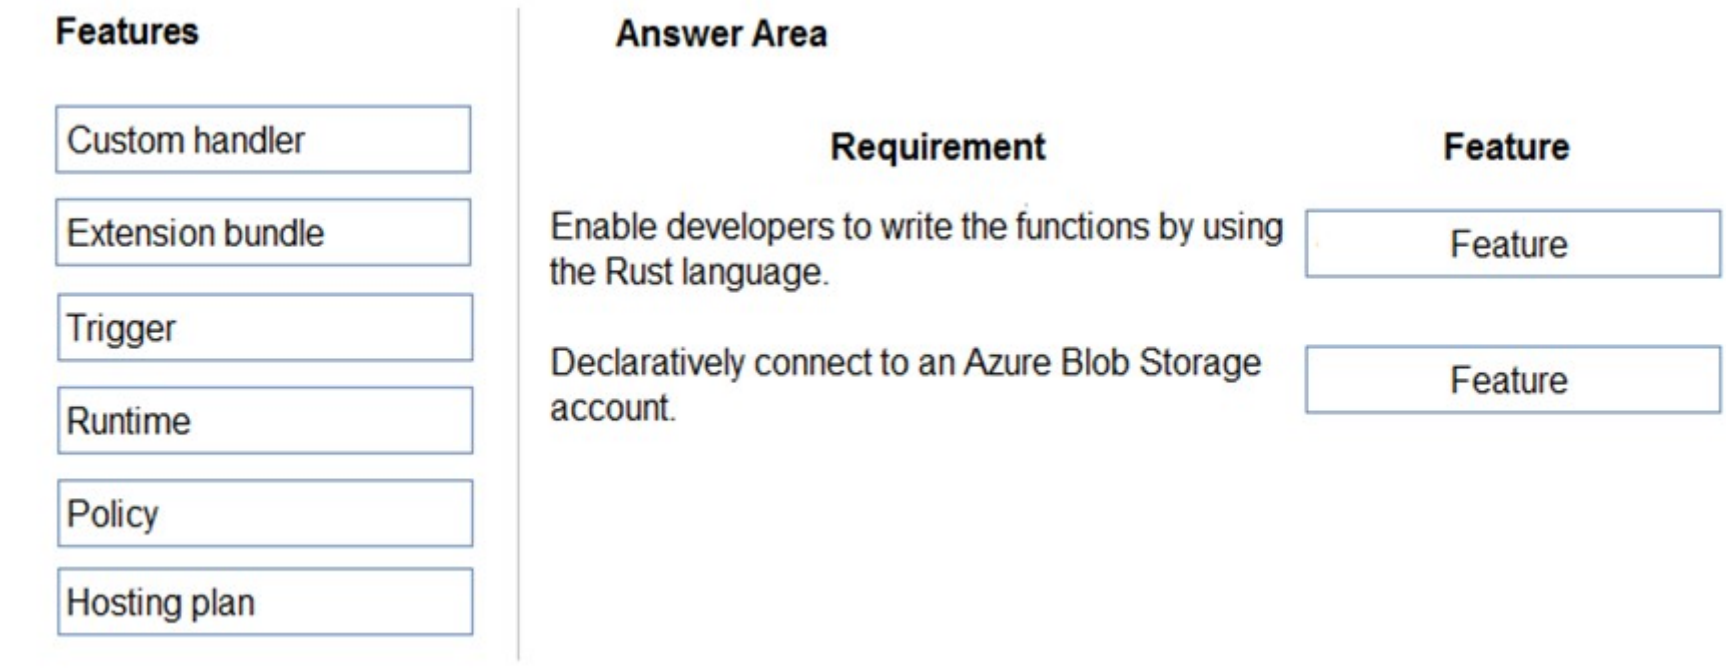

Correct Answer:

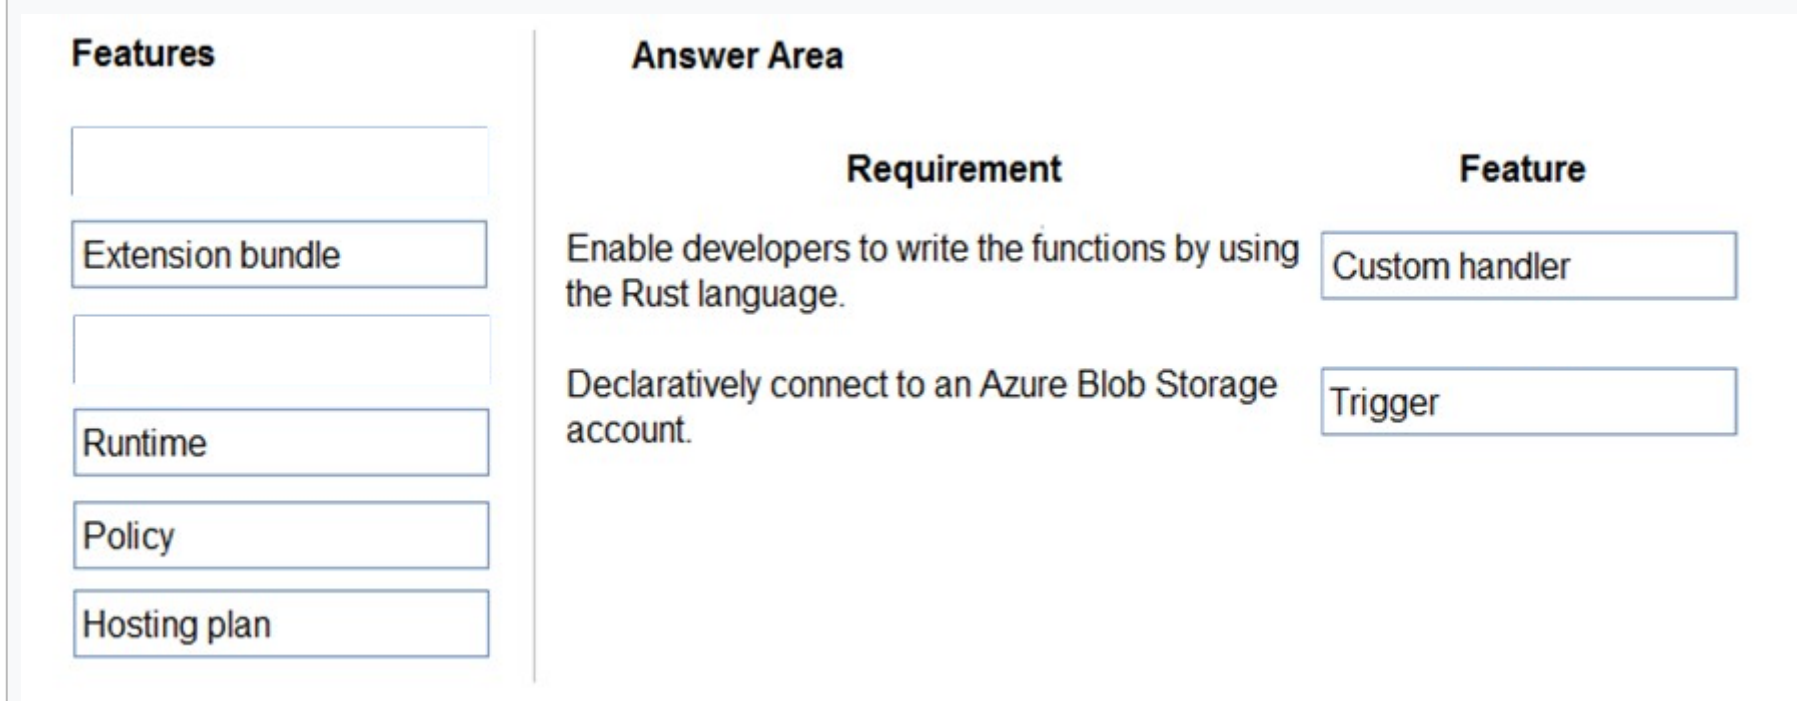

Box 1: Custom handler -

Custom handlers can be used to create functions in any language or runtime by running an HTTP server process, for example Go or Rust.

Box 2: Trigger -

Functions are invoked by a trigger and can have exactly one. In addition to invoking the function, certain triggers also serve as bindings. You may also define multiple bindings in addition to the trigger. Bindings provide a declarative way to connect data to your code.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/create-first-function-vs-code-other https://docs.microsoft.com/en-us/dotnet/architecture/serverless/azure-functions

abdou1987 Highly Voted 2 years, 4 months ago

Box 1: Custom handler
Custom handlers can be used to create functions in any language or runtime by running an HTTP server process, for example Go or Rust.

Box 2: extension bundles
is needed to support the bindings and triggers that you use
https://docs.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers?WT.mc_id=thomasmaurer-blog-thmaure#bindings-support

upvoted 76 times

sbsd 8 months, 1 week ago

https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers
- With custom handlers, you can use triggers and input and output bindings via extension bundles.
So: Custom Handler and Extension bundles.

upvoted 1 times

mlantonis Highly Voted 2 years, 5 months ago

Answer is correct.

Box 1: Custom handler
Custom handlers can be used to create functions in any language or runtime by running an HTTP server process, for example Go or Rust.

Box 2: Trigger
Functions are invoked by a trigger and can have exactly one. In addition to invoking the function, certain triggers also serve as bindings. You may also define multiple bindings in addition to the trigger. Bindings provide a declarative way to connect data to your code.

Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/create-first-function-vs-code-other
https://docs.microsoft.com/en-us/dotnet/architecture/serverless/azure-functions

upvoted 67 times

mmdex 8 months, 2 weeks ago

Box 2: Extension Bundles

Yes, you could use a Trigger, but then you have only ONE trigger per function. The question says nothing about how the function is being triggered, only that you need to connect to Blob Storage. What if you wanted to trigger the function by HTTP request and then perform some action in Blob Storage? You can't do that with Blob Storage trigger.

upvoted 3 times

Esward 9 months ago

Agreed

upvoted 1 times

TakumaK 2 years, 4 months ago

incorrect!

upvoted 2 times

stevenwu 2 years, 4 months ago

What's your answer? You can just point out incorrect without any updates.

upvoted 16 times

azurelearner666 2 years, 4 months ago

your incorrect is incorrect X2.
The answer is correct, Custom Handler for supporting additional languages in azure functions and a trigger for adding a binding to a function.

upvoted 10 times

aruni_mishra 11 months, 1 week ago

BOX 2: extension bundles

from MS:
With custom handlers, you can use triggers and input and output bindings via extension bundles.

https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers?WT.mc_id=thomasmaurer-blog-thmaure#bindings-support:~:text=With%20custom%20handlers%2C%20you%20can%20use%20triggers%20and%20input%20and%20output%20bindings%20via%20extension%20bundles.

upvoted 6 times

shekhar11 Most Recent 1 month, 1 week ago

Question in my exam on 22sept 2023

upvoted 2 times

Tarajee 1 month, 2 weeks ago

Got this on 2023sept

upvoted 3 times

Kanasan 2 months ago

For the second I vote for 'Trigger':
https://learn.microsoft.com/en-us/azure/azure-functions/functions-triggers-bindings?tabs=isolated-process%2Cpython-v2&pivots=programming-language-csharp#add-bindings-to-a-function

upvoted 1 times

kvtborad 2 months, 3 weeks ago

I got this question on 6th August 2023. chose highly voted. passed with 904. I got Case study: city. All questions are from ExamTopics.

upvoted 1 times

RikinPatel 1 month ago

There are 2 most voted ans.. so can you please specify which most voted you have selected?

upvoted 1 times

Firo 3 months, 1 week ago

Can anyone mentioned what the answer is for box2, there are 2 answers that got most voted which one is the correct one.

upvoted 2 times

JH81 4 months ago

Got this on 6/28/2023 and passed with 850. Went with highly voted answer.

upvoted 1 times

tom112 5 months, 3 weeks ago

For Box 2: Trigger. (But Binding would make more sense if there is.)

Extension bundles allow non-.NET functions apps to use the same bindings without having to deal with the .NET infrastructure.
-> so it's not for general trigger or bindings setting.

Reference: Register Azure Functions binding extensions
https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-register

upvoted 1 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

FYI:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

adilkhan 7 months, 2 weeks ago

Box 1: Custom handler
Box 2: extension bundles CHATGPT

upvoted 1 times

saravanasanthosh 10 months ago

Got this on exam 12/30/2022

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Box 2: extension bundles or Triggers? which one is correct ? both answers are highly voted with valid comments

upvoted 2 times

guchao2000 9 months, 1 week ago

Trigger
https://learn.microsoft.com/en-us/azure/azure-functions/create-first-function-vs-code-other?tabs=rust%2Cwindows

upvoted 1 times

[Removed] 1 year, 1 month ago

It has to be custom handler and extension bundle if the question has to be read together. But another of those questions that does not make much sense. People who want to create in RUST should use an app service instead of using a function app. It does not add a lot of value to play into these dark areas. Microsoft just adds features that are half baked instead of focusing on fixing the issues or making things consistent. Bad product management.

upvoted 4 times

Knightie 1 year, 2 months ago

Custom Handler , Extension Bundle instead of native triggers.

upvoted 2 times

xRiot007 1 year, 4 months ago

The answer is incorrect.
Box 1 : Custom Handler - to code a function in Rust or Go, you will need one of these.
Box 2: Extension Bundles - to use a standard trigger with a function coded in Rust, you need an extension bundle

upvoted 7 times

Eltooth 1 year, 4 months ago

After reviewing comments I think we can all agree
1. Custom handlers
2. Extension Bundles (most popular answer)

upvoted 8 times

Question #36 *Topic 2*

HOTSPOT -

You are developing an ASP.NET Core web application. You plan to deploy the application to Azure Web App for Containers.

The application needs to store runtime diagnostic data that must be persisted across application restarts. You have the following code:

```
public void SaveDiagData(string data)
{
      var path = Environment.GetEnvironmentVariable("DIAGDATA");
      File.WriteAllText(Path.Combine(path, "data"), data);
}
```

You need to configure the application settings so that diagnostic data is stored as required.

How should you configure the web app's settings? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

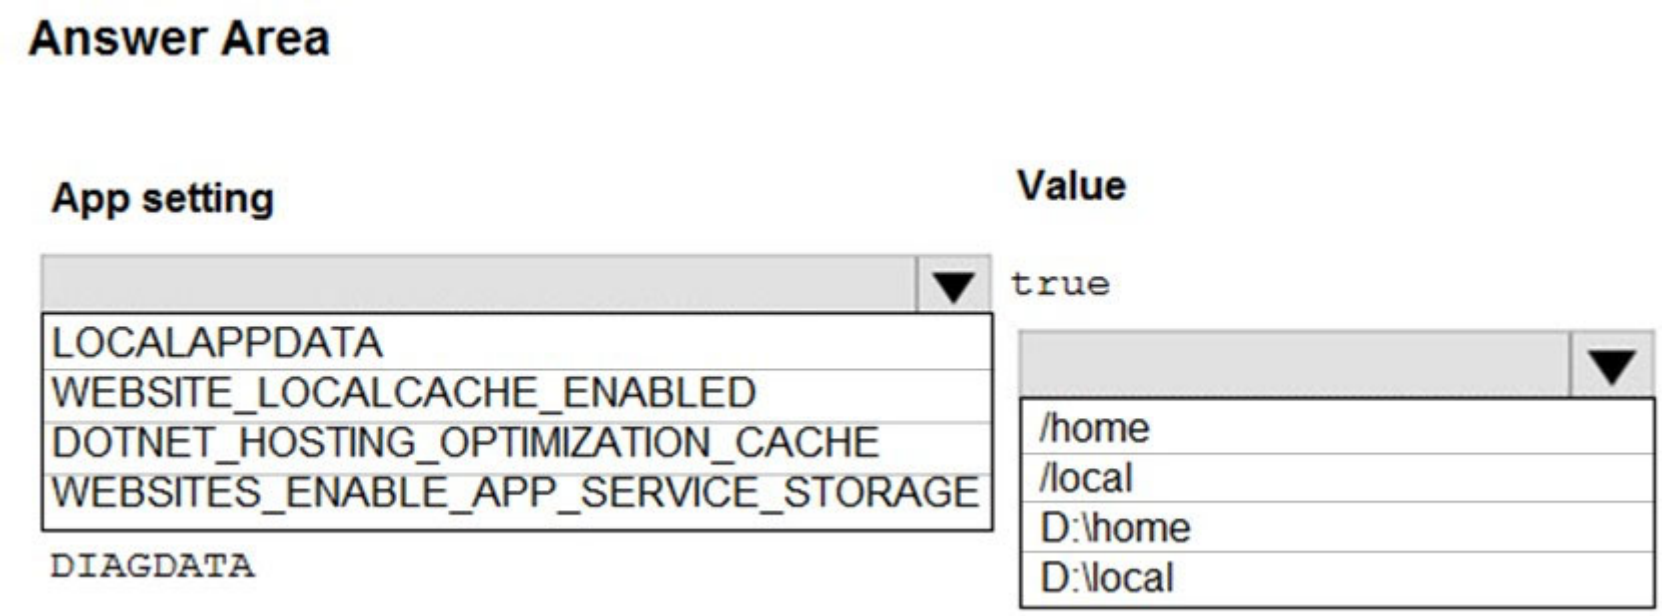

Correct Answer:

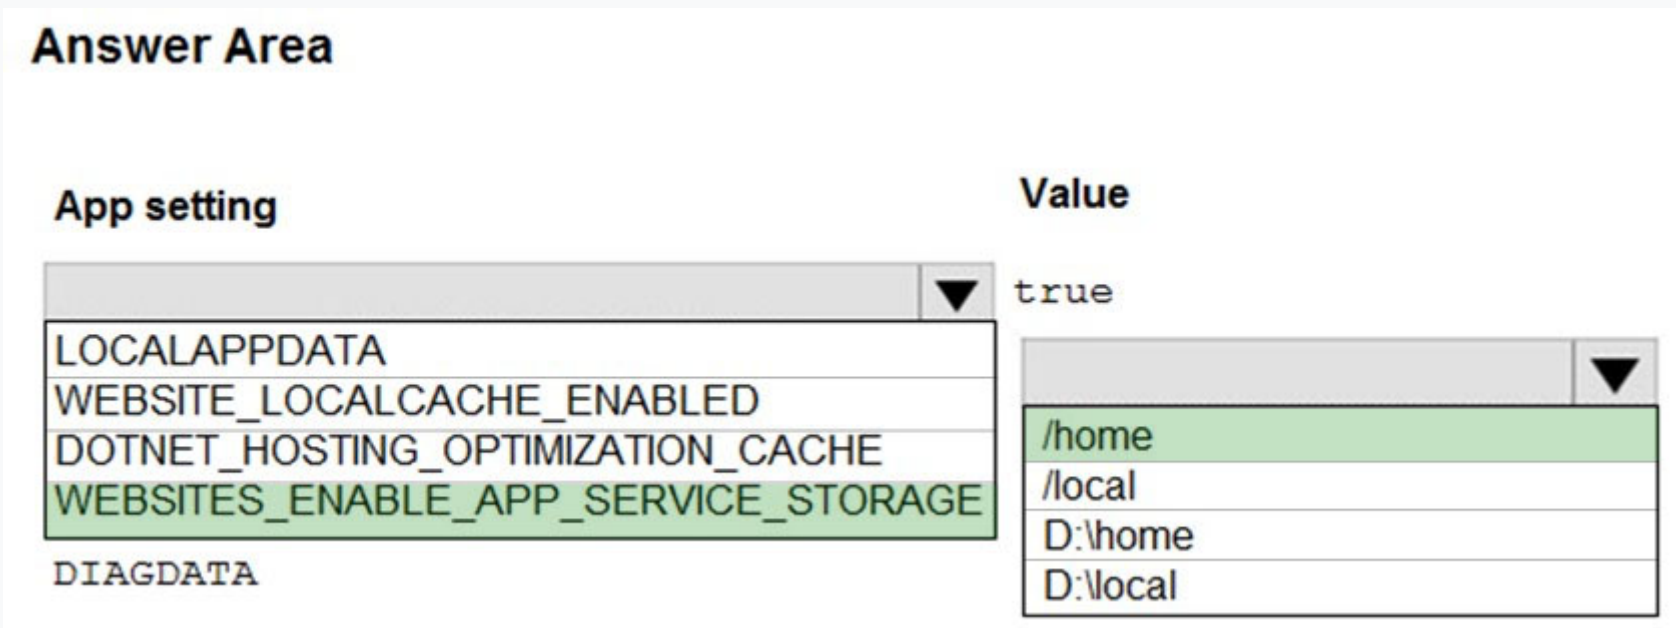

Box 1: If WEBSITES_ENABLE_APP_SERVICE_STORAGE

If WEBSITES_ENABLE_APP_SERVICE_STORAGE setting is unspecified or set to true, the /home/ directory will be shared across scale instances, and files written will persist across restarts

Box 2: /home -

Reference:

https://docs.microsoft.com/en-us/azure/app-service/containers/app-service-linux-faq

zinza Highly Voted 2 years, 9 months ago

I think the printed image is creating confusion. The correct answers (according to the image) are:
WEBSITES_ENABLE_APP_SERVICE_STORAGE=true
DIAGDATA=/home

upvoted 94 times

azurelearner666 2 years, 4 months ago

yep, correct. Same as the "official solution" :)

upvoted 5 times

lugospod 1 year, 9 months ago

How do you know it is Linux and not Windows?

upvoted 1 times

ScubaDiver123456 1 year, 9 months ago

For Windows,it would be C:\Home, which is not a provided option.. but /home is provided.

https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-windows

upvoted 12 times

GCMan Highly Voted 2 years, 12 months ago

Given answer correct.

upvoted 24 times

Esward 9 months ago

Agreed

upvoted 1 times

Kanasan Most Recent 2 months ago

The answer is correct, but explanation is not:
If WEBSITES_ENABLE_APP_SERVICE_STORAGE setting is unspecified or set to false, the /home/ directory will not be shared across scale instances, and files written will not persist across restarts.
https://learn.microsoft.com/en-us/troubleshoot/azure/app-service/faqs-app-service-linux#i-m-using-my-own-custom-container--i-want-the-platform-to-mount-an-smb-share-to-the---home---directory-

upvoted 2 times

micro9000 11 months, 3 weeks ago

The answer is correct
https://learn.microsoft.com/en-us/azure/app-service/reference-app-settings?tabs=kudu%2Cdotnet#custom-containers

upvoted 1 times

[Removed] 1 year, 1 month ago

Custom container itself is a dark area of app service and expecting human beings to remember the property names and the possible values for these settings is very remarkable.

upvoted 6 times

[Removed] 1 year, 1 month ago

I have to just remember all these properties? Are developers not allowed to use documentation in real life? Does search work in Redmond? Microsoft should really think what exactly they achieve from these random questions without a thought on what they want the developers to know or remember.

upvoted 10 times

Eltooth 1 year, 4 months ago

WEBSITES_ENABLE_APP_SERVICE_STORAGE
/home

upvoted 4 times

SivajiTheBoss 1 year, 7 months ago

Answer is correct:
WEBSITES_ENABLE_APP_SERVICE_STORAGE=true
DIAGDATA=/home

upvoted 2 times

upadhyayavi 1 year, 7 months ago

Can't D:/home and /home both be correct in this case??

upvoted 1 times

Ami_Nou 1 year, 5 months ago

For anyone confused, in Windows it is C:\Home

upvoted 4 times

upadhyayavi 1 year, 7 months ago

Answer to both the options: https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-windows#use-persistent-shared-storage

upvoted 1 times

petitbilly 1 year, 7 months ago

Here there is the specific documentation for this question: https://docs.microsoft.com/en-us/azure/app-service/faq-app-service-linux#i-m-using-my-own-custom-container--i-want-the-platform-to-mount-an-smb-share-to-the---home---directory-

upvoted 1 times

qiw 1 year, 9 months ago

WEBSITES_ENABLE_APP_SERVICE_STORAGE=true
By default, persistent storage is disabled on custom containers and the setting is exposed in the app settings. To enable it, set the WEBSITES_ENABLE_APP_SERVICE_STORAGE app setting value to true
https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-linux

DIAGDATA=/home
You can use the /home directory in your custom container file system to persist files across restarts and share them across instances. The /home directory is provided to enable your custom container to access persistent storage.
https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-linux
upvoted 4 times

mlantonis 2 years, 5 months ago
Correct Answer:

If WEBSITES_ENABLE_APP_SERVICE_STORAGE setting is unspecified or set to false, the /home/ directory will not be shared across scale instances, and files written will not persist across restarts. Explicitly setting WEBSITES_ENABLE_APP_SERVICE_STORAGE to true will enable the mount.

Box 1: WEBSITES_ENABLE_APP_SERVICE_STORAGE

Box 2: /home

Reference:

https://docs.microsoft.com/en-us/azure/app-service/containers/app-service-linux-faq
upvoted 16 times

glam 2 years, 5 months ago
WEBSITES_ENABLE_APP_SERVICE_STORAGE=true
DIAGDATA=/home
upvoted 3 times

JoeInOregon 2 years, 5 months ago
glam has the answers, y'all.
upvoted 3 times

Tom87 2 years, 5 months ago
On Windows the directory would be C:\home. Since this is not mentioned in possible answers, we are on Linux environment and the answer is correct.
https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-windows#use-persistent-shared-storage
upvoted 4 times

glam 2 years, 5 months ago
WEBSITES_ENABLE_APP_SERVICE_STORAGE=true
DIAGDATA=/home
upvoted 3 times

CAguys 2 years, 7 months ago
First box is correct but the second box should be 'true'. Please refer the link which is given in the explanation!
upvoted 1 times

Question #37 *Topic 2*

You are developing a web app that is protected by Azure Web Application Firewall (WAF). All traffic to the web app is routed through an Azure Application
Gateway instance that is used by multiple web apps. The web app address is contoso.azurewebsites.net.
All traffic must be secured with SSL. The Azure Application Gateway instance is used by multiple web apps.
You need to configure the Azure Application Gateway for the web app.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting.

B. Convert the web app to run in an Azure App service environment (ASE).

C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application Gateway.

D. In the Azure Application Gateway's HTTP setting, set the value of the Override backend path option to contoso22.azurewebsites.net.

Correct Answer: *AD*
D: The ability to specify a host override is defined in the HTTP settings and can be applied to any back-end pool during rule creation.
The ability to derive the host name from the IP or FQDN of the back-end pool members. HTTP settings also provide an option to dynamically pick the host name from a back-end pool member's FQDN if configured with the option to derive host name from an individual back-end pool member.
A (not C): SSL termination and end to end SSL with multi-tenant services.
In case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway.
Therefore, there is no need to add any authentication certificates.

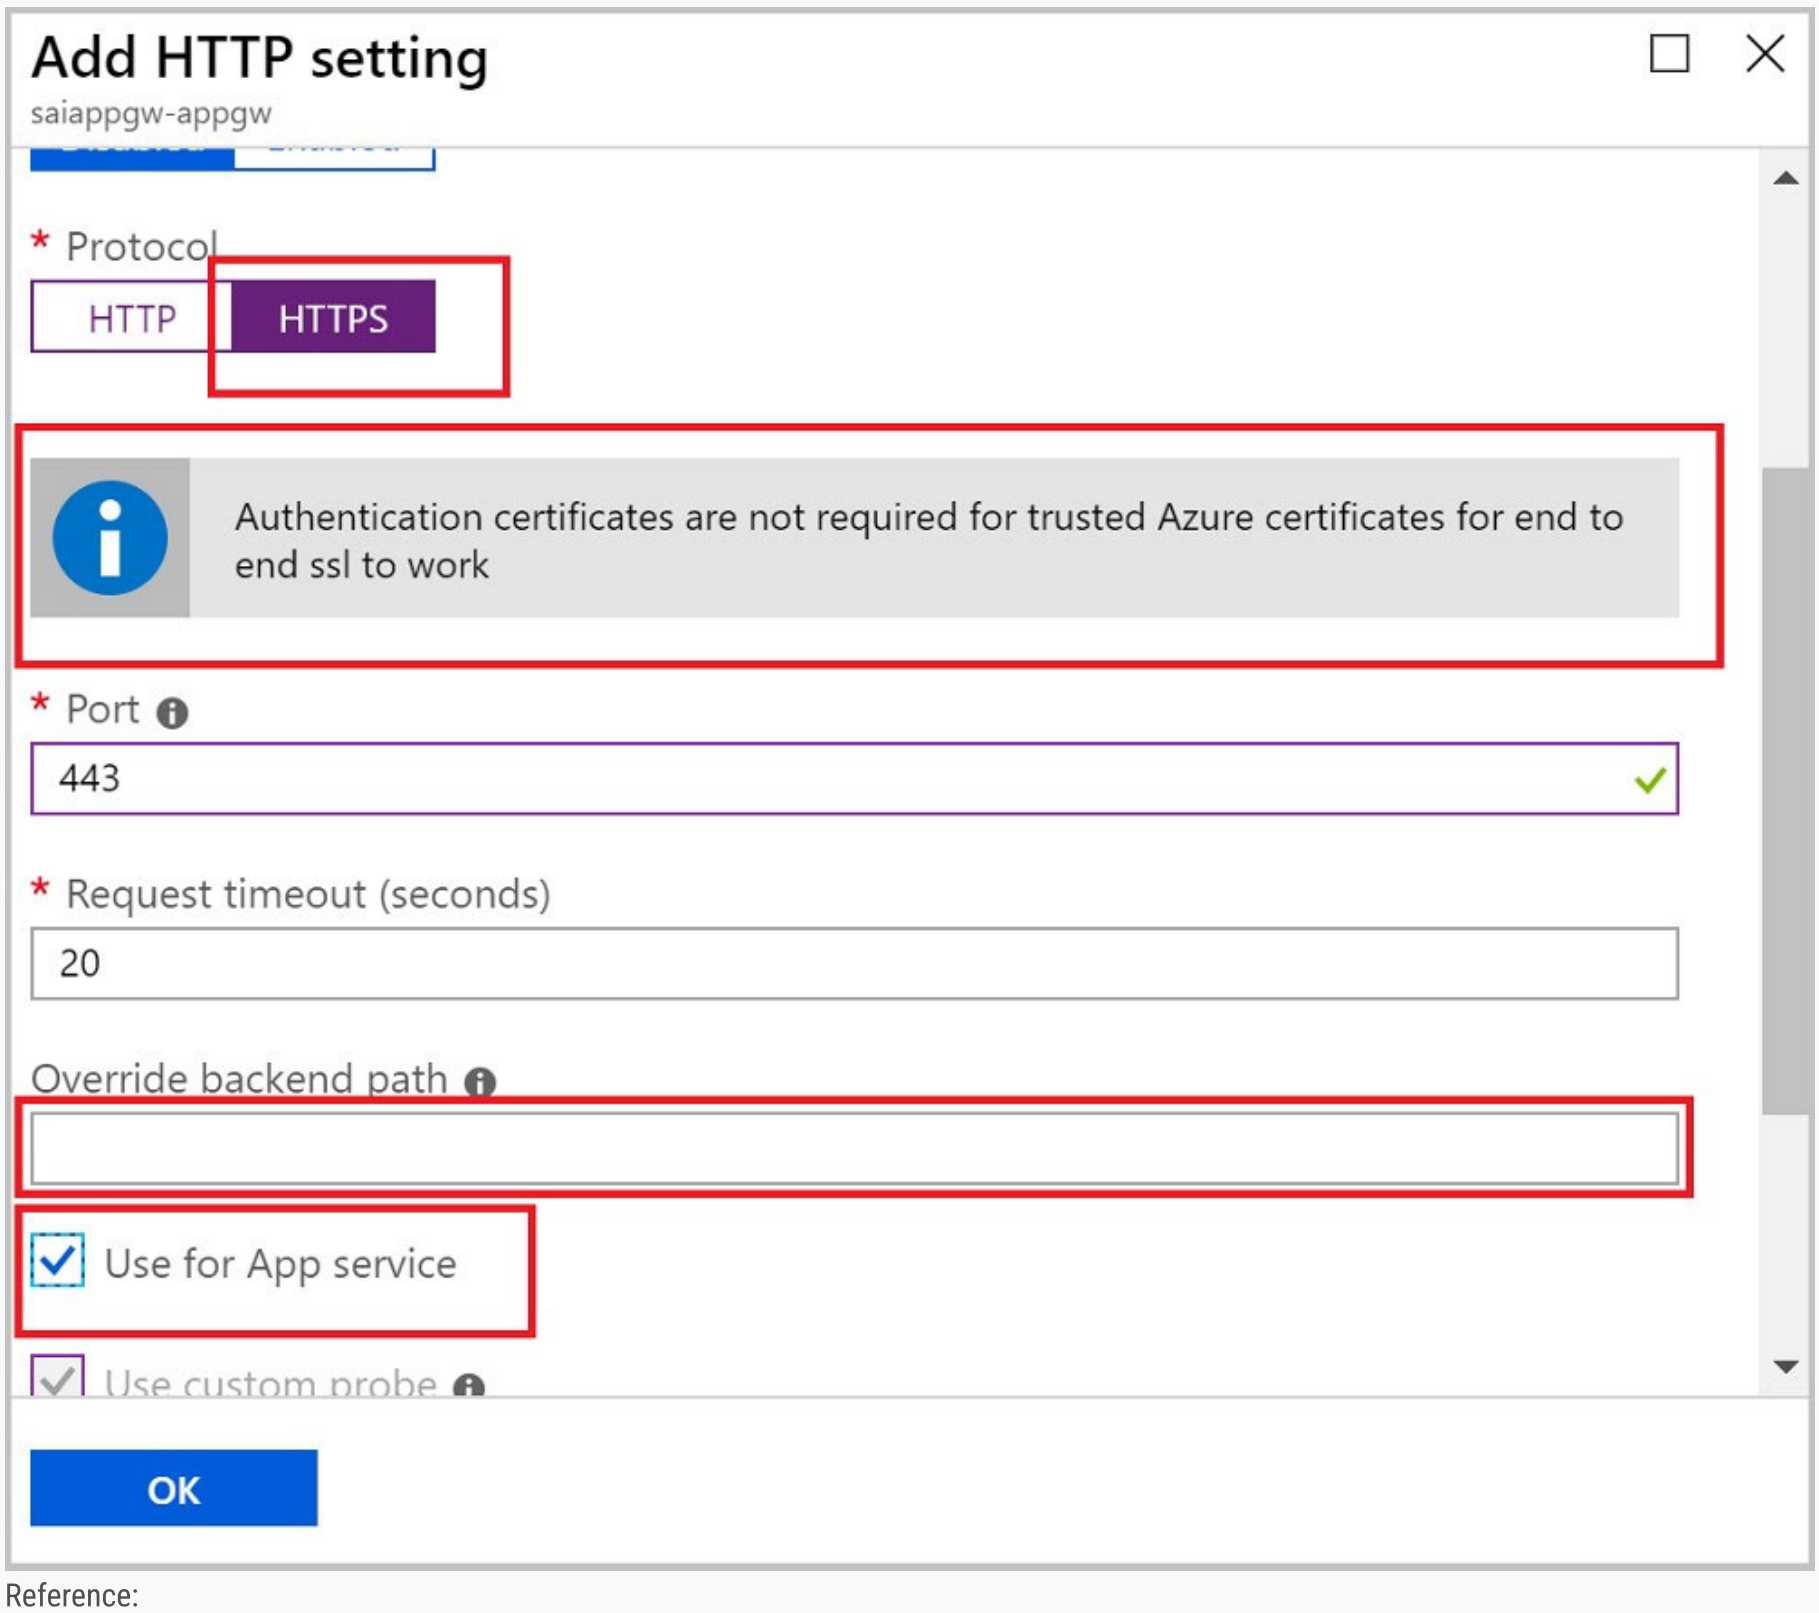

Reference:
https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-web-app-overview

Community vote distribution

AD (64%) AC (36%)

dirtygooback Highly Voted 2 years, 11 months ago

A&D are correct. The question was not copied correctly from AZ-203

https://www.examtopics.com/discussions/microsoft/view/21243-exam-az-203-topic-24-question-22-discussion/

upvoted 47 times

[Removed] 3 months, 3 weeks ago

No. AC is correct

upvoted 1 times

Bartimaeus 1 year, 8 months ago

Nope, there isn't anything about App Service - you can't use it with normal web app.
C & D - you need both

upvoted 1 times

clarionprogrammer 2 years, 6 months ago

Agreed A & D are correct.
Note: [T]here is no requirement to add authentication certificates when App service is selected as backend.
https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-web-app-overview#tls-termination-and-end-to-end-tls-with-multi-tenant-services

upvoted 6 times

samraw83 1 year, 4 months ago

Quite ood this should be AZ-104 question, It should be C & D, you have to update the secure certificate and have to update the backend pools as updated here:
https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-linux

upvoted 1 times

mrqc Highly Voted 2 years, 9 months ago

Is this a valid question for AZ204? I don't see this topic in the skills measured?

upvoted 15 times

NicolasMous Most Recent 1 month, 2 weeks ago

Selected Answer: AC

A, C, Bings exaplanation: The correct answers are:

A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting. C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application Gateway.

Explanation:

Option A: The Azure Application Gateway needs to be configured to use the App service setting. This is because the traffic is being routed through an Azure Application Gateway instance that is used by multiple web apps1.
Option C: Since all traffic must be secured with SSL, an authentication certificate for contoso.azurewebsites.net needs to be added to the Azure Application Gateway1.
Option B is not correct because converting the web app to run in an Azure App service environment (ASE) is not necessary for configuring the Azure Application Gateway for the web app1.

Option D is not correct because changing the backend path to a different web app (contoso22.azurewebsites.net) does not align with the requirement of routing traffic to contoso.azurewebsites.net1.

upvoted 1 times

Vladimir_Gajinov 1 month, 3 weeks ago

Selected Answer: AC

The right answers are A and C.
To configure the Azure Application Gateway for the web app hosted at contoso.azurewebsites.net and ensure that all traffic is secured with SSL, you should perform the following actions:

A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting.
This setting ensures that the Azure Application Gateway can communicate with the Azure App Service (web app) using the appropriate settings and protocols.

C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application Gateway.
This step is essential to enable SSL termination at the Azure Application Gateway. You need to upload and configure the SSL certificate for contoso.azurewebsites.net so that the gateway can securely handle HTTPS traffic for your web app.

The other options, B and D, are not necessary for securing traffic with SSL.

upvoted 1 times

MikeAWS 4 weeks ago

is this explanation from ChatGPT?

upvoted 1 times

tarek0811 3 months ago

it should be A and C, its in the doc:
The SSL certificate is needed to secure the traffic to the web app by encrypting the data between the client and the server.

If you do not add an SSL certificate to the Azure Application Gateway, then all traffic to the web app will be unencrypted. This means that anyone who can intercept the traffic can see the data that is being transmitted, including passwords, credit card numbers, and other sensitive information.

https://learn.microsoft.com/en-us/azure/application-gateway/end-to-end-ssl-portal

upvoted 1 times

[Removed] 3 months, 3 weeks ago

D is definitely incorrect

upvoted 1 times

Mark16 4 months, 1 week ago

Selected Answer: AC

Should be AC:

The correct actions to configure the Azure Application Gateway for the web app are:

A. In the Azure Application Gateway's HTTP setting, enable the "Use for App service" setting. This ensures that the Application Gateway is configured to work properly with the web app running in Azure App Service.

C. Add an authentication certificate for "contoso.azurewebsites.net" to the Azure Application Gateway. This certificate is necessary to enable SSL/TLS encryption for securing the traffic between the clients and the Application Gateway.

Therefore, the correct options are A and C.

upvoted 2 times

Mark16 4 months, 1 week ago

Also for the people that answered D:
Option D, which suggests setting the value of the "Override backend path" option to "contoso22.azurewebsites.net," is not necessary or relevant in this scenario. The goal is to configure the Azure Application Gateway to secure the traffic for the web app hosted at "contoso.azurewebsites.net," not "contoso22.azurewebsites.net." Therefore, option D is not a correct action for this particular requirement.

upvoted 2 times

uffuchsi 8 months, 1 week ago

Received this in my exam today (22/02/2023). Selected A and D. Score 927.

upvoted 3 times

jrnjerin 10 months, 2 weeks ago

Nowhere it is mentiones contoso22

upvoted 2 times

hubekpeter 11 months, 2 weeks ago

Selected Answer: AD

A. yes, app is using default appservice domain azurewebsites.net, so my assumption is, it's hosted as appservice, B - it's not needed as it's already running as App Service, therefore C - is not needed as well, D - yes you need either pool or redirection, in this case we need to redirect the traffic, there's 100% typo in hostname

upvoted 1 times

hubekpeter 11 months, 2 weeks ago

Here's an article where the answer D is explained in detail. https://learn.microsoft.com/en-us/azure/application-gateway/configure-web-app?tabs=customdomain%2Cazure-portal

upvoted 3 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: AD

A&D are correct.

upvoted 1 times

coffecold 1 year ago

D is incorrect by all means here : nowhere a override backend path is requested, contoso22 is never mentioned anywhere.
I think A (Https for backend) and C (https from frontend to service) is correct here.

upvoted 3 times

systherm 1 year, 7 months ago

Selected Answer: AD

correct

upvoted 1 times

avipi 1 year, 10 months ago

Selected Answer: AD

correct

upvoted 1 times

sawipef270 1 year, 10 months ago

Selected Answer: AD

correct

upvoted 1 times

karthik2123 1 year, 11 months ago

Selected Answer: AD

A&D are correct.

upvoted 2 times

MohmmadFayez 2 years ago

B & C are the correct answers

D is wrong because is telling override back-end path but we need to override the host header and since no option to override host hear are included at the answer we should select B ASE
A is wrong because we are going to use ASE which need a SLL certificate to be provided to the App gateway.

upvoted 2 times

Bartimaeus 1 year, 7 months ago

Yeah, that's actually only possible combination.
D - override backend path is used to redirect from 1 path to 2nd, e.g.
app-gateway.com/-> contoso.azurewebsites.net/images
would require to set override path to `images`.
It can't be hostname - it would just result in
contoso.azurewebsites.net/contoso.azurewebsites.net

So if D is invalid, then we choose between A,B,C.
AC -> wrong, if you use Use for App service setting, you don't need the cert
AB -> wrong, you can't use A setting with ASE environment, it requires its own cert.
-> BC.

upvoted 1 times

Question #38 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure
Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Use the Azure Blob Storage change feed to trigger photo processing.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
The change feed is a log of changes that are organized into hourly segments but appended to and updated every few minutes. These segments are created only when there are blob change events that occur in that hour.
Instead catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

Community vote distribution
B (100%)

finnishr Highly Voted 2 years, 2 months ago
The answer is 100% correct.

Change feed publishes records to the log within few minutes of the change. The process to produce a mobile-friendly version of the image must start in less than one minute.

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal
upvoted 21 times

Vladimir_Gajinov Most Recent 2 months, 2 weeks ago
Selected Answer: B
Simply NO!
upvoted 1 times

Hogwarts 6 months, 4 weeks ago
Got this in exam today (5 April 2023)
upvoted 3 times

chirandev 7 months, 1 week ago
Selected Answer: B
Change feed is for logging.
upvoted 2 times

Esward 9 months ago
Given answer is correct!
upvoted 1 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: B
correct.
upvoted 1 times

gmishra88 1 year ago
There is blob storage events through event grid and the azure-function-triggers as alternatives for faster processing
upvoted 1 times

gmishra88 1 year ago

https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal#should-i-use-the-change-feed-or-storage-events

upvoted 1 times

Ganesh1983 1 year, 5 months ago

Selected Answer: B

change feed is more to do with audit requirements.

upvoted 1 times

Freidrich 1 year, 8 months ago

Selected Answer: B

The correct answer is B: No.

upvoted 1 times

edengoforit 1 year, 9 months ago

same question here and the answer is No according to the question below
https://www.examtopics.com/discussions/microsoft/view/26699-exam-az-204-topic-1-question-19-discussion/

upvoted 2 times

r__k 1 year, 11 months ago

Important info related to this question. https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp

upvoted 1 times

ucsdmiami2020 1 year, 11 months ago

Selected Answer: B

The answer is correct. Per the provided referenced URL
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal
"Change feed provides a durable, ordered log model of the changes that occur to a blob. Changes are written and made available in your change feed log within an order of a few minutes of the change. If your application has to react to events much quicker than this, consider using Blob Storage events instead. Blob Storage Events provides real-time one-time events which enable your Azure Functions or applications to quickly react to changes that occur to a blob."

upvoted 2 times

nhontran 1 year, 11 months ago

I guess the answer is 'No', requirement is less than one minute, change feed takes more than that.
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal
"The change feed publishes records to the log within few minutes of the change and also guarantees the order of change operations per blob. Storage events are pushed in real time and might not be ordered."

upvoted 2 times

CodePoet 2 years, 2 months ago

The answer is correct

upvoted 3 times

Amit_kk 2 years, 3 months ago

It Should be Yes

upvoted 2 times

nombuso 2 years, 3 months ago

why yes?

upvoted 1 times

Question #39 *Topic 2*

You are developing a web application that runs as an Azure Web App. The web application stores data in Azure SQL Database and stores files in an Azure
Storage account. The web application makes HTTP requests to external services as part of normal operations.
The web application is instrumented with Application Insights. The external services are OpenTelemetry compliant.
You need to ensure that the customer ID of the signed in user is associated with all operations throughout the overall system.
What should you do?

A. Add the customer ID for the signed in user to the CorrelationContext in the web application

B. On the current SpanContext, set the TraceId to the customer ID for the signed in user

C. Set the header Ocp-Apim-Trace to the customer ID for the signed in user

D. Create a new SpanContext with the TraceFlags value set to the customer ID for the signed in user

Correct Answer: *A*
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/correlation

Community vote distribution
A (81%) B (19%)

willchenxa Highly Voted 1 year, 1 month ago

Selected Answer: A
A SpanContext represents the portion of a Span which must be serialized and propagated along side of a Baggage.
https://open-telemetry.github.io/opentelemetry-js-api/interfaces/spancontext.html
upvoted 5 times

warchoon 8 months, 1 week ago
Agree https://learn.microsoft.com/en-us/azure/azure-monitor/app/opentelemetry-enable?tabs=java#set-the-user-id-or-authenticated-user-id
upvoted 1 times

applepie Most Recent 3 months ago
got this question today, answer A - 7/30/2023, score 895/1000
upvoted 3 times

imanonion 9 months, 3 weeks ago

Selected Answer: A
Answer by ChatGPT, which I think looks legit:
"I would choose option A, Add the customer ID for the signed in user to the CorrelationContext in the web application.

The CorrelationContext is a way to associate contextual information with a request as it flows through the system. It allows you to track a request as it passes through different components of the system, and to identify related log entries and telemetry data. By adding the customer ID to the CorrelationContext in the web application, you can ensure that it is associated with all operations throughout the overall system. This will allow you to track the request and identify related log entries and telemetry data for a specific customer.
upvoted 4 times

imanonion 9 months, 3 weeks ago
Option B, On the current SpanContext, set the TraceId to the customer ID for the signed in user, is not correct because the TraceId is a unique identifier for a trace, and it is not meant to be used to store contextual information about the request.

Option C, Set the header Ocp-Apim-Trace to the customer ID for the signed in user, is not correct because the Ocp-Apim-Trace header is used by Azure API Management to track requests as they pass through the API gateway, and it is not related to application instrumentation or tracing.

Option D, Create a new SpanContext with the TraceFlags value set to the customer ID for the signed in user, is not correct because the TraceFlags value is used to indicate the state of the trace and is not meant to be used to store contextual information about the request."
upvoted 5 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: A
A. Add the customer ID for the signed in user to the CorrelationContext in the web application
upvoted 2 times

Akiu 1 year ago
OpenTelemetry uses correlation context and it fits the requirements:
Correlation context carries user-defined properties. These properties are typically data that you would like to eventually aggregate for correlation

analysis or use to filter your trace data, such as a customer identifier, ...
https://lightstep.com/opentelemetry/context-propagation

upvoted 3 times

chessace2000 1 year ago

Selected Answer: B

It has to be OpenTelemetry compliant. This means using W3C correlation spec. CorrelationContext (A) is legacy correlation spec not W3C.

upvoted 2 times

warchoon 8 months, 1 week ago

tracestate: Carries system-specific tracing context.
It is a part of correlation headers using W3C TraceContext.
Correlation-Context is deprecated. But this is the same as the Request-Context which is currently in use.

upvoted 1 times

finnishr 1 year, 1 month ago

Selected Answer: B

The correct answer is B.

upvoted 1 times

DonH 4 months, 2 weeks ago

This answer is useless without any explanation. Come on guys this is annoying and NOT helpfull. Please clarify WHY it is the correct answer. Now we have a group of people stating A is correct and an other group stating B is correct. How are we suppose to know what is correct?

upvoted 3 times

lhy904 1 year, 5 months ago

Why not B?

upvoted 2 times

warchoon 8 months, 1 week ago

"The value of trace-id SHOULD be globally unique. One recommended method to ensure global uniqueness, as well as to address some privacy and security considerations, to a satisfactory degree of certainty is to randomly (or pseudo-randomly) generate the trace-id. Implementers SHOULD use a trace-id generation method which randomly (or pseudo-randomly) generates at least the right-most 7 bytes of the ID." by specs
ChatGPT is lying (as usual) that you can't use it to request, but you can't use TraceId with all operations throughout the overall system

upvoted 2 times

nt_ar 1 year, 6 months ago

Selected Answer: A

That's correct

upvoted 2 times

Question #40 *Topic 2*

HOTSPOT -

You are developing an Azure Function App. You develop code by using a language that is not supported by the Azure Function App host. The code language supports HTTP primitives.

You must deploy the code to a production Azure Function App environment.

You need to configure the app for deployment.

Which configuration values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Configuration parameter	Configuration value
Publish	Code / Docker Container
Runtime stack	Node.js / Python / PowerShell Core / Custom Handler
Version	14 LTS / 7.0 / custom

Correct Answer:

Answer Area

Configuration parameter	Configuration value
Publish	Code / **Docker Container**
Runtime stack	Node.js / Python / **PowerShell Core** / Custom Handler
Version	14 LTS / **7.0** / custom

Box 1: Docker container -
A custom handler can be deployed to every Azure Functions hosting option. If your handler requires operating system or platform dependencies (such as a language runtime), you may need to use a custom container. You can create and deploy your code to Azure Functions as a custom Docker container.

Box 2: PowerShell core -
When creating a function app in Azure for custom handlers, we recommend you select .NET Core as the stack. A "Custom" stack for custom handlers will be added in the future.
PowerShell Core (PSC) is based on the new .NET Core runtime.

Box 3: 7.0 -
On Windows: The Azure Az PowerShell module is also supported for use with PowerShell 5.1 on Windows.
On Linux: PowerShell 7.0.6 LTS, PowerShell 7.1.3, or higher is the recommended version of PowerShell for use with the Azure Az PowerShell module on all platforms.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-function-linux-custom-image https://docs.microsoft.com/en-us/powershell/azure/install-az-ps?view=azps-7.1.0

SaintBahamut Highly Voted 1 year, 6 months ago
answer is wrong: if you pick docker container you cant specify stack and version,
So ill go with
>Code
>Custom Handler
>custom (only option when you pick Custom Handler)
upvoted 74 times

mihailos 2 months ago
Got it in exam on 28/08/23. Went with this answer. Scored 912
upvoted 2 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000
upvoted 8 times

adilkhan 7 months, 1 week ago
all questions were from exam topics?
upvoted 2 times

Esward 9 months ago

Agreed
upvoted 1 times

JustHereToLearn 10 months, 3 weeks ago
To further support this answer see this video tutorial from Microsoft official channel where they also select "Custom Handler" and do not use Docker: https://www.youtube.com/watch?v=roWbjuo-c30
upvoted 2 times

EricCao Highly Voted 1 year ago
I checked it on Azure portal and the answer should be
Code > Custom Handler > custom
upvoted 14 times

warchoon 8 months, 1 week ago
Did you use a language that is NOT supported by the Azure Function App host?
upvoted 2 times

JamieS 9 months ago
Just tried it to, he's right., Code > Custom Handler > Custom
upvoted 2 times

richardc7 Most Recent 1 month ago
got in the exam 28/09/2023
upvoted 1 times

RuffBoii 1 month, 1 week ago
Had this on my exam today.
upvoted 1 times

nikipediaa 1 month ago
me too, today
upvoted 1 times

Tarajee 1 month, 2 weeks ago
Got this on exam 2023sept
upvoted 2 times

nlk0 2 months ago
Got this today 26/08/2023, went with the most voted. 880
Case study was contoso -> 9 questions
+ 35 question, two series.
3 of them were not from examptopics, all others were
upvoted 3 times

kvtborad 2 months, 3 weeks ago
I got this question on 6th August 2023. chose highly voted. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.
Code, Custom Handler, Custom
upvoted 3 times

macobuzi 2 months, 2 weeks ago
Thank you for informing us! Glad to hear questions are from here.
upvoted 2 times

juanckar 3 months, 3 weeks ago
This was on the exam (July 2023). Went with highly voted. Scored 917
upvoted 2 times

MysticalSam 4 months, 3 weeks ago
This question was in today's exam at 10-June-2023
upvoted 3 times

red0400 7 months, 2 weeks ago
Answer is -
Code
Custom Handler
Custom
upvoted 1 times

sivaj 7 months, 3 weeks ago
A custom handler can be deployed to every Azure Functions hosting option. If your handler requires operating system or platform dependencies (such as a language runtime), you may need to use a custom container.

When creating a function app in Azure for custom handlers, we recommend you select .NET Core as the stack. A "Custom" stack for custom handlers will be added in the future.
The given answer is correct
upvoted 1 times

Kluk_Kluk 7 months, 3 weeks ago

Custom handlers are lightweight web servers that receive events from the Azure Functions host. Any language that supports HTTP primitives can implement a custom handler. https://learn.microsoft.com/en-us/azure/azure-functions/supported-languages

upvoted 1 times

Defa 7 months, 3 weeks ago

Going for:
> Code
> Custom
> Custom

upvoted 1 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23
score: 817

upvoted 1 times

Priya0703 8 months, 1 week ago

Got this question on exam today 20-02-2023

upvoted 1 times

NikolaPops 10 months ago

I think Docker is correct because you can easily have updated versions of your runtime stack using pre-built docker images. But I'd say its Docker/custom/custom

upvoted 1 times

warchoon 8 months, 1 week ago

It's not logical, but MS recommends using Core
https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers#deploying

upvoted 1 times

hubekpeter 11 months, 2 weeks ago

Code/Custom Handler/Custom - Custom handlers are lightweight web servers that receive events from the Azure Functions host. Any language that supports HTTP primitives can implement a custom handler. This means that custom handlers can be used to create functions in languages that aren't officially supported.

upvoted 2 times

Question #41 *Topic 2*

DRAG DROP -

You provision virtual machines (VMs) as development environments.

One VM does not start. The VM is stuck in a Windows update process. You attach the OS disk for the affected VM to a recovery VM.

You need to correct the issue.

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

Select and Place:

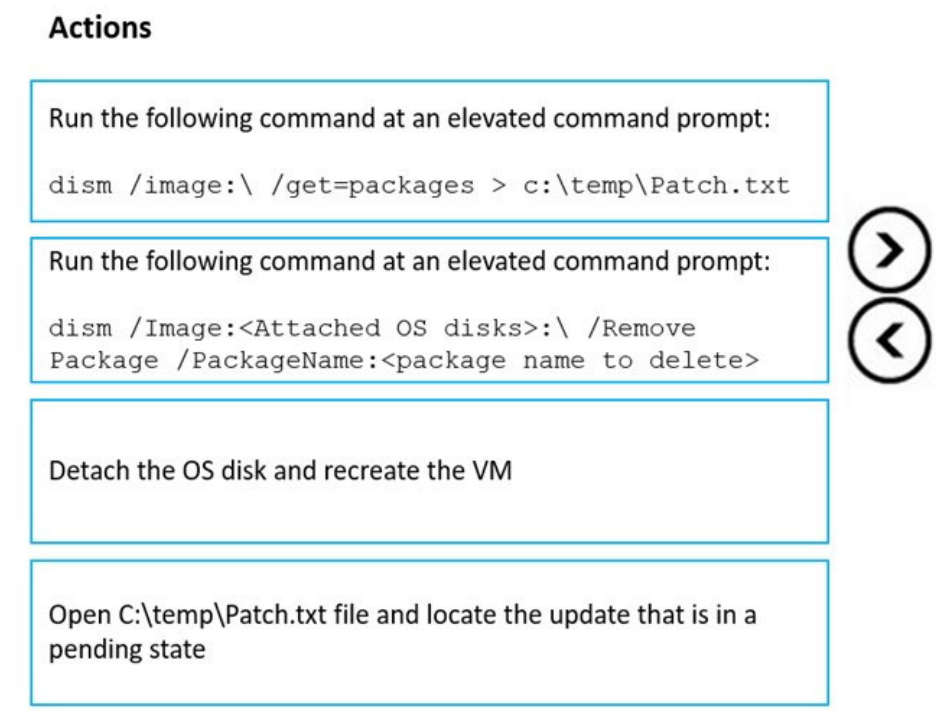

Correct Answer:

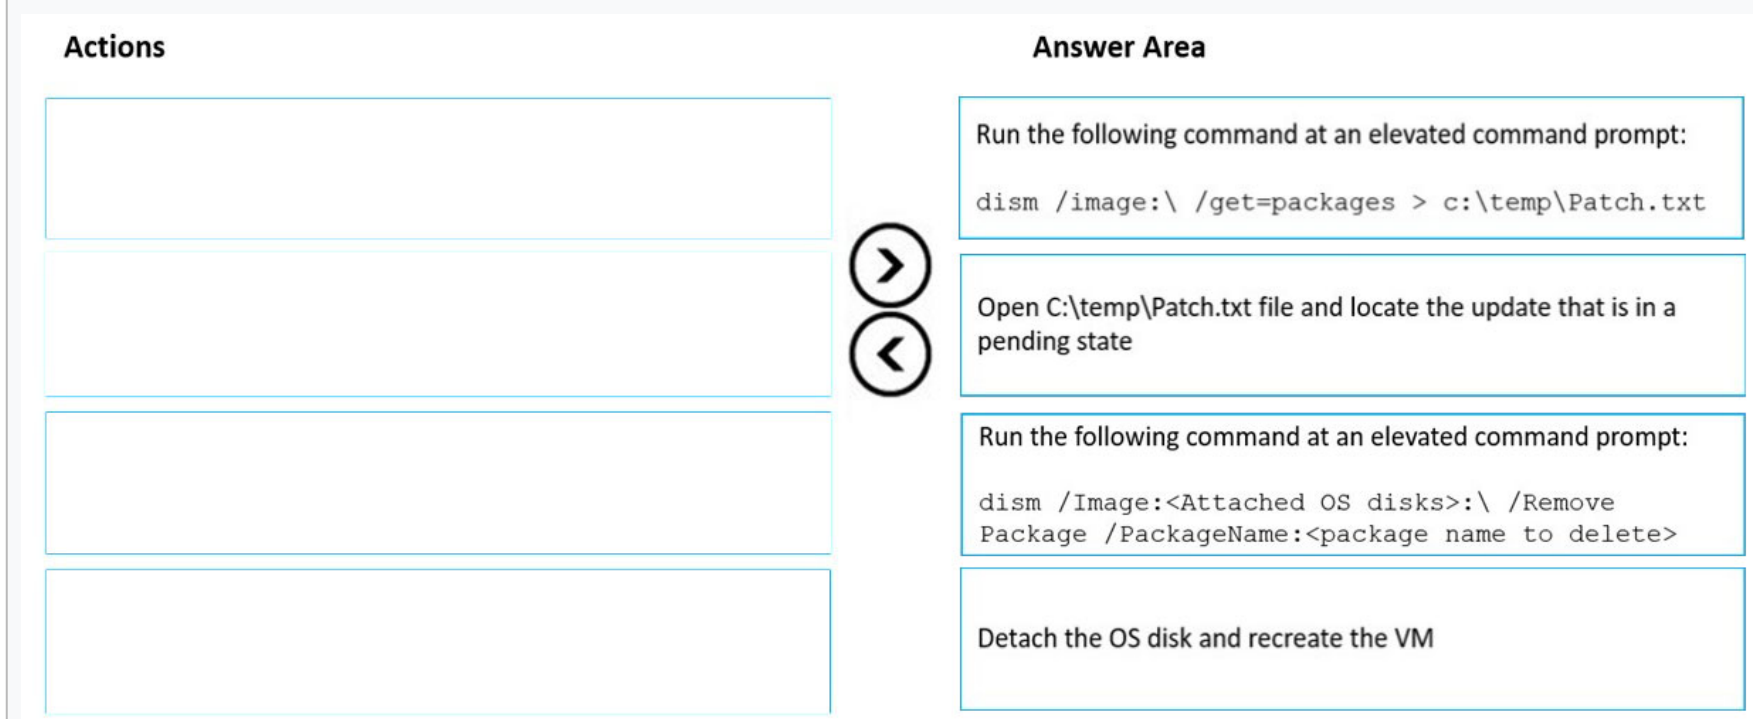

Remove the update that causes the problem

1. Take a snapshot of the OS disk of the affected VM as a backup.
2. Attach the OS disk to a recovery VM.
3. Once the OS disk is attached on the recovery VM, run diskmgmt.msc to open Disk Management, and ensure the attached disk is ONLINE.
4. (Step 1) Open an elevated command prompt instance (Run as administrator). Run the following command to get the list of the update packages that are on the attached OS disk: dism /image:<Attached OS disk>:\ /get-packages > c:\temp\Patch_level
5. (Step 2) Open the C:\temp\Patch_level.txt file, and then read it from the bottom up. Locate the update that's in Install Pending or Uninstall Pending state.
6. Remove the update that caused the problem:

dism /Image:<Attached OS disk>:\ /Remove-Package /PackageName:<PACK

7. (Step 4) Detach the OS disk and recreate the VM. Then check whether the issue is resolved.

Reference:

https://docs.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshoot-stuck-updating-boot-error

nedlo Highly Voted 7 months, 2 weeks ago

sure i will use this knowledge for my developer work...

upvoted 26 times

85e4c91 2 weeks, 5 days ago

I guess developers don't have access to google nowadays

upvoted 1 times

br00net 5 months ago

clearly, every developer needs to memorize such knowledge instead of searching in the documentation

upvoted 3 times

grudzin 3 months, 4 weeks ago

Question says "The VM is stuck in a Windows update process". I ques it's common ;)

upvoted 4 times

Magis Highly Voted 11 months, 4 weeks ago

Correct.

All steps are listed here - https://docs.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshoot-stuck-updating-boot-error

upvoted 10 times

Lucifer14 Most Recent 10 months ago

In 28-12-2022 exam

upvoted 3 times

memadk 1 year ago

Is this a part of AZ-204? Then I must have missed something....

upvoted 3 times

martijn44 1 year ago

I doubt it.. it seems many of these questions were copied over from AZ-203.

upvoted 1 times

finnishr 1 year, 1 month ago

100% correct!!

upvoted 4 times

SoftSol 1 year, 6 months ago

Correct

upvoted 2 times

Question #42 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data.
You need to ensure the app does not time out and processes the blob data.
Solution: Update the functionTimeout property of the host.json project file to 10 minutes.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:
Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the
HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices

Community vote distribution
B (100%)

Ganesh1983 Highly Voted 1 year, 5 months ago
Selected Answer: B
HTTP triggers on function apps can timeout in 230 seconds. Hence B is the answer.
upvoted 5 times

surprise0011 6 months, 2 weeks ago
received 2023-04-17 went given answer, score 926
upvoted 2 times

Nitin23 Most Recent 1 month, 2 weeks ago
Answer is A as per chatgpt.
Its a valid solution
upvoted 1 times

MDWPartners 1 month, 1 week ago
Do not use chatGPT as a source, is really faulty with this kind of questions
upvoted 2 times

uffuchsi 8 months, 1 week ago
Received this in my exam today (22/02/2023). Selected B. Score 927.
upvoted 2 times

Esward 9 months ago
B is correct answer
https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout
upvoted 1 times

OPT_001122 11 months ago
Selected Answer: B
https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices
upvoted 1 times

stfnmrr 1 year, 1 month ago
Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout

upvoted 4 times

sghaha 1 year, 6 months ago

Why not a?

upvoted 1 times

Question #43 *Topic 2*

HOTSPOT -

You are developing an Azure Durable Function based application that processes a list of input values. The application is monitored using a console application that retrieves JSON data from an Azure Function diagnostic endpoint.

During processing a single instance of invalid input does not cause the function to fail. Invalid input must be available to the monitoring application.

You need to implement the Azure Durable Function and the monitoring console application.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
[FunctionName("App")]
public static async Task<List<string>> RunOrchestrator(
  [OrchestrationTrigger] IDurableOrchestrationContext context) {
  EntityId[] input = . . .
  int errIndex =. . .
    [                                                    ▼ ];
}
  using (var client = new HttpClient())
  {
    while (true)
    {
      var response = await client.GetAsync(". . .");
      response.EnsureSuccessStatusCode();
      var json = await response.Content.ReadAsStringAsync();
      dynamic result = JsonConvert.DeserializeObject(json);
      if (result.runtimeStatus == "                    ▼ ")
        {
          return result.                                 ▼ ;
        }
    }
}
```

- context.SetOutput(input[errIndex])
- context.SetCustomStatus(input[errIndex])
- context.SignalEntity(input[errIndex], "error")
- await context.CallEntityAsync(input[errIndex], "error")

- Failed
- Awaited
- Listening
- Completed

- input
- output
- runtimeStatus
- customStatus

Answer Area

Correct Answer:

Box 1: await context.CallEntityAsync(input[errindex],"error")

Orchestration signals and calls an entity

Orchestrator functions can access entities by using APIs on the orchestration trigger binding.

Example:

```
[FunctionName("CounterOrchestration")]
public static async Task Run(
[OrchestrationTrigger] IDurableOrchestrationContext context)
{
var entityId = new EntityId(nameof(Counter), "myCounter");
// Two-way call to the entity which returns a value - awaits the response int currentValue = await context.CallEntityAsync<int>(entityId, "Get");
```

Box 2: Failed -

During processing a single instance of invalid input does not cause the function to fail.

Note: RuntimeStatus: One of the following values:

Failed: The instance failed with an error.

Completed: The instance has completed normally.

Terminated: The instance was stopped abruptly.

Pending: The instance has been scheduled but has not yet started running.

Running: The instance has started running.

ContinuedAsNew: The instance has restarted itself with a new history. This state is a transient state.

Box 3: Input -

Invalid input must be available to the monitoring application.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-entities https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-instance-management

RochaG2 Highly Voted 1 year, 2 months ago
I've not seen any of this stuff in Microsoft's Learning Path's for AZ-204
upvoted 26 times

gmishra88 1 year ago
The learning path and the preparation videos are showing things you need not know for the exam.
upvoted 16 times

OPT_001122 1 year, 1 month ago
yes. there are a few questions added in last month. not sure if all of them are of AZ-204
upvoted 2 times

Akiu Highly Voted 1 year ago
I'd say:
SetOutput()
== Completed
result.output
The function needs to return the one invalid input, so you could use either SetOutput or SetCustomStatus, but custom status is not available after the function completes.
The one invalid input does not cause the function to fail, therefore Completed.
upvoted 18 times

regux 2 weeks, 2 days ago
async functions needs an await part. SetOutput() won't qualify for that as function.
upvoted 1 times

macobuzi 2 months, 1 week ago
This one seems right!
upvoted 1 times

POOOJAAAAAAAAAA 10 months, 1 week ago
is this correct ?
upvoted 4 times

longnguyendh Most Recent 3 months ago
I think that the anwser is
SetCustomStatus
Failed
result.customStatus
upvoted 2 times

MysticalSam 4 months, 3 weeks ago
This question was in today's exam at 10-June-2023
upvoted 7 times

NavinD 3 months, 3 weeks ago
How much percent questions did you get from here? Do you have contributor access?
upvoted 1 times

Hogwarts 6 months, 4 weeks ago
Got this in exam today (5 April 2023)
upvoted 4 times

PP2015 7 months ago
Did my exam on 3/29/2023. This question was on it. Selected context.Signal(input[errIndex], "error"), Completed, output. Score 850
upvoted 2 times

red0400 7 months, 2 weeks ago
Answer - setOutput, completed, output
upvoted 5 times

pkolodziej 7 months, 3 weeks ago
For those who are wondering if it's a SignalEntity, it's not:

SignalEntity(EntityId, DateTime, String, Object)
Signals an operation to be performed by an entity at a specified time. Any result or exception is ignored (fire and forget).

We're clearly waiting for completion here, so have to use CallEntity,

upvoted 1 times

pkolodziej 7 months, 3 weeks ago

I correct my mistake: it should be CustomStatus

https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-custom-orchestration-status?tabs=csharp#querying-custom-status-with-http

upvoted 3 times

uffuchsi 8 months, 1 week ago

Received this in my exam today (22/02/2023). Selected context.Signal(input[errIndex], "error"), Completed, output. Score 927.

upvoted 9 times

memobed890 7 months, 2 weeks ago

De donde saco la información para decir que son esas respuestas?

upvoted 4 times

RealRaymond 8 months, 2 weeks ago

Answer should be CustomStatus.

Ref: https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-custom-orchestration-status?tabs=csharp#querying-custom-status-with-http

upvoted 2 times

manopeydakon 8 months, 3 weeks ago

Nigros the answer is correct since we are after the failed input values.

upvoted 3 times

dimsok 8 months, 3 weeks ago

Seem
s correct, but the code in the answer area is missing stuff

upvoted 1 times

Esward 9 months ago

I think, given answers are correct!

upvoted 1 times

TaoLu 9 months, 1 week ago

The answer is correct, guys!

upvoted 1 times

deepbh 10 months ago

Looks to me like: SetOutput() / == Failed / result.output
I agree on using output instead of customresult, same as Akiu's comment, but disagree on their rationale about using Completed in box2. It's the job of the code we're writing to make sure we don't fail at invalid input. We will encounter a Failed status for any such invalid inputs and if we do, we make sure return the saved info (which we know now to be invalid) to the monitoring application.

Or I might be wrong, who knows. AZ-204 is a dark art.

upvoted 2 times

edmond 10 months, 3 weeks ago

SetCustomStatus
Failed
result.customStatus

Durable Functions doesn't have SetOutput so it shouldn't be output.
Input[errorIndex] is simply an EntityId, which means it doesn't make sense to call error on it.

upvoted 6 times

edmond 10 months, 2 weeks ago

Ignore my last comment, I think the correct answer should be:
SetOutput()
== Completed
result.output

Please refer to Akiu's comment

upvoted 2 times

sabbebabbe 12 months ago

I actually had this question last week when I did the az-204 exam

upvoted 3 times

[Removed] 3 months, 3 weeks ago

And what? What did you select? Otherwise your comment is absolutely useless

upvoted 2 times

AjeyAkshay 11 months, 2 weeks ago

Which options did you select?

upvoted 1 times

Question #44 *Topic 2*

You are developing an Azure Durable Function to manage an online ordering process.
The process must call an external API to gather product discount information.
You need to implement the Azure Durable Function.
Which Azure Durable Function types should you use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Orchestrator

B. Entity

C. Client

D. Activity

Correct Answer: *AB*
The Durable Functions extension exposes a set of built-in HTTP APIs that can be used to perform management tasks on orchestrations, entities, and task hubs.
These HTTP APIs are extensibility webhooks that are authorized by the Azure Functions host but handled directly by the Durable Functions extension.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-http-api

Community vote distribution
AD (92%) 8%

serpevi Highly Voted 1 year, 1 month ago
Got this in 09/22 , went with Orchestrator and Entity, score 927.
upvoted 18 times

[Removed] 9 months, 3 weeks ago
Entity is for storing state which isn't called out as requirement. An activity can return/work with the HTTP req. Like mentioned before this, Dur Func 2.0 supports DurableHttpResponse. (Grats/Great score)
upvoted 3 times

OPT_001122 1 year, 1 month ago
Thanks for mentioning the exam date in which it was asked.
upvoted 2 times

Chief2110 Highly Voted 1 year, 1 month ago
It should be orchestrator, client & activity right?
upvoted 16 times

Esward 9 months, 3 weeks ago
orchestrator and activity
upvoted 5 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000
upvoted 4 times

JustHereToLearn 9 months, 1 week ago
Yes I think orchestrator and activity is correct.
Activity for this reason "Unlike orchestrator functions, activity functions aren't restricted in the type of work you can do in them. Activity functions are frequently used to make network calls or run CPU intensive operations. An activity function can also return data back to the orchestrator function."
Reference: https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview#activity-functions
upvoted 2 times

coffecold 1 year ago
yes, see https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview
upvoted 1 times

Kanasan Most Recent 2 months ago
Selected Answer: AD
A and D:
https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview#activity-functions

Activity functions are frequently used to make network calls or run CPU intensive operations. An activity function can also return data back to the orchestrator function.

upvoted 3 times

Kanasan 2 months ago

A and D:
https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview#activity-functions
Activity functions are frequently used to make network calls or run CPU intensive operations. An activity function can also return data back to the orchestrator function.

upvoted 1 times

longnguyendh 3 months ago

It should be orchestrator & activity
Activity Function: This type of function is responsible for performing individual tasks, such as calling an external API to gather product discount information

upvoted 1 times

RayD7 5 months, 1 week ago

Durable functions are out of scope as of 28th April 2023

upvoted 3 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

My case:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

Raja419 6 months ago

Selected Answer: AD

The "Orchestrator" function type is responsible for coordinating the workflow and managing the execution of the "Activity" function type. It receives inputs, calls external APIs, and makes decisions based on the results of the API calls.

The "Activity" function type is responsible for performing specific actions or tasks, such as calling an external API to gather product discount information. It receives inputs from the "Orchestrator" function and returns results back to it.

upvoted 2 times

Saluk_DE 7 months ago

Question was in Exam 2023-03-30

upvoted 2 times

maqsoodshah 7 months ago

Correct Answers:
A. Orchestrator
D. Activity

Activity functions
Activity functions are the basic unit of work in a durable function orchestration. Activity functions are the functions and tasks that are orchestrated in the process. For example, you might create an orchestrator function to process an order. The tasks involve checking the inventory, charging the customer, and creating a shipment. Each task would be a separate activity function. These activity functions may be executed serially, in parallel, or some combination of both.

Entity functions
Entity functions define operations for reading and updating small pieces of state. We often refer to these stateful entities as durable entities. Like orchestrator functions, entity functions are functions with a special trigger type, entity trigger. They can also be invoked from client functions or from orchestrator functions

https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview

upvoted 2 times

adilkhan 7 months, 1 week ago

A,D chatGPT

upvoted 1 times

DonH 4 months, 2 weeks ago

Again stop with this ChatGPT nonsense respons. ChatGPT is NOT a source rather than a search engine. It can also give incorrect answers. So stating its from ChatGPT doesn't say a thing about it's validity.

upvoted 6 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23
score: 817

upvoted 2 times

Priya0703 8 months, 1 week ago

On exam today 20-02-2023
upvoted 2 times

alexein74 9 months, 1 week ago

Selected Answer: AD

A. Orchestrator
B. Activity
An Azure Durable Function that manages an online ordering process should use the Orchestrator type to control the flow of the process and coordinate the activities that need to be executed. The Orchestrator is the main entry point of the Durable Function and it is responsible for managing the state and the execution of the other functions.
Activity functions should be used to perform the external API call to gather product discount information. An activity function is a function that performs a single unit of work and it can be invoked by the Orchestrator.
upvoted 2 times

[Removed] 9 months, 3 weeks ago

All options here are considered durable function types. A client is required to call/kick-off an Orchestration. There's two ways I see here. Durable Functions 2.0 can have orchestration call HTTP without requiring Activity type. In which case AC. Technically you can also do ACD, its unclear as this ultimately depends on the finer details of the code. Entity is for storing state, which isn't called out as a requirement (the discount might be gathered at execution, used and thrown away). Coin flip for me between AC and ACD
upvoted 1 times

deepbh 10 months ago

A (Orchestrator) is always needed. C (Client) is used to call Orchestrator so it's not valid here. B (Entity) is used to manipulate small pieces of information hence also not valid.

Calling the external API should be the task of an Activity function. "Unlike orchestrator functions, activity functions aren't restricted in the type of work you can do in them. Activity functions are frequently used to make network calls or run CPU intensive operations." See https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview#activity-functions

So answer is AD.
upvoted 5 times

edmond 10 months, 2 weeks ago

Selected Answer: AD

Orchestrator and Activity are the required function types here
upvoted 5 times

Question #45 *Topic 2*

DRAG DROP -

You are authoring a set of nested Azure Resource Manager templates to deploy multiple Azure resources.

The templates must be tested before deployment and must follow recommended practices.

You need to validate and test the templates before deployment.

Which tools should you use? To answer, drag the appropriate tools to the correct requirements. Each tool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Tools

- Parameter file
- Template function
- Azure Resource Manager test toolkit
- User-defined function
- What-if operation
- Azure Deployment Manager

Answer Area

Requirement	Tool
Determine whether the templates follow recommended practices.	Tool
Test and validate changes that templates will make to the environment.	Tool

Correct Answer:

Tools

- Parameter file
- Template function
- Azure Resource Manager test toolkit
- User-defined function
- What-if operation
- Azure Deployment Manager

Answer Area

Requirement	Tool
Determine whether the templates follow recommended practices.	Azure Resource Manager test toolkit
Test and validate changes that templates will make to the environment.	What-if operation

Box 1: Azure Resource Manager test toolkit

Use ARM template test toolkit -

The Azure Resource Manager template (ARM template) test toolkit checks whether your template uses recommended practices. When your template isn't compliant with recommended practices, it returns a list of warnings with the suggested changes. By using the test toolkit, you can learn how to avoid common problems in template development.

Box 2: What-if operation -

ARM template deployment what-if operation

Before deploying an Azure Resource Manager template (ARM template), you can preview the changes that will happen. Azure Resource Manager provides the what-if operation to let you see how resources will change if you deploy the template. The what-if operation doesn't make any changes to existing resources.

Instead, it predicts the changes if the specified template is deployed.

Reference:

https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/test-toolkit https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/deploy-what-if

Billabongs Highly Voted 1 year, 1 month ago

Looks Correct

https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/test-toolkit
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/deploy-what-if?tabs=azure-powershell

upvoted 17 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went given answer, score 926

upvoted 7 times

Esward 9 months ago

Agreed

upvoted 1 times

drhuu Highly Voted 1 year ago

on exam oct 2022

upvoted 6 times

seijumi 1 year ago

Whats ur score?

upvoted 1 times

MikeAWS Most Recent 3 weeks, 6 days ago

The Azure Resource Manager template (ARM template) test toolkit checks whether your template uses recommended practices.
https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/test-toolkit

Before deploying an Azure Resource Manager template (ARM template), you can preview the changes that will happen. Azure Resource Manager provides the what-if operation to let you see how resources will change if you deploy the template. The what-if operation doesn't make any changes to existing resources. Instead, it predicts the changes if the specified template is deployed.
https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/deploy-what-if?tabs=azure-powershell

upvoted 1 times

longnguyendh 3 months ago

Correct

upvoted 1 times

Saluk_DE 7 months ago

Question was in Exam 2023-03-30

upvoted 1 times

sarmaria 7 months, 2 weeks ago

Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 1 times

mabdo 8 months, 2 weeks ago

Was om exam 02/2023

upvoted 1 times

kociuba 10 months ago

got it on my exam 30-12-2022
score: 818

upvoted 4 times

OPT_001122 12 months ago

is the given answer correct?

upvoted 1 times

Parkruta 1 year ago

on exam 10/28/2022

upvoted 2 times

vcfvct 1 year ago

Got this on 10/21/2022., score 850

upvoted 3 times

gmishra88 1 year ago

Don't forget which one falls into which one. It can be confusing

upvoted 2 times

coffecold 1 year ago

Correct
https://4bes.nl/2020/08/09/testing-arm-templates/

upvoted 2 times

Question #46 *Topic 2*

You develop Azure Durable Functions to manage vehicle loans.

The loan process includes multiple actions that must be run in a specified order. One of the actions includes a customer credit check process, which may require multiple days to process.

You need to implement Azure Durable Functions for the loan process.

Which Azure Durable Functions type should you use?

A. orchestrator

B. client

C. entity

D. activity

Correct Answer: *A*

Community vote distribution

A (100%)

imanonion Highly Voted 9 months, 3 weeks ago

Selected Answer: A

An orchestrator function is the appropriate type of Azure Durable Function to use in this scenario, because it allows you to define the overall flow of the loan process and call other functions or activities as needed. The credit check process can be implemented as a separate activity function, which can be called by the orchestrator function and run in parallel with other actions in the loan process.

Entity functions are designed for use cases where you need to perform operations on a shared piece of state in a reliable and atomic way, such as a distributed queue or counter. In this scenario, it does not appear that there is a need to use entity functions.

There is no such thing as a "client" function in Azure Durable Functions. "Activity" functions are called by orchestrator functions to perform specific tasks, but the orchestrator function is the one that defines the overall flow of the loan process, so it is the correct answer in this case.

upvoted 19 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went given answer, score 926

upvoted 3 times

Saluk_DE Most Recent 7 months ago

Question was in Exam 2023-03-30

upvoted 1 times

sarmaria 7 months, 2 weeks ago

Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 1 times

NombreFalso 8 months, 2 weeks ago

Selected Answer: A

If we can only choose one then A, would also include D if possible

upvoted 3 times

Esward 9 months, 3 weeks ago

Orchestrator
https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-sequence?tabs=csharp

upvoted 2 times

Question #47 *Topic 2*

HOTSPOT
-

You are developing an Azure Function app.

All functions in the app meet the following requirements:

• Run until either a successful run or until 10 run attempts occur.
• Ensure that there are at least 20 seconds between attempts for up to 15 minutes.

You need to configure the host.json file.

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

NOTE: Each correct selection is worth one point.

Answer Area

Correct Answer:

Answer Area

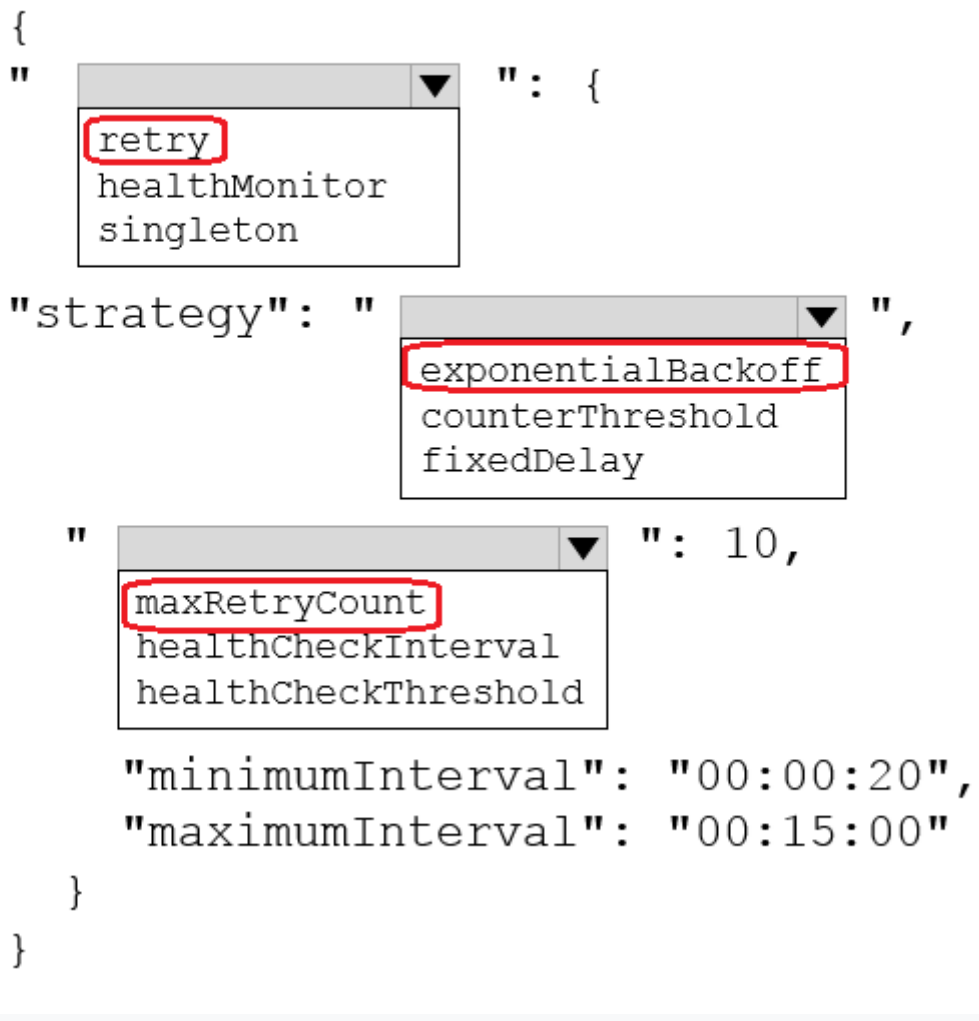

imanonion Highly Voted 9 months, 3 weeks ago

Answer seems correct.

Did some research about exponential backoff:
The "exponential backoff" retry strategy is a technique for retrying failed operations in a manner that avoids overloading the system being accessed. It works by increasing the amount of time that is waited between each retry attempt, using an exponential function to calculate the wait time.

For example, with a coefficient of 2.0 and an initial retry interval of 1 second, the wait times between retries might be 1 second, 2 seconds, 4 seconds, 8 seconds, and so on. This allows the system being accessed to recover from any failures or load spikes before the next retry attempt is made, reducing the likelihood of further failures.

upvoted 11 times

the_odd_one 2 weeks, 2 days ago

Where is then the coefficient argument in the code? or is this an implicit value?

upvoted 2 times

dom271219 1 month, 3 weeks ago

So, based on your explanations, exponentialBackoff may conflict with the minimumInterval (20s) and maximumInterval (15m).

upvoted 1 times

uffuchsi Highly Voted 8 months, 1 week ago

Received this in my exam today (22/02/2023). Selected retry, exponentialBackoff, and maxRetry. Score 927.

upvoted 9 times

85e4c91 Most Recent 2 weeks, 5 days ago

I don't get it, Exponential Backoff would increase the time between retrys, so it would not be 20 seconds anymore and it is clearly stated that the time between retrys should be 20 seconds for the 15 minutes. Both Fixed delay and Exponential backoff have Max retry count.

upvoted 1 times

85e4c91 2 weeks, 5 days ago

Per documentation on https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages?tabs=exponential-backoff%2Cisolated-process&pivots=programming-language-csharp

Fixed Delay: A specified amount of time is allowed to elapse between each retry.
Exponential Backoff: The first retry waits for the minimum delay. On subsequent retries, time is added exponentially to the initial duration for each retry, until the maximum delay is reached. Exponential back-off adds some small randomization to delays to stagger retries in high-throughput scenarios.

upvoted 1 times

leviatas 1 month ago

Got this quesiton in examn, went with answer. - 2023.09.25. Got Case Study Contoso

upvoted 1 times

p2006 1 month, 2 weeks ago

https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages?tabs=exponential-backoff%2Ccsharp-script&pivots=programming-language-csharp#retry-examples

upvoted 1 times

p2006 1 month, 2 weeks ago

"Ensure that there are at least 20 seconds between attempts for up to 15 minutes." ==> exponentialBackoff

upvoted 1 times

longnguyendh 3 months ago

the answer is correct
retry
exponentialBackoff
maxRetryCount

upvoted 1 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 3 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 3 times

PP2015 7 months ago

Received this in my exam today (29/02/2023). Selected retry, exponentialBackoff, and maxRetry. Score 850.

upvoted 3 times

mmdex 8 months, 2 weeks ago

retry
exponentialBackoff
maxRetryCount

https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages?tabs=exponential-backoff%2Ccsharp-script&pivots=programming-language-csharp#retry-strategies

Only fixedDelay and exponentialBackoff strategies are possible and fixedDelay has different parameters than shown here.

upvoted 4 times

Esward 8 months, 2 weeks ago

It was there in 13 Feb 2023

upvoted 2 times

Esward 9 months, 3 weeks ago

retry, fixedDelay, maxretryy
FixedDelay: A specified amount of time is allowed to elapse between each retry.
Exponentional: The first retry waits for the minimum delay. On subsequent retries, time is added exponentially to the initial duration for each retry, until the maximum delay is reached. Exponential back-off adds some small randomization to delays to stagger retries in high-throughput scenarios.

https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages?tabs=fixed-delay%2Cin-process&pivots=programming-language-csharp#retries

upvoted 4 times

Esward 9 months, 3 weeks ago

sorry my bad, given answares are correct.
retry, exponential backoff, maxRetryCount

upvoted 3 times

Sukon_Desknot 9 months, 1 week ago

Why exponential, question says "Ensure that there are at least 20 seconds between attempts for up to 15 minutes". Seems like a fixed time

upvoted 2 times

warchoon 8 months, 1 week ago

10 times 20 seconds less than 4 minutes instead of 15

upvoted 2 times

chettir01 9 months, 3 weeks ago

true :
strategy Required. The retry strategy to use. Valid values are fixedDelay or exponentialBackoff.
maxRetryCount Required. The maximum number of retries allowed per function execution. -1 means to retry indefinitely.
delayInterval The delay that is used between retries when using a fixedDelay strategy. Specify as a string with the format HH:mm:ss.
minimumInterval The minimum retry delay when using an exponentialBackoff strategy. Specify as a string with the format HH:mm:ss.
maximumInterval The maximum retry delay when using exponentialBackoff strategy. Specify as a string with the format HH:mm:ss.

upvoted 3 times

Question #48 *Topic 2*

You develop Azure Web Apps for a commercial diving company. Regulations require that all divers fill out a health questionnaire every 15 days after each diving job starts.

You need to configure the Azure Web Apps so that the instance count scales up when divers are filling out the questionnaire and scales down after they are complete.

You need to configure autoscaling.

What are two possible auto scaling configurations to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Recurrence profile

B. CPU usage-based autoscaling

C. Fixed date profile

D. Predictive autoscaling

Correct Answer: *CD*

Community vote distribution

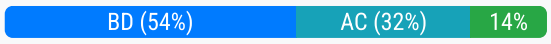

Osip21 Highly Voted 9 months ago

Selected Answer: BD

I think it should
B. CPU usage-based autoscaling
D. Predictive autoscaling

A. Recurrence profile is used to schedule the scaling of resources at specific times or dates, but it does not meet the requirement to scale up when divers are filling out the questionnaire and scale down after they are complete. It only triggers scaling based on a set schedule, not based on actual usage.

C. Fixed date profile is used to specify the number of instances at a specific date and time, but it also does not meet the requirement to dynamically scale based on actual usage. It only sets a fixed number of instances and does not adjust based on changing workloads.

upvoted 18 times

RikinPatel 1 month ago

Predictive autoscale - Support is only available for virtual machine scale sets.
D is incorrect options
A - Recurrence profile is correct because you can specify every 15 days.

upvoted 1 times

Shayman204 3 weeks, 4 days ago

You have many drivers and they dont always fall on same 15day cycle? correct? I think B and D is right also because they mentioned "scaled UP" not scale "OUT" which also validates further its Azure web apps on a VM

upvoted 1 times

imanonion Highly Voted 9 months, 3 weeks ago

Selected Answer: AC

I think it should be A and C instead because:
Options "B. CPU usage-based autoscaling" and "D. Predictive autoscaling" are not relevant to this scenario, because they are based on resource utilization or predictive analytics rather than a fixed schedule.

upvoted 11 times

Teit_examptopics Most Recent 1 month ago

https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-predictive

Only applies to VM Scale sets

upvoted 1 times

katrang 1 month, 1 week ago

Selected Answer: BD

The usage schedule is regular, but we don't know the actual dates on which this happens and if all users fill them out on the same day. But the CPU will go up on those days (A) and the system can work out patterns after it has consumed enough usage data to predict when the spikes will

happen (D)
upvoted 1 times

74gjd_37 1 month, 2 weeks ago

Selected Answer: BD

Each driver needs to fill a profile every 15 days after that driver's job starts. Since the company has many drivers, there are no fixed dates or no recurring date patterns. Therefore, options linked to dates are incorrect. We can only use CPU-based or predictive autoscaling.
upvoted 1 times

Kanasan 2 months ago

Selected Answer: BD

B. CPU usage-based autoscaling
This is one of the most common ways to autoscale an Azure Web App. You can set a threshold for CPU usage, and when the CPU usage goes above that threshold, Azure will automatically add more instances. Similarly, when the CPU usage drops below a certain level, Azure will remove instances. This can be particularly useful if filling out the questionnaires is a resource-intensive operation.
D. Predictive autoscaling
Predictive autoscaling uses machine learning algorithms to predict future demand based on past usage patterns. If you have historical data that shows spikes in usage around the times when divers are filling out the questionnaire, predictive autoscaling can preemptively scale up before the demand hits and scale down afterward.
upvoted 2 times

longnguyendh 3 months ago

I think It should be A & C
upvoted 1 times

[Removed] 3 months, 2 weeks ago

Definitely not A
upvoted 1 times

grudzin 4 months, 2 weeks ago

Selected Answer: BD

Question says "every 15 days after each diving job starts" so it's surely not C (fixed date profile) cause this scale profile is executed only once, not after each time.
There is no information if those jobs starts in fixed time intervals so A (recurrence profile) is also not correct, cause this only works if event happens in strict intervals.
Option B (CPU usage-based) will work correctly every time as it only scales if service is used.
I'm not sure about D (predictive autoscaling) cause it uses AI to determinate scaling (and you never know if AI does things as you want ;-)) but I guess it could work fine (also you need to select 2 answers and A,C are surely not correct).

https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-understanding-settings#autoscale-profiles
upvoted 2 times

kleranti 4 months, 4 weeks ago

Selected Answer: AC

AC as we have a fixed time period of 15 days after the job starts
upvoted 2 times

DonH 4 months, 2 weeks ago

No this is wrong. Because you don't know WHEN they dived AND multiple groups can dive each day of the week which make Recurrence Profile (A) incorrect. Because then it should just scale up everyday for the entire day. The same goes for (C), you don't know when or how many groups are diving at which point in time. So you don't know any Fixed date
upvoted 1 times

br00net 5 months ago

BD
Predictive autoscale uses machine learning to help manage and scale Azure Virtual Machine Scale Sets with cyclical workload patterns.
Predictive autoscale supports scale-out only. Configure standard autoscale to manage scaling in.
upvoted 2 times

br00net 5 months ago

definitively not C, there is no info that a diving job starts on a regular basis at a specific date
upvoted 2 times

dy0917 6 months ago

BD is the correct answer.
upvoted 2 times

lyggwtwtczxnhznebw 6 months, 1 week ago

Selected Answer: BD

ChatGPT thinks B & D:
B. CPU usage-based autoscaling:
This configuration uses the CPU usage of the web app to determine when to scale up or down. As divers fill out the health questionnaire, the web app will experience increased CPU usage, triggering autoscaling to increase the number of instances. Once the questionnaires are complete, CPU usage will decrease, triggering autoscaling to decrease the number of instances.

D. Predictive autoscaling:

This configuration uses historical data and machine learning algorithms to predict when to scale up or down. As the diving company collects more data on the health questionnaire process, the predictive autoscaling algorithm can use that data to predict when to scale up or down based on expected questionnaire submissions. This can help ensure that there are enough instances to handle the questionnaire submissions without overprovisioning and wasting resources.

upvoted 3 times

br00net 5 months ago

Stop using ChatGPT - there is no intelligence, no logic, only more advanced guessing the order of words based on statistics.

upvoted 7 times

macobuzi 2 months, 1 week ago

There is nothing wrong with checking on ChatGPT as long as you make sure to verify the information it gives.

upvoted 1 times

MDWPartners 1 month, 1 week ago

most people don't do that in this comment section

upvoted 1 times

IvanIco 4 months, 1 week ago

why are you getting mad, why do you care if someone uses chatGPT or not

upvoted 1 times

MDWPartners 1 month, 1 week ago

ChatGPT is really faulty when it answer questions like this. So, using that as "proof" is really useless.

upvoted 1 times

ENGs 3 months, 3 weeks ago

This is a comment section for discussing technical questions. Everyone can use ChatGPT independently; therefore, answers from ChatGPT are mostly irrelevant, wrong and/or outdated, and are not really welcome.

upvoted 2 times

IvanIco 4 months, 1 week ago

and statistics are correct more often then not so i would say it's good to check on more places for the correct options then just in the comment section

upvoted 1 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 1 times

shadownin 7 months, 1 week ago

A and C
I was actually trying to work this out for ages and realised the answer. The question doesn't mention anything about automation. It just says there needs to be instances made when filling out. Not automatically done. Meaning you would use recurrance or fixed date by setting the instance to scale out and scale in at the required time.

upvoted 3 times

surprise0011 6 months, 2 weeks ago

make sense

upvoted 1 times

Kluk_Kluk 7 months, 3 weeks ago

Answer is correct C and D
D because Predictive autoscaling supports only available for the Azure Commercial Cloud. Predictive autoscale support only scale-out. Configure standard autoscale to manage scalling in.
C because Fixed date profile supports a single date an time.

upvoted 1 times

Question #49 *Topic 2*

HOTSPOT
-

You are developing an online game that allows players to vote for their favorite photo that illustrates a word. The game is built by using Azure Functions and uses durable entities to track the vote count.

The voting window is 30 seconds. You must minimize latency.

You need to implement the Azure Function for voting.

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

Answer Area

```
[FunctionName("Vote")]
public static async Task<HttpResponseMessage> Run(
  [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
  [                                    ] c,
    CallEntityAsync
    SignalEntityAsync
    [DurableClient] IDurableEntityClient
    [DurableClient] IDurableOrchestrationClient
  string id)
{
  var eid = new EntityId("pic", id);
  await c.[                            ](eid, "vote");
    CallEntityAsync
    SignalEntityAsync
    [DurableClient] IDurableEntityClient
    [DurableClient] IDurableOrchestrationClient

  return req.CreateResponse(HttpStatusCode.OK);
}
```

Correct Answer:

Answer Area

```
[FunctionName("Vote")]
public static async Task<HttpResponseMessage> Run(
  [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
  [                                    ] c,
    CallEntityAsync
    SignalEntityAsync
    [DurableClient] IDurableEntityClient
    [DurableClient] IDurableOrchestrationClient
  string id)
{
  var eid = new EntityId("pic", id);
  await c.[                            ](eid, "vote");
    CallEntityAsync
    SignalEntityAsync
    [DurableClient] IDurableEntityClient
    [DurableClient] IDurableOrchestrationClient

  return req.CreateResponse(HttpStatusCode.OK);
}
```

12jan37 Highly Voted 6 months ago

Answer is correct: https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-dotnet-entities

upvoted 7 times

Teit_examptopics Most Recent 1 month ago

https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-dotnet-entities#example-client-signals-entity

upvoted 1 times

p2006 1 month, 2 weeks ago

https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-dotnet-entities#example-client-signals-entity

upvoted 1 times

adilkhan 6 months, 1 week ago

Answers are correct. Also Chat GPT generated the same code.

upvoted 2 times

d79 6 months, 1 week ago

IDurableEntityClient is correct to works with entities.
Signal is used for 1-way
Call would be used for 2-way

upvoted 2 times

notedo 6 months, 2 weeks ago

This seems to be a new question. It's not here last month. Will exam next week

upvoted 2 times

pelsky 6 months, 1 week ago

Same, I'm writing next week

upvoted 1 times

Question #50 *Topic 2*

HOTSPOT
-

You have an App Service plan named asp1 based on the Free pricing tier.

You plan to use asp1 to implement an Azure Function app with a queue trigger. Your solution must minimize cost.

You need to identify the configuration options that will meet the requirements.

Which value should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Configuration option	Value
Azure App Service feature	Always On Managed identity Continuous deployment
Azure App Service pricing tier	Basic Shared Standard

Correct Answer:

Answer Area

Configuration option	Value
Azure App Service feature	Always On **Managed identity** Continuous deployment
Azure App Service pricing tier	**Basic** Shared Standard

[Removed] Highly Voted 3 months, 2 weeks ago

1. Always on
If you run on an App Service plan, you should enable the Always on setting so that your function app runs correctly. On an App Service plan, the functions runtime goes idle after a few minutes of inactivity, so only HTTP triggers will "wake up" your functions. The Always on setting is available only on an App Service plan. On a Consumption plan, the platform activates function apps automatically.
Even with Always On enabled, the execution timeout for individual functions is controlled by the functionTimeout setting in the host.json project file.

2. Basic
The dedicated App Service plans supported for function app hosting include Basic, Standard, Premium, and Isolated SKUs. Free and Shared tier App Service plans aren't supported by Azure Functions. Basic is the correct answer since it is cheaper than Standard.

https://learn.microsoft.com/en-us/azure/azure-functions/dedicated-plan

upvoted 7 times

RikinPatel 1 month ago

In your ans. you are talking about HTTP trigger but as per que. requirement is "queue trigger"

upvoted 1 times

lord_okello Most Recent 2 weeks ago

Answer is correct. https://learn.microsoft.com/en-us/azure/azure-functions/dedicated-plan

upvoted 1 times

Dixavado 1 month ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 1 times

leviatas 1 month ago

Got this quesiton in examn, went with answer. - 2023.09.25. Got Case Study Contoso

upvoted 1 times

longnguyendh 3 months ago

1. Always on
2. Basic

upvoted 3 times

AlexMMReis 3 months, 3 weeks ago

1. Always on
2. Basic

"If you run on an App Service plan, you should enable the Always on setting so that your function app runs correctly. On an App Service plan, the functions runtime goes idle after a few minutes of inactivity, so only HTTP triggers will "wake up" your functions. The Always on setting is available only on an App Service plan."

https://learn.microsoft.com/en-us/azure/azure-functions/dedicated-plan

upvoted 1 times

AlexMMReis 3 months, 3 weeks ago

Regarding why basic:
Is the the cheapest tier with dedicated compute type

upvoted 1 times

ayush66203 1 week, 2 days ago

but there is not any requirement of having a dedicated computation? i think shared would shared pricing tier will be good here.

upvoted 1 times

75daltonic 3 months, 3 weeks ago

1. Always on
2. Free

upvoted 1 times

75daltonic 3 months, 3 weeks ago

Sorry, 2. Basic :-)

upvoted 1 times

Roel1 3 months, 3 weeks ago

Managed identity?

upvoted 3 times

Question #51 *Topic 2*

DRAG DROP
-

You are developing several microservices to run on Azure Container Apps.

The microservices must allow HTTPS access by using a custom domain.

You need to configure the custom domain in Azure Container Apps.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions

- Validate the custom domain name.
- Enable ingress.
- Bind the certificate.
- Add DNS records to the domain provider.
- Add the custom domain name.

Answer Area

Correct Answer:

Answer Area

- Enable ingress.
- Add the custom domain name.
- Validate the custom domain name.
- Bind the certificate.
- Add DNS records to the domain provider.

Papee Highly Voted 3 months ago

1. Enable ingress: Enable ingress for the Azure Container Apps. This allows external traffic to reach the microservices.
2. Add the custom domain name: Add the custom domain name that you want to use for HTTPS access to your Azure Container Apps.
3. Add DNS records to the domain provider: After adding the custom domain name in Azure Container Apps, you need to add the required DNS records (such as CNAME or A records) to your domain provider's DNS settings. This step is essential for directing traffic from the custom domain to your Azure Container Apps.
4. Validate the custom domain name (after adding the dns records to the provider): After adding the DNS records to the domain provider, you need to validate the custom domain name in Azure Container Apps. This step ensures that the DNS records are correctly configured and the domain is pointing to your Azure Container Apps.
5. Bind the certificate (only enabled after finishing the validation): Once the custom domain name is validated, you can bind the SSL/TLS certificate to enable HTTPS access for the custom domain.

upvoted 14 times

andrewnm Highly Voted 3 months, 1 week ago

from my understanding of https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates

1) Enable Ingress == step 2
2) Add the custom domain name == step 4
3) Bind certificate == steps 5 to 11
4) Add DNS records to the domain provider == 12 to 14
5) Validate the custom domain name == step 15

Using similar questions as a basis, I'm assuming that "Add the custom.." means click the button and not step 16 "Once validation succeeds, select the Add button"

Not very clear.

upvoted 12 times

75daltonic 3 months, 1 week ago

This question is $%&#*!!!
If you follow the steps in https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates, you have to validate the domain

name twice.
In some videos, as @Desniskam has posted, you can see that first of all you fill the domain name and after it, you click on validate button. Then, other configuration options appears. After filling all other configurations, you have to click on validate button again.
My doubt is that I have watched some videos and it is different for app services, web app, etc. I haven't found a container example, so I will follow microsoft learn steps, with validating as last step.

upvoted 1 times

andrewnm 3 months, 1 week ago

Occurs to me that steps 5-11 adding the cert maybe irrelevant and binding the cert is actually the last step, based on what you click in the UI.. so I'd revise that to:

1) Enable Ingress
2) Add the custom domain name
3) Add DNS records to the domain provider
4) Validate the custom domain name
5) Bind certificate,

Still not clear!

upvoted 2 times

Venusmappi Most Recent 2 weeks, 6 days ago

1) Enable Ingress
2) Add the custom domain name
3) Add DNS records to the domain provider
4) Validate the custom domain name
5) Bind certificate

upvoted 1 times

HimGarg 3 weeks, 6 days ago

Why fixed order is required. You can do it in different ways.

upvoted 1 times

NicolasMous 1 month, 2 weeks ago

The correct answer is:
1. Enable ingress
2. Add the custom domain name
3. Bind the certificate
4. Add DNS records to the domain provider
5. Validate the custom domain name
Explanation:

-Enable ingress: This is the first step as it allows HTTP access to your container app.
-Add the custom domain name: After enabling ingress, you can add the custom domain name to your container app.
-Bind the certificate: Once the custom domain name is added, you can bind the certificate to it. This can be a certificate you upload or a free managed certificate provided by Azure.
-Add DNS records to the domain provider: After binding the certificate, you need to add DNS records to your domain provider. This points the domain to your container app and verifies that you own it.
-Validate the custom domain name: Finally, you validate the custom domain name. This ensures that the DNS records have been correctly configured and that the domain correctly points to your container app.

upvoted 3 times

Hasti123 1 month, 1 week ago

100% right

upvoted 1 times

Kanasan 2 months ago

1. Enable ingress
2. Add the custom domain name
3. Bind the certificate
4. Add DNS records to the domain provider
5. Validate the custom domain name
https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates
First validation is for certification, not for a custom domain name.
So, I put the validation step as the last one, because we validate the custom domain at the end.

upvoted 2 times

longnguyendh 3 months ago

https://learn.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-custom-domain?tabs=root%2Cazurecli
I think that It should be
1) Add the custom domain name
2) Bind certificate
3) Validate the custom domain name
4) Add DNS records to the domain provider
5) Enable Ingress

upvoted 2 times

[Removed] 3 months, 2 weeks ago

Given answer is 100% incorrect. This feature is quite new and i never itied it before.
But it should be:

1) Enable Ingress
2) Add the custom domain name
3) Bind certificate
4) Add DNS records to the domain provider
5) Validate the custom domain name

https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates

upvoted 6 times

Brandon902 3 months, 2 weeks ago

I get the same steps when following that documentation.

upvoted 3 times

tarek0811 3 months, 2 weeks ago

1. Enable ingress
2. Add the custom domain name
3. Add DNS records to the domain provider
4. Validate the custom domain name (after adding the dns records to the provider)
5. Bind the certificate (only enabled after finishing the validation)

demo:
https://www.youtube.com/watch?v=BrJtpCV_foU

upvoted 5 times

AbidooKing 3 months, 2 weeks ago

This is the one!

upvoted 1 times

tarek0811 3 months ago

but i am still not sure, because the video is not for Azure Container Apps, which could differ

upvoted 1 times

tarek0811 3 months, 2 weeks ago

1. Enable ingress
2. Add the custom domain name
3. Validate the custom domain name
4. Bind the certificate
5. Add DNS records to the domain provider

https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates

upvoted 3 times

junix_ 3 months, 3 weeks ago

1. Enable ingress
2. Add the custom domain name
3. Validate the custom domain name
4. Add DNS records to the domain provider
5. Bind the certificate

upvoted 1 times

AbidooKing 3 months, 3 weeks ago

Im trying to make sense of this question but fail. I get chicken or the egg scenario when reading the documentation and trying to add the alternatives in a "correct order" seems impossible.. Or am I missing something?

https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates

upvoted 1 times

Potar 3 months, 3 weeks ago

Enable Ingress
Bind certificate
Add DNS records to the domain Provider
Validate Custom domain name
Add Custom domain name
https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates

upvoted 1 times

Roel1 3 months, 3 weeks ago

Maybe validate before adding

upvoted 1 times

Question #52 *Topic 2*

You are developing several microservices to run on Azure Container Apps. External HTTP ingress traffic has been enabled for the microservices.

The microservices must be deployed to the same virtual network and write logs to the same Log Analytics workspace.

You need to deploy the microservices.

What should you do?

A. Enable single revision mode.

B. Use a separate environment for each container.

C. Use a private container registry image and single image for all containers.

D. Use a single environment for all containers.

E. Enable multiple revision mode.

Correct Answer: *A*

Community vote distribution

D (100%)

MikeM27 Highly Voted 3 months, 3 weeks ago

The correct answer should be D

upvoted 8 times

Steffexx 3 months, 3 weeks ago

That is what I thought as well. The single revision would not solve anything here.

upvoted 2 times

p2006 Most Recent 1 month, 2 weeks ago

Selected Answer: D

https://learn.microsoft.com/en-us/azure/container-apps/environment#single-environment

upvoted 1 times

renierdbruyn 1 month, 4 weeks ago

Selected Answer: D

D is correct

upvoted 1 times

Kanasan 2 months ago

Selected Answer: D

D. Use a single environment for all containers.
By deploying all microservices in a single environment, you can ensure that they are part of the same virtual network. This makes it easier to manage network policies and allows the services to communicate more easily with each other. Additionally, using a single environment can simplify the configuration for writing logs to the same Log Analytics workspace.

upvoted 1 times

AashirwadSinha 2 months, 2 weeks ago

My vote goes to D

upvoted 1 times

JeyTlenJey 2 months, 3 weeks ago

I think option D is correct.
A Container Apps environment is a secure boundary around groups of container apps that share the same virtual network and write logs to the same logging destination.
https://learn.microsoft.com/en-us/azure/container-apps/environment

upvoted 2 times

AlexMMReis 3 months, 3 weeks ago

Selected Answer: D

seems the only one to make sense D

upvoted 2 times

lyggwtwtczxnhznebw 3 months, 3 weeks ago

Selected Answer: D

voting for D as well

upvoted 2 times

hch59934 3 months, 3 weeks ago

Selected Answer: D

vote D

upvoted 3 times

75daltonic 3 months, 3 weeks ago

Selected Answer: D

I think D

upvoted 2 times

Question #53 *Topic 2*

HOTSPOT
-

You are developing several microservices to run on Azure Container Apps. External HTTP ingress traffic has been enabled for the microservices.

A deployed microservice must be updated to allow users to test new features. You have the following requirements:

• Enable and maintain a single URL for the updated microservice to provide to test users.
• Update the microservice that corresponds to the current microservice version.

You need to configure Azure Container Apps.

Which features should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Feature
Single URL for test users	Revision label / Revision mode / Container image / Container registry
Current microservice activation	Revision label / Revision mode / Container image / Container registry

Correct Answer:

Answer Area

Requirement	Feature
Single URL for test users	Revision label / Revision mode / Container image / Container registry
Current microservice activation	Revision label / Revision mode / Container image / Container registry

edlima Highly Voted 1 month, 4 weeks ago

1. Revision Label
Labels are useful for testing new revisions. For example, when you want to give access to a set of test users, you can give them the label's URL. Then when you want to move your users to a different revision, you can move the label to that revision.
2. Revision Mode:
The revision mode controls whether only a single revision or multiple revisions of your container app can be simultaneously active. You can set your app's revision mode from your container app's Revision management page in the Azure portal, using Azure CLI commands, or in the ARM template.

https://learn.microsoft.com/en-us/azure/container-apps/revisions

upvoted 5 times

Mpho7 Highly Voted 3 months ago

Correct.

Revision labels are most useful when the app is in multiple revision mode.

Revision mode controls whether only a single revision or multiple revisions of your container app can be simultaneously active.
upvoted 5 times

shekhar11 Most Recent 1 month, 1 week ago
Question in my exam 22sept 2023
upvoted 3 times

nikipediaa 1 month ago
me too, today
upvoted 2 times

macobuzi 2 months, 1 week ago
1. Revision Label:
For container apps with external HTTP traffic, labels are a portable means to direct traffic to specific revisions. A label provides a unique URL that you can use to route traffic to the revision that the label is assigned.
2. Revision Mode
The revision mode controls whether only a single revision or multiple revisions of your container app can be simultaneously active.

https://learn.microsoft.com/en-us/azure/container-apps/revisions
upvoted 3 times

Topic 3 - Question Set 3

Question #1 *Topic 3*

HOTSPOT -

You are developing a solution that uses the Azure Storage Client library for .NET. You have the following code: (Line numbers are included for reference only.)

```
01 CloudBlockBlob src = null;
02 try
03 {
04    src = container.ListBlobs().OfType<CloudBlockBlob>().FirstOrDefault();
05    var id = await src.AcquireLeaseAsync(null);
06    var dst = container.GetBlockBlobReference(src.Name);
07    string cpid = await dst.StartCopyAsync(src);
08    await dst.FetchAttributeAsync();
09    return id;
10 }
11 catch (Exception e)
12 {
13    throw;
14 }
15 finally
16 {
17    if (src != null)
18    await src.FetchAttributesAsync();
19    if (src.Properties.LeaseState != LeaseState.Available)
20    await src.BreakLeaseAsync(new TimeSpan(0));
21 }
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statement	Yes	No
The code creates an infinite lease	◯	◯
The code at line 06 always creates a new blob	◯	◯
The finally block releases the lease	◯	◯

Correct Answer:

Answer Area

Statement	Yes	No
The code creates an infinite lease	◯	◯
The code at line 06 always creates a new blob	◯	◯
The finally block releases the lease	◯	◯

Box 1: Yes -

AcquireLeaseAsync does not specify leaseTime.

leaseTime is a TimeSpan representing the span of time for which to acquire the lease, which will be rounded down to seconds. If null, an infinite lease will be acquired. If not null, this must be 15 to 60 seconds.

Box 2: No -

The GetBlockBlobReference method just gets a reference to a block blob in this container.

Box 3: Yes -

The BreakLeaseAsync method initiates an asynchronous operation that breaks the current lease on this container.
Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.getblockblobreference
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync

Tealon Highly Voted 2 years, 10 months ago

I think the answer is correct:

Optional. Version 2012-02-12 and newer. For a break operation, this is the proposed duration of seconds that the lease should continue before it is broken, between 0 and 60 seconds. This break period is only used if it is shorter than the time remaining on the lease. If longer, the time remaining on the lease is used. A new lease will not be available before the break period has expired, but the lease may be held for longer than the break period. If this header does not appear with a break operation, a fixed-duration lease breaks after the remaining lease period elapses, and an infinite lease breaks immediately.
From: https://docs.microsoft.com/en-us/rest/api/storageservices/lease-blob

upvoted 45 times

AzureFailure 2 years, 4 months ago

It talks about the header values only ReleaseLeaseasync release lease

upvoted 2 times

AzureAz204Fan 2 years, 5 months ago

Last few words "infinite lease breaks immediately" is the key for this context.

upvoted 4 times

Figa Highly Voted 2 years, 11 months ago

I think Box 3 should be no insetad of yes. BreakRelease don`t release the lease directly.

You use ReleaseLease to do this:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.releaselease?view=azure-dotnet-legacy

upvoted 23 times

Figa 2 years, 11 months ago

Ah I think modele is right:
BreakLeaseAsync (TimeSpan? breakPeriod)

breakPeriod
Nullable<TimeSpan>
A TimeSpan representing the amount of time to allow the lease to remain, which will be rounded down to seconds. If null, the break period is the remainder of the current lease, or zero for infinite leases.

upvoted 4 times

Chang401 12 months ago

ya!!! an infinite lease breaks immediately

upvoted 1 times

dom271219 Most Recent 1 month, 3 weeks ago

So, when you provide a TimeSpan value of 0, it essentially means "break the lease right now without any delay." This behavior aligns with the statement you provided: "If null, the break period is the remainder of the current lease, or zero for infinite leases." When you explicitly set it to 0, it ensures an immediate lease termination. It is a matter of knowing the English Language.

upvoted 1 times

kabbas 3 months, 2 weeks ago

to break the lease it will have to enter in the if block ..and since the lease ¡state will be always be available since it is infinite ...it will never enter that block...so the answer is no

upvoted 1 times

Esward 9 months, 3 weeks ago

Yes, No, Yes

upvoted 1 times

Esward 9 months, 3 weeks ago

AcquireLease(null)-->Creates infinite lease

upvoted 2 times

Esward 9 months, 3 weeks ago

AcquireLease(-1) will also create infinite lease

upvoted 1 times

coffecold 1 year ago

version 11 is legacy.
Similar code for version 12 see
https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-copy?tabs=dotnet

upvoted 1 times

coffecold 1 year ago

This is Microsoft.Azure.Storage.Blob v11.1.0 code, it is a legacy library. Did anyone got this on the exam recently?

upvoted 2 times

strongmmc 1 year, 1 month ago

i'm looking for an explanation on what a "lease" is, or what is intended by it, and cannot find any explanatory resource. Does anyone have a reference to it ?

upvoted 3 times

strongmmc 1 year, 1 month ago

https://www.c-sharpcorner.com/article/implementing-blob-leasing-understanding-blob-storage-part-9/

The Lease Blob operation establishes and manages a lock on a blob for write and delete operations. The lock duration can be 15 to 60 seconds or it can be infinite. Once you have the lease, you can update the Blob or delete the Blob without worrying about another process changing it underneath you. When a Blob is leased, other processes can still read it but any attempt to update it will fail. You can also update Blobs without taking a lease first, but you do run the chance of another process also attempting to modify it at the same time.

upvoted 4 times

winston_45 1 year, 2 months ago

Correct answer:
YES
NO
YES

upvoted 3 times

gds03 1 year, 2 months ago

If the finally code releases the lease, then 1st is wrong. The code doesn't create an infinite lease. Answer is NO.

upvoted 3 times

sasha_gud 1 year, 7 months ago

Tried to execute the code sample:
1) If acquires lease on the blob with null parameter, means the infinite lease is acquired [1]
2) On line 6 it gets reference to existing source blob, so no new blob is created there
3) It fails on start copy operation, because it tries to copy onto itself and the lease is already acquired
4) Goes to catch and then finally block
5) In finally breaks the lease with (Zero) parameter that means the immediate lease break [2]

[1] https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync?view=azure-dotnet-legacy#microsoft-azure-storage-blob-cloudblobcontainer-acquireleaseasync(system-nullable((system-timespan))-system-string)
[2] https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync?view=azure-dotnet-legacy#microsoft-azure-storage-blob-cloudblobcontainer-breakleaseasync(system-nullable((system-timespan)))

upvoted 15 times

Manivannan19 1 year, 8 months ago

Answer to the third statement is 'Yes'.
BreakLeaseAsync(mew TimsSpan(0)) will break/release an infinite lease (which is true in this scenario).
Source: https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync?view=azure-dotnet-legacy

upvoted 2 times

lugospod 1 year, 10 months ago

Regarding the third question. The answer is YES.

Here is to link: https://github.com/Azure/azure-storage-net/blob/master/Test/WindowsRuntime/Blob/LeaseTests.cs

Inside this file one can see a unit test at line 255 which does the same scenario and the expected result is the ID of the BROKEN infinite lease.

```
/// <summary>
/// Puts the lease on the given blob in a broken state due to a break period of zero.
/// </summary>
/// <param name="blob">The blob with the lease.</param>
/// <returns>The lease ID of the broken lease.</returns>
internal static async Task<string> SetInstantBrokenStateAsync(CloudBlob blob)
{
string leaseId = await SetLeasedStateAsync(blob, null /* infinite lease */);
await blob.BreakLeaseAsync(TimeSpan.Zero);
return leaseId;
}
```

upvoted 2 times

lugospod 1 year, 10 months ago

tdadadadadadad

upvoted 3 times

amanpritkaur 1 year, 11 months ago

"The final block releases the lease" --> should not this be NO?

"BreakLeaseAsync", breaks the lease but "ReleaseLeaseAsync", releases the lease.

Release, to free the lease if it is no longer needed so that another client may immediately acquire a lease against the blob.
Break, to end the lease but ensure that another client cannot acquire a new lease until the current lease period has expired.

upvoted 4 times

Freidrich 1 year, 8 months ago

So then, since the lease period is infinite, no other client may ever acquire a new lease?

upvoted 2 times

ning 2 years, 2 months ago

The last one should be "NO", this will break the lease immediately. However, you will need a release statement to release the locker, confirmed by running codes.

upvoted 7 times

wolf_lu 2 years, 4 months ago

I agree with the answer

upvoted 2 times

Question #2 *Topic 3*

You are building a website that uses Azure Blob storage for data storage. You configure Azure Blob storage lifecycle to move all blobs to the archive tier after 30 days.
Customers have requested a service-level agreement (SLA) for viewing data older than 30 days.
You need to document the minimum SLA for data recovery.
Which SLA should you use?

A. at least two days

B. between one and 15 hours

C. at least one day

D. between zero and 60 minutes

Correct Answer: *B*
The archive access tier has the lowest storage cost. But it has higher data retrieval costs compared to the hot and cool tiers. Data in the archive tier can take several hours to retrieve depending on the priority of the rehydration. For small objects, a high priority rehydrate may retrieve the object from archive in under 1 hour.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal

Community vote distribution
B (100%)

homimi6115 Highly Voted 2 years, 11 months ago
Answer is correct
upvoted 43 times

surprise0011 6 months, 2 weeks ago
received 2023-04-17 went given answer, score 926
upvoted 4 times

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: B

- Standard priority: The rehydration request will be processed in the order it was received and may take up to 15 hours.
- High priority: The rehydration request will be prioritized over Standard requests and may finish in under 1 hour for objects under ten GB in size.

Reference:

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal#archive-access-tier

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-rehydration?tabs=azure-portal
upvoted 25 times

JH81 Most Recent 4 months ago
Selected Answer: B
Got this on 6/28/2023 and passed with 850. Went with answer.
upvoted 1 times

aragones 5 months, 3 weeks ago
Got this 2023-05-12.
Easy one.
My case:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada
upvoted 1 times

Saluk_DE 7 months ago
Question was in Exam 2023-03-30
upvoted 1 times

sarmaria 7 months, 2 weeks ago
Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.
upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: B

B. between one and 15 hour

upvoted 1 times

elequiel 11 months, 2 weeks ago

Selected Answer: B

Correct Answer: B,

upvoted 1 times

winston_45 1 year, 2 months ago

This question is sooo open for interpretation: "For small objects, a high priority rehydrate may retrieve the object from archive in under 1 hour.".

Where is this 15 hours coming from?

upvoted 4 times

AZAdam22 1 year, 3 months ago

Selected Answer: B

B - Because it takes between one and 15 hours to recover data from the archive tier.

upvoted 2 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.
Gen 1 storage used to be up to 24 hours, now gen 2 is up to 15 hours.

https://docs.microsoft.com/en-us/azure/storage/blobs/archive-rehydrate-overview#rehydration-priority

upvoted 2 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose B. between one and 15 hours

upvoted 1 times

KiranAtShinde 1 year, 9 months ago

Selected Answer: B

Correct Answer is between 1-15 hours

upvoted 1 times

sertes 1 year, 9 months ago

Selected Answer: B

Answer is correct

upvoted 1 times

ucsdmiami2020 1 year, 11 months ago

Selected Answer: B

The answer is correct. Per the provided reference URL https://docs.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview?tabs=azure-portal

Archive tier-An offline tier optimized for storing data that is rarely accessed, and that has flexible latency requirements, on the order of hours. Data in the Archive tier should be stored for a minium of 108 days.

upvoted 1 times

UnknowMan 2 years, 5 months ago

Answer is correct

upvoted 1 times

glam 2 years, 5 months ago

correct.

upvoted 3 times

Question #3 *Topic 3*

HOTSPOT -

You are developing a ticket reservation system for an airline.

The storage solution for the application must meet the following requirements:

- Ensure at least 99.99% availability and provide low latency.
- Accept reservations even when localized network outages or other unforeseen failures occur.
- Process reservations in the exact sequence as reservations are submitted to minimize overbooking or selling the same seat to multiple travelers.
- Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window.

You provision a resource group named airlineResourceGroup in the Azure South-Central US region.

You need to provision a SQL API Cosmos DB account to support the app.

How should you complete the Azure CLI commands? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
resourceGroupName='airlineResourceGroup'
name='docdb-airline-reservations'
databaseName='docdb-tickets-database'
collectionName='docdb-tickets-collection'
consistencyLevel= [ ▼ ]
```

- Strong
- Eventual
- ConsistentPrefix
- BoundedStaleness

```
az cosmosdb create \
--name $name \
[ ▼ ]
```

- --enable-virtual-network true \
- --enable-automatic-failover true \
- --kind 'GlobalDocumentDB' \
- --kind 'MongoDB' \

```
--resource-group $resourceGroupName \
--max-interval 5 \
[ ▼ ]
```

- --locations 'southcentralus'
- --locations 'eastus'
- --locations 'southcentralus=0 eastus=1 westus=2'
- --locations 'southcentralus=0'

```
--default-consistency-level = $consistencylevel
```

Correct Answer:

Answer Area

```
resourceGroupName='airlineResourceGroup'
name='docdb-airline-reservations'
databaseName='docdb-tickets-database'
collectionName='docdb-tickets-collection'
consistencyLevel=
```

- Strong
- Eventual
- ConsistentPrefix
- **BoundedStaleness**

```
az cosmosdb create \
--name $name \
```

- --enable-virtual-network true \
- **--enable-automatic-failover true **
- --kind 'GlobalDocumentDB' \
- --kind 'MongoDB' \

```
--resource-group $resourceGroupName \
--max-interval 5 \
```

- --locations 'southcentralus'
- --locations 'eastus'
- **--locations 'southcentralus=0 eastus=1 westus=2'**
- --locations 'southcentralus=0'

```
--default-consistency-level = $consistencylevel
```

Box 1: BoundedStaleness -

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is,

"updates") of an item or by "T" time interval. In other words, when you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (T) by which the reads might lag behind the writes

Incorrect Answers:

Strong -

Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Box 2: --enable-automatic-failover true\

For multi-region Cosmos accounts that are configured with a single-write region, enable automatic-failover by using Azure CLI or Azure portal.

After you enable automatic failover, whenever there is a regional disaster, Cosmos DB will automatically failover your account.

Question: Accept reservations event when localized network outages or other unforeseen failures occur.

Box 3: --locations'southcentralus=0 eastus=1 westus=2

Need multi-region.

Reference:

https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/manage-with-cli.md

sasisang Highly Voted 2 years, 11 months ago

Answer is correct.

upvoted 67 times

cbn 2 years, 8 months ago

Last option is correct as well
https://stackoverflow.com/questions/51197375/creating-cosmosdb-with-azure-cli-in-powershell

upvoted 6 times

cbn 2 years, 8 months ago

I guess there is a typo on quotes though.

upvoted 3 times

ning Highly Voted 2 years, 2 months ago

No need to overthink here, max-interval, indicates this must be bounded-slateness, enable-automatic-failover, indicated this must be multi-region

upvoted 50 times

OPT_001122 10 months, 3 weeks ago

this is best explanation

upvoted 2 times

jakobste Most Recent 2 months, 2 weeks ago

It's important to carefully read the entire question.

upvoted 3 times

Dianahu 3 months, 1 week ago

I think that enable-automatic-failover = true option does not make any sense without other regions specified but given syntax for regions and priorities is wrong and does not work

upvoted 1 times

Dats1987 6 months, 2 weeks ago

I think ans should be using chat gpt:
1. Eventual : maintain high availability and low latency
2. enable-automatic-failover = true
3. locations "South Central US"=0 : This is primary region

upvoted 2 times

IvanIco 4 months ago

Eventual is definitely not correct

upvoted 4 times

AzSimuz 7 months, 3 weeks ago

BoundedStaleness doesn't guarantee the order of the order, that's session consistency. Where am I wrong?

upvoted 1 times

Esward 8 months, 2 weeks ago

It was there in 13 Feb 2023 exam

upvoted 5 times

mdg3501 10 months, 3 weeks ago

got this on 2022-12-7

upvoted 2 times

leandrocantiero 1 year, 7 months ago

esse aqui ta tudo errado

upvoted 3 times

glam 2 years, 5 months ago

correct

upvoted 2 times

crepatata 2 years, 6 months ago

The strong consistency doesn't work here:
1) there is a --max-interval property which is being used with bounded staleness only
2) there is a requirement "Accept reservations event when localized network outages or other unforeseen failures occur." which points us to multiple writes for multiple regions which is not being supported by the Strong consistency.
--locations syntax looks like obsolete as for late march 2021

upvoted 9 times

Spooky7 2 years, 5 months ago

Ad. 2 - it is not true that it implies writes to multiple regions. In this case there is single write region with automatic failover configured (which provides write availability). Besides with write to multiple regions BoundedStaleness minimum time is 300 seconds, while requirement is 5 seconds (which is actually a minimum for BoundedStaleness with single write region)

upvoted 3 times

Ralph 2 years, 6 months ago

Answer is wrong .. "strong" instead of "boundedStaleness"

upvoted 4 times

TakumaK 2 years, 5 months ago

Are you sure? If then, elaborate why. Does this option --max-interval work for strong?

upvoted 7 times

[Removed] 2 years, 9 months ago

I think this should be "locations southcentralus". The requirement for 99.99% availability is covered by a single region and the syntax for providing multiple regions is wrong.

https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability
https://docs.microsoft.com/en-us/cli/azure/cosmosdb?view=azure-cli-latest#az_cosmosdb_check_name_exists-optional-parameters

upvoted 3 times

agcertif 2 years, 9 months ago

I agree with you.... 99,99 can be reached with only one region.
Why it's not the right solution (someone can helps please) ?

upvoted 3 times

Brak 2 years, 9 months ago

Because of the need for high availability during local network outages. Hence the auto failover selection.

upvoted 4 times

NikB 2 years, 8 months ago

Agreed, also according to the documentation, multiple regions have a min T of 300s:
"For a single region account, the minimum value of K and T is 10 write operations or 5 seconds. For multi-region accounts the minimum value of K and T is 100,000 write operations or 300 seconds."

upvoted 2 times

RaviKS 2 years, 9 months ago

Given Answer is Correct

upvoted 4 times

tomis 2 years, 9 months ago

--locations : Add a location to the Cosmos DB database account.
Usage: --locations KEY=VALUE [KEY=VALUE ...]
Required Keys: regionName, failoverPriority
Optional Key: isZoneRedundant
Default: single region account in the location of the specified resource group.
Failover priority values are 0 for write regions and greater than 0 for read regions. A failover priority value must be unique and less than the total number of regions.
Multiple locations can be specified by using more than one `--locations` argument.

upvoted 8 times

diligent176 2 years, 9 months ago

You are correct about this - none of the given options will work (all wrong syntax).
But I think they "meant" to use the 3 regions answer.

upvoted 4 times

Cornholioz 2 years, 10 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 2 times

Tealon 2 years, 10 months ago

If the last answer is correct (multi-region) than the consistency level cannot be strong (not possible for multi-region), and not Bounded Staleness because multi-region minimal values are 100.000 writes or 300 seconds. But then again, Bounded Staleness is the best option you can select, because the other options are worse (given that strong is not possible).

upvoted 1 times

diligent176 2 years, 9 months ago

the CLI command did not use "--enable-multiple-write-locations", therefore this will configure a single write region for this CosmosDB. Others will be read regions.

upvoted 3 times

Question #4 *Topic 3*

HOTSPOT -

You are preparing to deploy a Python website to an Azure Web App using a container. The solution will use multiple containers in the same container group. The

Dockerfile that builds the container is as follows:

```
FROM python:3
ADD website.py
CMD [ "python", "./website.py"]
```

You build a container by using the following command. The Azure Container Registry instance named images is a private registry.

```
docker build –t images.azurecr.io/website:v1.0.0
```

The user name and password for the registry is admin.

The Web App must always run the same version of the website regardless of future builds.

You need to create an Azure Web App to run the website.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
az configure --defaults web=website
az configure --defaults group=website
az appservice plan create --name websitePlan [ ▼ ]
```

- --sku SHARED
- --tags container
- --sku B1 --hyper-v
- --sku B1 --is-linux

```
az webapp create --plan websitePlan [ ▼ ]
```

- --deployment-source-url images.azurecr.io/website:v1.0.0
- --deployment-source-url images.azurecr.io/website:latest
- --deployment-container-image-name images.azurecr.io/website:v1.0.0
- --deployment-container-image-name images.azurecr.io/website:latest

```
az webapp config [ ▼ ]
```

- set --python-version 2.7 --generic-configurations user=admin password=admin
- set –python-version 3.6 --generic-configurations user=admin password=admin
- container set --docker-registry-server-url https://images.azurecr.io –u admin –p admin
- container set --docker-registry-server-url https://images.azurecr.io/website -u admin -p admin

Correct Answer:

Answer Area

```
az configure --defaults web=website
az configure --defaults group=website
az appservice plan create --name websitePlan [ ▼ ]
```

- --sku SHARED
- --tags container
- --sku B1 --hyper-v
- --sku B1 --is-linux

```
az webapp create --plan websitePlan [ ▼ ]
```

- --deployment-source-url images.azurecr.io/website:v1.0.0
- --deployment-source-url images.azurecr.io/website:latest
- --deployment-container-image-name images.azurecr.io/website:v1.0.0
- --deployment-container-image-name images.azurecr.io/website:latest

```
az webapp config [ ▼ ]
```

- set --python-version 2.7 --generic-configurations user=admin password=admin
- set –python-version 3.6 --generic-configurations user=admin password=admin
- container set --docker-registry-server-url https://images.azurecr.io –u admin –p admin
- container set --docker-registry-server-url https://images.azurecr.io/website -u admin -p admin

Box 1: --SKU B1 --hyper-v -

--hyper-v

Host web app on Windows container.

Box 2: --deployment-source-url images.azurecr.io/website:v1.0.0

--deployment-source-url -u

Git repository URL to link with manual integration.

The Web App must always run the same version of the website regardless of future builds.

Incorrect:

--deployment-container-image-name -i

Linux only. Container image name from Docker Hub, e.g. publisher/image-name:tag.

Box 3: az webapp config container set -url https://images.azurecr.io -u admin -p admin az webapp config container set

Set a web app container's settings.

Paremeter: --docker-registry-server-url -r

The container registry server url.

The Azure Container Registry instance named images is a private registry.

Example:

az webapp config container set --docker-registry-server-url https://{azure-container-registry-name}.azurecr.io

Reference:

https://docs.microsoft.com/en-us/cli/azure/appservice/plan

glam Highly Voted 2 years, 5 months ago

--sku B1 --is-linux
--deployment-container-image-name images.azurecr.io/website:v1.0.0
-- container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin

upvoted 136 times

applepie 3 months ago

got this question today, go with this answer - 7/30/2023, score 895/1000

upvoted 3 times

Esward 9 months, 3 weeks ago

https://www.linkedin.com/pulse/how-quickly-create-micro-service-azure-webapp-fastapi-bonnet-?trk=pulse-article_more-articles_related-content-card

upvoted 2 times

aperez1979 Highly Voted 2 years, 6 months ago

"use multiple containers in the same container group" this not is possible in windows.
Solution is:
--is-linux
--deployment-container-image-name

upvoted 40 times

Tom87 2 years, 6 months ago

That's true.
"Multi-container groups are currently restricted to Linux containers."
https://docs.microsoft.com/en-us/azure/container-instances/container-instances-multi-container-group

upvoted 8 times

azurelearner666 2 years, 4 months ago

True! also last response is correct with
-- container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin

upvoted 5 times

leviatas Most Recent 1 month ago

Got this quesiton in examn, went with answer. - 2023.09.25. Got Case Study Contoso

upvoted 1 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 1 times

surprise0011 6 months, 2 weeks ago

microsoft promote good practises for -u and -p (°ʖ°)

upvoted 8 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 5 times

uffuchsi 8 months, 1 week ago

Received this in my exam today (22/02/2023). Selected linux, --deployment-container...:v1.0.0, and container set ... in --u admin --p admin. Score 927.

upvoted 4 times

Esward 8 months, 2 weeks ago

It was there in 13 Feb 2023 exam

upvoted 4 times

Tito28 10 months ago

this answer are not wrong, if you must know, these question and answer are giving by microsoft , so how can it be wrong if microsoft support those answer?

upvoted 1 times

Net_IT 8 months, 2 weeks ago

The answers are not provided by Microsoft but by the community/Examtopics, sometimes given answers are wrong. But the discussions are great to help you with that! In this case I believe the highest voted answers are correct (from aperez1979).

upvoted 4 times

Dani_ac7 1 year, 4 months ago

Linux for multiple containers

upvoted 3 times

Pize 1 year, 4 months ago

Another wrong answer - this site riddled with wrong answers

upvoted 6 times

Eltooth 1 year, 4 months ago

--sku B1 --is-linux
--deployment-container-image-name images.azurecr.io/website:v1.0.0
-- container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin

upvoted 4 times

yevgen91 1 year, 10 months ago

The suggested answer is for Windows, but Python is deprecated for Windows, as well as multiple containers are not possible in Windows.
So it must be Linux, and based on this article https://docs.microsoft.com/en-us/azure/app-service/tutorial-custom-container?pivots=container-linux
the answer must be:
1. --sku B1 --is-linux
2. --deployment-container-image-name ... :v1.0.0
3. container set ... https://images.azurecr.io ...

upvoted 4 times

Molte 1 year, 10 months ago

for answer 3: Can somebody tell me why its images.azurecr.io and not images.azurecr.io/website?

upvoted 4 times

ScubaDiver123456 1 year, 9 months ago

The command requires the URL of the registry, not the image name. That is appended later.
"--docker-registry-server-url
The container registry server url."

upvoted 9 times

BeshoyRomany 2 years, 1 month ago

Hi All the Correct answer is
--sku B1 --is-linux
--deployment-container-image-name images.azurecr.io/website:v1.0.0
-- container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin

Multi-container groups currently support only Linux containers. For Windows containers, Azure Container Instances only supports deployment of a single container instance. While we are working to bring all features to Windows containers, you can find current platform differences in the service

https://docs.microsoft.com/en-us/azure/container-instances/container-instances-container-groups

upvoted 5 times

shai0946 2 years, 3 months ago

Multi-container groups currently support only Linux containers. For Windows containers, Azure Container Instances only supports deployment of a single container instance.
https://docs.microsoft.com/en-us/azure/container-instances/container-instances-container-groups

upvoted 2 times

bssehra 2 years, 3 months ago

I know that Multi-container groups are currently restricted to Linux containers but here we are talking about multiple containers in same container group. I think it is possible in windows. Is there any helpful link which says its not possible on Windows?

upvoted 2 times

JamieS 9 months ago

It's the same thing.
https://learn.microsoft.com/en-us/azure/container-instances/container-instances-multi-container-group

"Azure Container Instances supports the deployment of multiple containers onto a single host using a container group."

"Multi-container groups are currently restricted to Linux containers."

""osType": "Linux","

upvoted 1 times

Question #5 *Topic 3*

HOTSPOT -

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

A rule already exists to scale up the App Service when the average queue length of unprocessed and valid queue messages is greater than 1000.

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

How should you configure the Scale rule? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

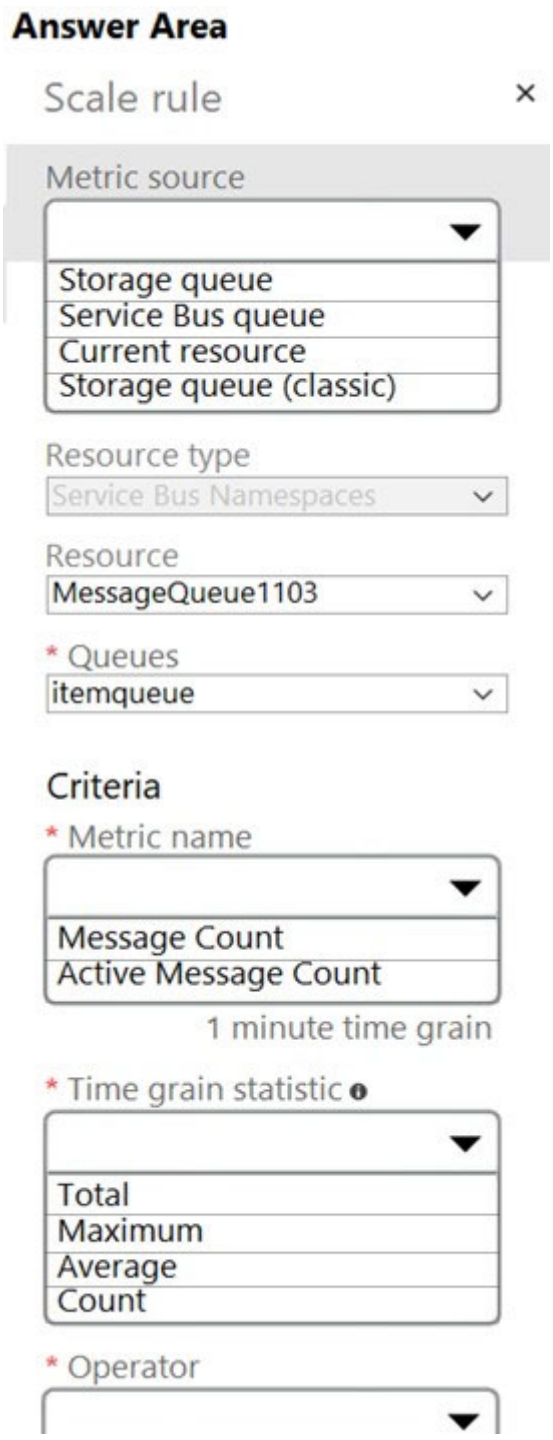

Correct Answer:

Answer Area

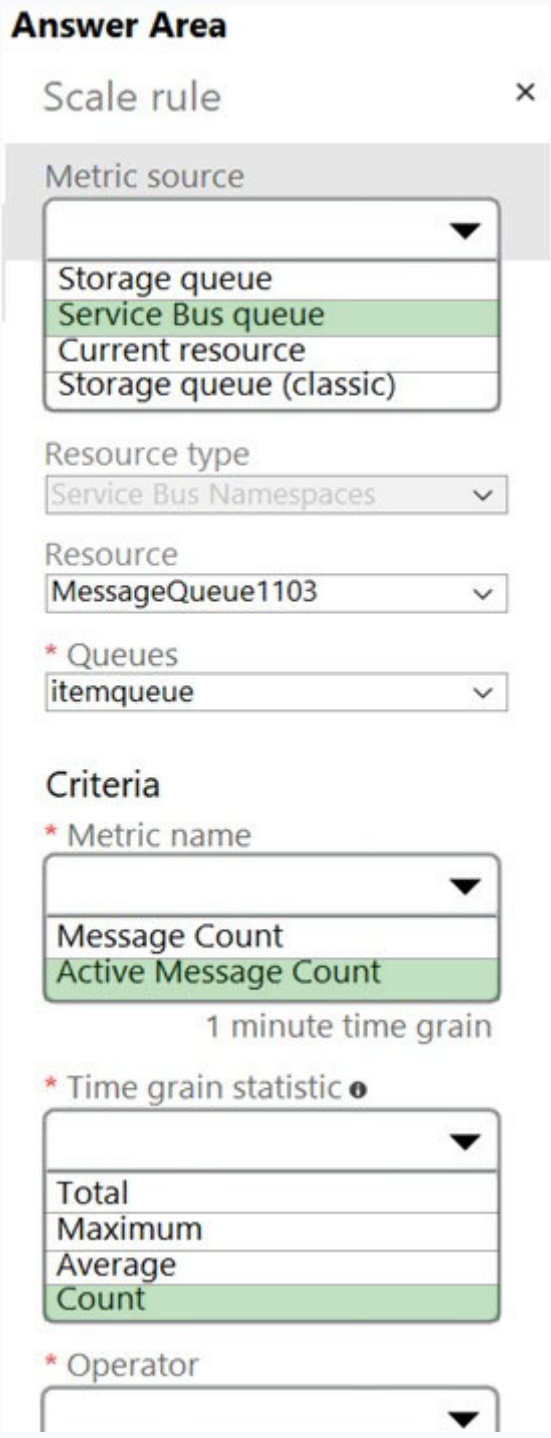

Box 1: Service bus queue -

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

Box 2: ActiveMessage Count -

ActiveMessageCount: Messages in the queue or subscription that are in the active state and ready for delivery.

Box 3: Count -

Box 4: Less than or equal to -

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

Box 5: Decrease count by

VijiVivek Highly Voted 2 years, 6 months ago

The correct answers are
1) Service bus queue
2) Active message count
3) Average
4) Less than or equal to
5) Decrease count by

upvoted 111 times

azurelearner666 2 years, 4 months ago

this is correct (the "official" solution is wrong)
Also check the proper image https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_062_Image_0001.jpg

upvoted 37 times

Evo_Morales 1 year, 6 months ago

How often are the responses scrubbed for updates/corrections?

upvoted 1 times

micro9000 11 months, 3 weeks ago

I think the #4 is should be "Less than" only, because you don't want to trigger the Scale up and Scale down at the same time.

upvoted 7 times

JustHereToLearn 9 months, 1 week ago

I won't trigger at same time as question says "average queue length of unprocessed and valid queue messages is greater than 1000"

upvoted 3 times

micro9000 11 months, 3 weeks ago

Here is a reference why I think it should be Less than,
https://learn.microsoft.com/en-us/training/modules/scale-apps-app-service/5-autoscale-best-practices?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wwl.create-azure-app-service-web-apps

upvoted 1 times

applepie 3 months ago

got this question today, go with this answer - 7/30/2023, score 895/1000

upvoted 1 times

mlantonis Highly Voted 2 years, 5 months ago

Full image: https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_062_Image_0001.jpg

upvoted 75 times

mlantonis 2 years, 5 months ago

Box 1: Service bus queue
You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

Box 2: ActiveMessage Count
ActiveMessageCount: Number of messages in the queue or subscription that are in the active state and ready for delivery.

Box 3: Average
For special metrics such as Storage or Service Bus Queue length metric, the threshold is the average number of messages available per current number of instances.

Box 4: Less than or equal to
You need to add a new rule that will continuously scale down the App Service, as long as the scale up condition is not met.

Box 5: Decrease count by

upvoted 37 times

mlantonis 2 years, 5 months ago

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-best-practices#considerations-for-scaling-threshold-values-for-special-metrics

https://docs.microsoft.com/en-us/azure/service-bus-messaging/message-counters

upvoted 9 times

warchoon 8 months, 1 week ago

https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-best-practices#considerations-for-scaling-when-multiple-rules-are-configured-in-a-profile
points Less than

upvoted 1 times

jay158 2 years, 5 months ago

You are providing invaluable help, by clarifying the questions,

upvoted 11 times

NightshadeRC Most Recent 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 2 times

AsuraX 9 months, 3 weeks ago

Got this on my exam, January 10, 2023 (I passed)
My answer:
Box 1: Service bus queue
Box 2: Active Message Count
Box 3: Average
Box 4: Less than or equal to
Box 5: Decrease count by

upvoted 3 times

Lucifer14 10 months ago

In 28-12-2022 exam, Answers.
Service Bus queue
active message count
count
less than equal to
decrease count

upvoted 1 times

at_sayali 11 months, 3 weeks ago

Got it on 10/2022, I chose as below:
Box 1: Service bus queue
Box 2: Active Message Count
Box 3: Average
Box 4: Less than or equal to
Box 5: Decrease count by

upvoted 2 times

OPT_001122 1 year, 1 month ago

Admin, Request to please update the full image. Current image is cut and options 4 and 5 are not visible

upvoted 3 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose as below:
Box 1: Service bus queue
Box 2: ActiveMessage Count
Box 3: Average
Box 4: Less than or equal to
Box 5: Decrease count by

upvoted 3 times

GigaCaster 2 years, 3 months ago

Just a question but how are you going to reduce the count if there is now count to reduce?

upvoted 2 times

glam 2 years, 5 months ago

The correct answers are
1) Service bus queue
2) Active message count
3) Average
4) Less than or equal to
5) Decrease count by

upvoted 4 times

Phenr 2 years, 7 months ago

I believe the answers are:
1st box: as given, since we're analyzing the state of messages in the Service Bus Queue
2nd box: as given, because Active Message Count will show us just the count of messages that are in the active state and ready for delivery, while Message Count would show us all the messages.
3rd box: Average, since we should stay consistent with the Scale Up time grain, which is also "Average" - this way, when we use "average" on both, we have consistent and clear conditions when to go up and when to go down

upvoted 14 times

hobob 2 years, 7 months ago

I think the 3rd option should be "Average" since the scale up rule is based off average, it makes sense to make the scale down rule work on the same principal.

upvoted 5 times

Mike_St 2 years, 7 months ago

The image is cut so i cant see the bottom 2 dropdowns but they sound alright in the answer explanation bellow.
Also the 2nd box should be Message Count.

Reason:
In the question it is said:
"when the average queue length of unprocessed AND valid queue messages"

If it was just valid Queue messages it would have been enough to set it on Active Message Count... but they want Valid AND Unproscessed messages... so i think it should be Message Count.
(i am not 100% sure but unprosessed message should not count as an active message)

upvoted 3 times

qerem 2 years, 7 months ago

Messagecount is false because it also contains deadlettered messages among others. Unprocessed messages are active because they are "ready to delivery". So i think "Active Message Count" is correct

upvoted 3 times

Question #6 *Topic 3*

DRAG DROP -

You have an application that uses Azure Blob storage.

You need to update the metadata of the blobs.

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

Select and Place:

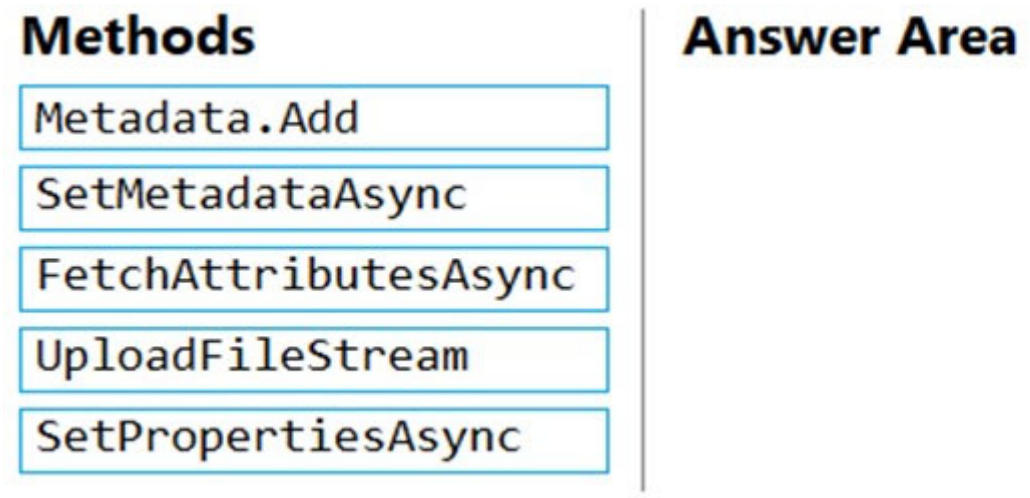

Correct Answer:

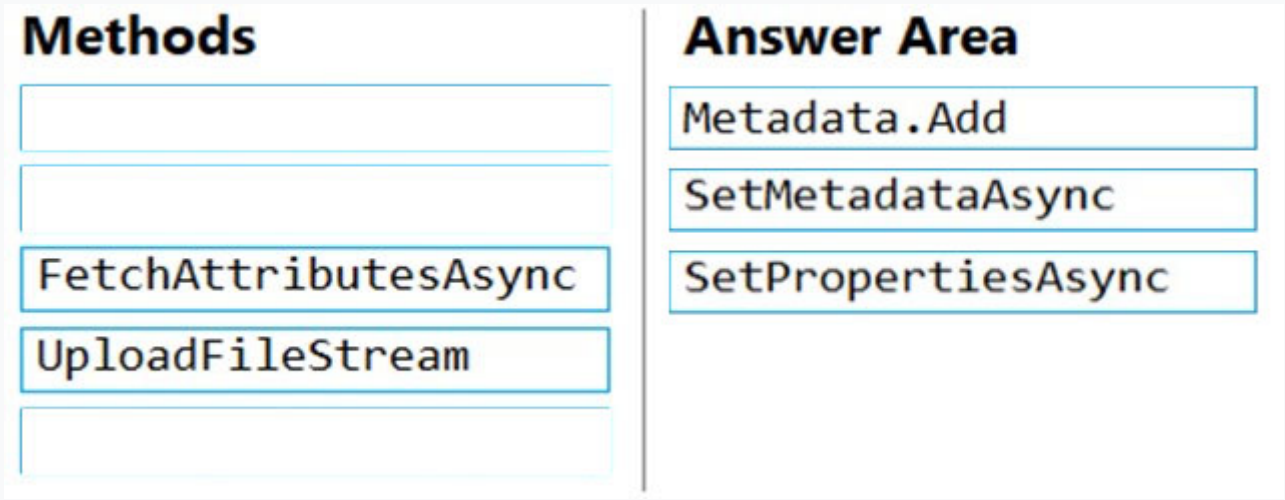

Metadata.Add example:

// Add metadata to the dictionary by calling the Add method

metadata.Add("docType", "textDocuments");

SetMetadataAsync example:

// Set the blob's metadata.

await blob.SetMetadataAsync(metadata);

// Set the blob's properties.

await blob.SetPropertiesAsync();

Reference:

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-properties-metadata

Phenr Highly Voted 2 years, 7 months ago

Since we're talking about updating the metadata,
- first we need to fetch it, to populate blob's properties and metadata (we want to update it - without fetching we would just set the new metadata):
FetchAttributesAsync
- second, we need to manipulate the metadatas to update them and the best fitting is
Metadata.Add
- third, we have to persist our changes. We can use a method that initiates an asynchronous operation to update the blob's metadata, which is SetMetadataAsync

upvoted 231 times

altafpatel1984 1 year, 11 months ago

This is correct answer. Others answers suggesting GetProperties/SetProperties are wrong, because GetProperties is no in options at all, and there is no SetProperties method.

upvoted 4 times

azurelearner666 2 years, 4 months ago

This is the best response (the "official" response is wrong)

upvoted 4 times

rdemontis 2 years, 7 months ago

correct!

upvoted 2 times

TakumaK 2 years, 5 months ago

nicely explained!

upvoted 3 times

Tom87 Highly Voted 2 years, 6 months ago

I suppose there are two correct answers, depending on the version of Azure.Storage.Blobs. For v11: FetchAttributesAsync, Metadata.Add, SetMetadataAsync.
For v12: GetPropertiesAsync, Metadata.Add, SetMetadataAsync.
Just look here, there are two tabs with source code. One for v11, one for v12:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-properties-metadata?tabs=dotnet#set-and-retrieve-metadata

upvoted 51 times

vtomy 2 years, 4 months ago

Correct. But GetPropertiesAsync not mentioned in answer options. We can go with FetchAttributesAsync

upvoted 13 times

xRiot007 1 year, 4 months ago

MS updated their pages.
Now, to set metadata, you need to add it either with metadata.Add, or metadata[key] = value.
After, you just need to save it, by calling SetMetadataAsync.
No other operations are required as we do not care what metadata is already on the blob set.

upvoted 1 times

macobuzi 2 months, 1 week ago

Yeah, But the question forces you to pick 3 options. So FetchAtrributeAsync is probably the closest one.

upvoted 1 times

ENGs Most Recent 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

RuffBoii 1 month, 1 week ago

Had this on my exam today.

upvoted 1 times

dddddd111 1 month, 1 week ago

I got this same question. Provided answers are correct. (Note: I failed the exam 20/Sept/23. I only scored 644 and I felt bad. I think because many questions here in Examtopics are not correct or already outdated. I suggest following the most voted answers and not rely on Examtopics answers. At the beginning of the exam, you will be asked which programming languages you want to use. C#/Python. I chose C#. Also, I just want to add that some questions here in the actual exams, but the choices are written and formatted differently. Please be aware of that. Goodluck. I feel bad for failing it, but I want to retake next month. I will try Python. T_T

upvoted 1 times

TechyNetty 1 month, 1 week ago

Sorry to hear that. I have exam soon. Were all the questions from ExamTopics> Any tips you would like to share please, Thanks

upvoted 1 times

longnguyendh 2 months, 4 weeks ago

It should be FetchAttributesAsync, Metadata.Add, SetMetadataAsync

upvoted 2 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with fetch/metadata.add/setMetadata. Scored 917

upvoted 4 times

DonH 4 months, 2 weeks ago

Just for information: I just had this question on my AZ204 exam - 16-jun-2023.
I barely made it (with only 767 points) so I can't inform anyony if this answer is correct or not, just stating that this is an actual exam question.

upvoted 3 times

saravanasanthosh 10 months ago

Got this in exam 12/30/2022

upvoted 2 times

OPT_001122 10 months, 3 weeks ago

thanks to all who have mentioned the exam dates

upvoted 1 times

OPT_001122 11 months ago

FetchAttributesAsync
Metadata.Add
SetMetadataAsync

upvoted 3 times

elequiel 1 year ago

Got it in exam 20/10/2022

upvoted 2 times

coffecold 1 year ago

FetchAttributesAsync is version 11 (legacy)
In version 12 is would be :
1. await blob.GetPropertiesAsync();
2. metadata.Add("docType", "textDocuments");
3. await blob.SetMetadataAsync(metadata);
https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-properties-metadata

upvoted 4 times

vinodshetty 1 year, 1 month ago

Got it on 03/09/2022

upvoted 3 times

Yoennis 1 year, 3 months ago

Correct Answer: FetchAttributesAsync, Metadata.Add and SetMetadataAsync

upvoted 5 times

Evo_Morales 1 year, 6 months ago

Whoever writes some of these questions should be taken out and flogged. They are terribly worded, ambiguous, and often irrelevant to helping someone architect a solution.

upvoted 7 times

gmishra88 1 year ago

True that. Similar feeling. It's but a blame on us making ourselves use products from Microsoft and more importantly think of doing this certification

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, chose as FetchAttributesAsync => Metadata.Add => SetMetadataAsync

upvoted 5 times

bazzopardi 1 year, 3 months ago

Was it correct?

upvoted 2 times

RikinPatel 1 month ago

no one know which ans. actually correct as per microsoft. because after exam you wil just get total your score.

upvoted 1 times

Question #7 *Topic 3*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce
2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Event Grid. Configure the machine identifier as the partition key and enable capture.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

Community vote distribution
B (100%)

bbijls Highly Voted 2 years, 7 months ago
I don't know who gave all the answer, but that person definitely failed the exam..
upvoted 135 times

basquiat 2 months, 3 weeks ago
I honestly think that this is on purpose so people actually discuss the questions. Reverse psychology. LOL
upvoted 3 times

AnuSubramaniam 2 years, 2 months ago
Its Event Hub not, Grid.
upvoted 22 times

kondapaturi 2 years, 4 months ago
what is the answer?
upvoted 1 times

Anil4Az 2 years, 5 months ago
Good one... :)
upvoted 4 times

Dinima Highly Voted 2 years, 7 months ago
I think event hub. Azure Event Hub can be used to get the messages from the various devices. Azure Event Hub capture can then be used to persist the events to Azure Blob storage.
upvoted 48 times

Tom87 2 years, 6 months ago
I agree. The solution says "enable capture". I have found capture only for Event Hub, not for Event Grid.
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-capture-overview
upvoted 8 times

Esward 9 months ago
Agreed
upvoted 1 times

Samuel316 Most Recent 2 weeks, 6 days ago
It should be EventHub, not EventGrid.
upvoted 1 times

ENGs 3 weeks, 4 days ago
Had this series of questions On my exam 2023-10 before the Update of the Exam
upvoted 1 times

macobuzi 2 months, 1 week ago

The answer is Azure Event Hub, not Event Grid.

upvoted 4 times

KarthikCS 6 months, 1 week ago

How come almost every answer marked by examtopic is incorrect. Whats the purpose of marking answers then if you have to follow discussion thread for each question?

upvoted 5 times

Kluk_Kluk 7 months, 3 weeks ago

"device data from 2,000 stores located throughout the world." It is a distributed data streaming. Answer is EventHub

upvoted 4 times

fkaracan 8 months ago

Selected Answer: B

Given solution does not meet the goal because you are using Event Grid instead of Event Hubs. Event Grid does not have a capture feature and it does not store data in Blob storage. You need to use Event Hubs as your source and enable capture with Blob storage as your destination.

upvoted 4 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. No.
Azure Event Grid is an event routing service that allows you to handle events from various Azure services and your own applications. It can be used to send events from an application to multiple subscribers, but it is not well suited for receiving data from thousands of devices and storing them in Azure Blob storage.
Azure Event Grid is also not meant for long-term data storage and it is not a good fit for this scenario, where large amounts of data need to be stored and correlated based on a device identifier.

upvoted 4 times

cemergen 10 months, 1 week ago

Selected Answer: B

B is the correct answer

upvoted 1 times

Knightie 1 year, 2 months ago

Selected Answer: B

The reason for No is due to the max no. of partition. The question demands 5x2000 = 10000 parition, which is more than all tiers available.

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-quotas#basic-vs-standard-vs-premium-vs-dedicated-tiers

upvoted 3 times

Satish_Babu 1 year, 2 months ago

Selected Answer: B

EventHub not EventGrid

upvoted 1 times

sahirl 1 year, 3 months ago

Selected Answer: B

B is the correct answer

upvoted 1 times

AZAdam22 1 year, 3 months ago

Selected Answer: B

No - Capture is a feature of EventHub, not EventGrid.

upvoted 2 times

Dani_ac7 1 year, 4 months ago

Selected Answer: B

is event hub, not grid

upvoted 2 times

xRiot007 1 year, 4 months ago

Selected Answer: B

Event Hub is the correct option. Sensors are sending data, they aren't just notifying you that something happened.

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.
Event Hub would be a solution here NOT Event Grid.

upvoted 1 times

Question #8 *Topic 3*

You develop Azure solutions.

A .NET application needs to receive a message each time an Azure virtual machine finishes processing data. The messages must NOT persist after being processed by the receiving application.

You need to implement the .NET object that will receive the messages.

Which object should you use?

A. QueueClient

B. SubscriptionClient

C. TopicClient

D. CloudQueueClient

Correct Answer: *D*

A queue allows processing of a message by a single consumer. Need a CloudQueueClient to access the Azure VM.

Incorrect Answers:

B, C: In contrast to queues, topics and subscriptions provide a one-to-many form of communication in a publish and subscribe pattern. It's useful for scaling to large numbers of recipients.

Reference:

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-queues-topics-subscriptions

Community vote distribution

A (100%)

mlantonis Highly Voted 2 years, 5 months ago

Correct Answer: A

Azure.Storage.Queues.QueueClient: .NET v12
Azure.Storage.Queues.CloudQueueClient: .NET v11 (Legacy)

So, the question is really about what kind of queue message tool you should use. And the key word here is that "message must NOT persist after being processed".

Azure.Storage.Queues.QueueClient supports "At-Most-Once" deliver mode, while Azure.Storage.Queues.CloudQueueClient doesn't.

Reference:

https://docs.microsoft.com/en-us/dotnet/api/azure.storage.queues.queueclient?view=azure-dotnet

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.queue.cloudqueueclient?view=azure-dotnet-legacy

upvoted 82 times

guchao2000 9 months, 1 week ago

It's QueueClient of Service Bus.

upvoted 2 times

MiraA 2 years ago

It seems the CloudQueueClient is a legacy class and is used for creation of QueueClient instances.
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.queue.cloudqueueclient.getqueuereference

upvoted 1 times

Spooky7 Highly Voted 2 years, 5 months ago

First of all - question is not precise as we don't know which QueueClient they are asking about. There are two options:
- Microsoft.AzureService.Bus.QueueClient?
- Azure.Storage.Queues.QueueClient?

I would say it is about Microsoft.AzureService.Bus.QueueClient as the difference between Azure.Storage.Queues.CloudQueueClient (v12) is just a legacy version of the Azure.Storage.Queues.QueueClient (v11)
So the question is really about what kind of queue message tool you should use. And the key word here is that "message must NOT persist after being processed". So correct answer would be Microsoft.AzureService.Bus.QueueClient (A) as it supports "At-Most-Once" deliver mode while Azure.Storage.Queues.CloudQueueClient doesn't.

upvoted 19 times

basquiat 2 months, 3 weeks ago

I agree. Plus that Azure.Storage.Queues.QueueClient AFAIK is more appropriate for data intensive scenarios.

upvoted 1 times

RikinPatel Most Recent 4 weeks, 1 day ago

Selected Answer: A

@Admin Please correct answer.
As per MS Doc and ChatGPT correct ans is A so please correct ans and its explanation

upvoted 2 times

alexein74 9 months, 1 week ago

Selected Answer: A

A. QueueClient

upvoted 1 times

OPT_001122 11 months ago

Selected Answer: A

A. QueueClient

upvoted 1 times

gmishra88 1 year ago

Microsoft creates all these random naming convention showing how disorganized they are in individual islands and they dare to ask questions along this fault lines.

upvoted 5 times

Perplex 1 year, 2 months ago

Selected Answer: A

A is correct, the new API calls it QueueClient. See also most upvoted answer.

upvoted 3 times

xRiot007 1 year, 4 months ago

A - version 12
D - "legacy" version 11

upvoted 1 times

Bartimaeus 1 year, 8 months ago

What about SeviceBus QueueClient in the old API?
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.servicebus.queueclient?view=azure-dotnet

upvoted 1 times

idroj 1 year, 8 months ago

Selected Answer: A

A is the correct answer

upvoted 3 times

Ccastan1 1 year, 8 months ago

So the explanation for why A is wrong is incorrect? It states you cannot access the VM using A

upvoted 1 times

prabhjot 2 years, 5 months ago

yes A is correct

upvoted 1 times

glam 2 years, 5 months ago

A. QueueClient

upvoted 2 times

aperez1979 2 years, 7 months ago

Why not A?

upvoted 4 times

Shion2009 2 years, 7 months ago

See the answer information:
A queue allows processing of a message by a single consumer. Need a CloudQueueClient to access the Azure VM.

upvoted 3 times

ivan0590 1 year, 4 months ago

I really don't know why you need to access the Azure VM, it makes no sense. Whatever the VM does, in the end, it add a message to a queue in Azure. And you don't need to access the VM to access the queue. This explanation is very awkward.
Also, is it true that CloudQueueClient can access a VM? That sounds really strange, it goes beyond the purpose of a "queue client". It's like using a fork as a screwdriver...
The only difference between CloudQueueClient and QueueClient I've found so far is that CloudQueueClient is the legacy version of the client (v11) and QueueClient is the newest version (v12).
So, I would choose A.

upvoted 2 times

MrZoom 2 years, 7 months ago

Agreed, A is the new-style API, and D is the old-style API, so IMHO, A is better.

Link (A): https://docs.microsoft.com/en-us/dotnet/api/azure.storage.queues.queueclient?view=azure-dotnet
Link (D): https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.queue.cloudqueueclient?view=azure-dotnet-legacy

upvoted 18 times

vladans 2 years, 7 months ago

You're right.
new version --> .NET v12 --> QueueClient
old version --> .NET v11 --> CloudQueueClient
Link --> https://docs.microsoft.com/en-us/azure/storage/queues/storage-dotnet-how-to-use-queues?tabs=dotnet#create-the-queue-storage-client
Therefore, the answer is QueueClient.

upvoted 7 times

Question #9 *Topic 3*

DRAG DROP -

You are maintaining an existing application that uses an Azure Blob GPv1 Premium storage account. Data older than three months is rarely used. Data newer than three months must be available immediately. Data older than a year must be saved but does not need to be available immediately.

You need to configure the account to support a lifecycle management rule that moves blob data to archive storage for data not modified in the last year.

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

Select and Place:

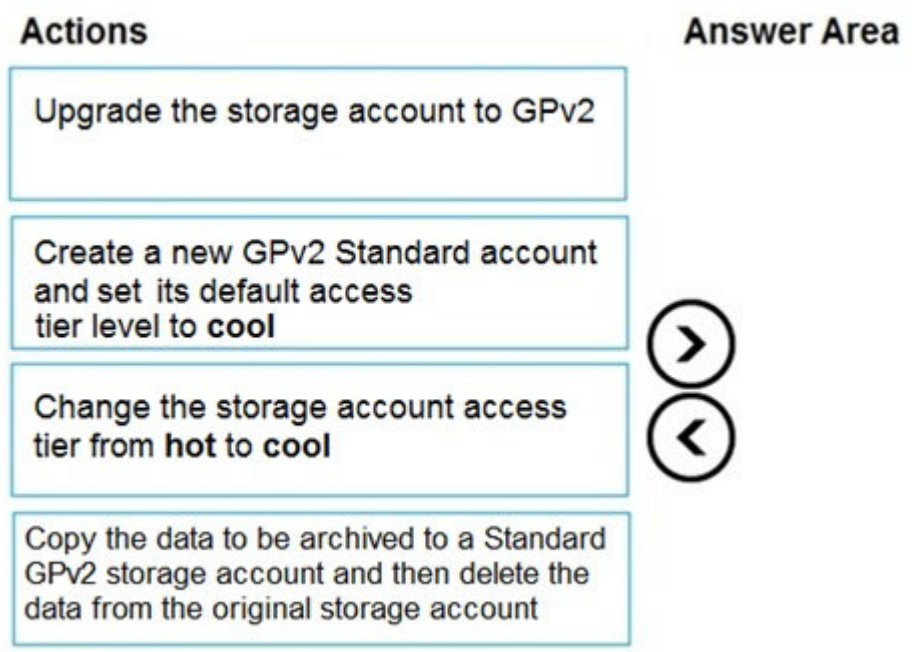

Correct Answer:

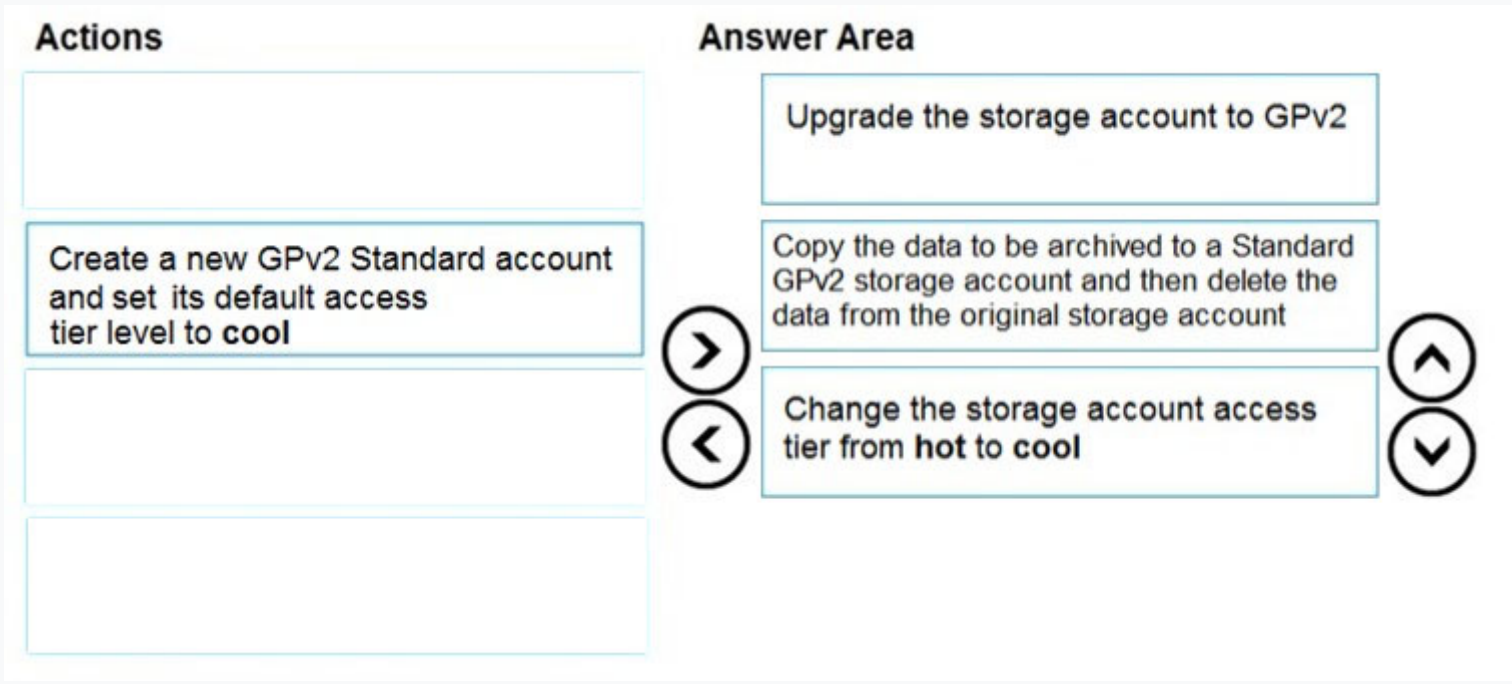

Step 1: Upgrade the storage account to GPv2

Object storage data tiering between hot, cool, and archive is supported in Blob Storage and General Purpose v2 (GPv2) accounts. General Purpose v1 (GPv1) accounts don't support tiering.

You can easily convert your existing GPv1 or Blob Storage accounts to GPv2 accounts through the Azure portal.

Step 2: Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account

Step 3: Change the storage account access tier from hot to cool

Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Only the hot and cool access tiers can be set at the account level. The archive access tier can only be set at the blob level.

Reference:

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers

sien Highly Voted 2 years, 6 months ago

Since we already have a premium P1 account with gpv1. Why not:
- Upgrade the existing one to GPv2
- Create a new GPV2 standard account with default access level to cool

- And then copy archive data to the GPV2 and delete the data from original storage account.
That makes sense to me.
upvoted 108 times

jay158 2 years, 4 months ago
Is there any requirement in question, which says set default access tier to COOL?
Please clarify.
upvoted 2 times

lighting 2 years, 4 months ago
I have 2 ideas:
1.
- One HOT for newer than 3 months
- One COOL for older than 3 months and Archive data.
2.
- One HOT for non-archived data. Can be accessed immediately. Because they just said "Older than 3 months data are rarely accessed" but didn't tell us anything about can it be accessed immediately.
- One COOL for archived data. Archived tier can just be set at blob level. "Only the hot and cool access tiers can be set at the account level. The archive access tier can only be set at the blob level" https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers
upvoted 3 times

Ram0202 2 years, 3 months ago
Why we need to Upgrade the existing one to GPv2 as we are creating a new one?
upvoted 14 times

warchoon 8 months, 1 week ago
To have additional account for the archive data?
upvoted 2 times

WillPassExam 2 years, 6 months ago
does this mean at the end, we have 2 GPv2 storage accounts, one access level is cool (for archive data) and the other is hot?
upvoted 2 times

sien 2 years, 6 months ago
I would say so.
upvoted 1 times

sien 2 years, 6 months ago
Also look at this blog:
https://www.apptio.com/blog/essential-guide-azure-blob-storage-pricing/
Only GPv2 and Blob storage accounts support tiering. If you are using GPv1, and you want to leverage tiering, convert your account to GPv2 through the Azure portal.
upvoted 4 times

ferut 2 years, 5 months ago
My choice as well
upvoted 1 times

mlantonis Highly Voted 2 years, 5 months ago
Step 1: Upgrade the storage account to GPv2
Object storage data tiering between hot, cool, and archive is supported in Blob Storage and General Purpose v2 (GPv2) accounts. General Purpose v1 (GPv1) accounts don't support tiering. You can easily convert your existing GPv1 or Blob Storage accounts to GPv2 accounts through the Azure portal.

Step 2: Create a new GPV2 standard account with default access level to cool

Step 3: Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account
upvoted 19 times

mlantonis 2 years, 5 months ago
Reference:

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers

https://docs.microsoft.com/en-us/azure/storage/common/storage-account-upgrade?tabs=azure-portal
upvoted 2 times

mlantonis 2 years, 5 months ago
Although Step 3 seems unusual and not necessary.
upvoted 4 times

ensa 2 years ago
every time with good explanation but step 3 needed because why not transfer the old data that needed to new one and delete the old one for saving cost
upvoted 1 times

tosm 1 year, 4 months ago

Agreed: 1. upgrade storage account 1 from GPv1 -> v2 (hot), 2. create storage account 2 GPv2 (cool) 3. copy data from account 1 to account 2. End result: account 1 (hot), account 2(cool and archive) both accounts can set lifecycle policy

upvoted 3 times

narenazure Most Recent 6 months, 2 weeks ago

Azure Blob storage lifecycle management offers a rich, rule-based policy for General Purpose v2 and Blob storage accounts.

upvoted 1 times

deepak_26 7 months, 1 week ago

Create , change , copy

Data stored in a premium block blob storage account cannot be tiered to Hot, Cool, or Archive using Set Blob Tier or using Azure Blob Storage lifecycle management. To move data, you must synchronously copy blobs from the block blob storage account to the Hot tier in a different account using the Put Block From URL API or a version of AzCopy that supports this API.

upvoted 1 times

motekim 7 months, 1 week ago

Just for reference: GPv2 supports Hot, Cool, and Archive tiers

upvoted 2 times

dimsok 8 months, 3 weeks ago

Upgrade-Create-Copy

upvoted 2 times

JamieS 9 months ago

Seems like we don't have the entire world view, and that there's already an existing standard v2 account elsewhere. If that's the case we're starting with v1 storage where the data currently is and we also have a default v2 storage that exists.

We upgrade the current v1 to v2 so we can access lifetime management
we copy the data to be archived to the standard v2 (that already existed, and by default has hot tier)
then we set the tier to be cool (where the archive data was just copied to).

I think the current answer is correct.

upvoted 2 times

micro9000 11 months, 3 weeks ago

The answer doesn't make any sense to me.
in the 2nd step, it says copy the data from old storage account to new one, it means there are 2 storage accounts, why you need to upgrade the existing?

upvoted 2 times

Akiu 1 year ago

Upgrade to GPv2
Copy data
Change tier to cool
You can't create the second account in a cool tier because of this:
Data stored in a premium block blob storage account cannot be tiered to hot, cool, or archive using Set Blob Tier or using Azure Blob Storage lifecycle management. To move data, you must synchronously copy blobs from the block blob storage account to the hot tier in a different account using ...
https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview#blob-lifecycle-management

upvoted 1 times

Azziet 1 year, 4 months ago

Upgrade
Create
Copy

upvoted 7 times

Evo_Morales 1 year, 6 months ago

Agree with voting/answer, but the question itself seems flawed. Need to learn not to read anything else into the questions and not add steps/requirements that yes, would make sense.

upvoted 3 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 5 times

massnonn 1 year, 8 months ago

for me not have sense - Upgrade the existing one to GPv2
So:
-create ..
-change..
-copy

upvoted 1 times

edengoforit 1 year, 9 months ago

Azure offers three types of storage accounts: General Purpose v2 (GPv2), General Purpose v1 (GPv1), and a dedicated object storage service—Blob Storage. GPv1 and GPv2 support the same storage types, but GPv2 supports Hot, Cool, and Archive tiers. Adopt GPv2 or Blob Storage to leverage tiering.

GPv2 offers up a variety of storage options (blob, files, queues, tables, and disks), performance tiers, and replication options not available with dedicated blob storage account types. Azure uses GPv2 to roll out new product enhancements (e.g., new redundancy options). GPv2 is the preferred storage account type if you want early access to new services. Azure storage account options are:

GPv2: Basic storage account type for blobs, files, queues, and tables. Recommended for most scenarios using Azure Storage.
GPv1: Legacy account type for blobs, files, queues, and tables. Use GPv2 accounts instead when possible.

upvoted 1 times

Yazhu 1 year, 9 months ago

Confusing here
What's the right answer??

upvoted 6 times

ning 2 years, 2 months ago

The goal should be
1. 0-90 days old --> hot
2. 91-365 days old --> cool
3. 366 or more --> archive

But I cannot see anything satisfy that

upvoted 9 times

j888 2 years, 2 months ago

The given answer is correct. I can only assume the second storage is in GPv2 Archived but wasn't been advised in the question.
It makes sense when you put this together.

1)So when you upgrade from GPV1 to Gpv2, it is defaulting to hot.
2)Take advantage of the hot tier pricing by moving your archive data to another Gpv2 storage that wasn't been advised in the question
3)No change this storage to cool tier to satisfy the requirement

upvoted 4 times

lms001 2 years, 3 months ago

since the solution is insisting on having three processes,,
1. Create a new gpv2 account
2. Copy archive data to the new gpv2 account and delete the data from the original account
3. Change the access tier from hot to cool

upvoted 8 times

Question #10 *Topic 3*

You develop Azure solutions.
You must connect to a No-SQL globally-distributed database by using the .NET API.
You need to create an object to configure and execute requests in the database.
Which code segment should you use?

A. new Container(EndpointUri, PrimaryKey);

B. new Database(EndpointUri, PrimaryKey);

C. new CosmosClient(EndpointUri, PrimaryKey);

Correct Answer: *C*
Example:
// Create a new instance of the Cosmos Client
this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)
//ADD THIS PART TO YOUR CODE
await this.CreateDatabaseAsync();
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql-api-get-started

Community vote distribution
C (100%)

Mr2302682 Highly Voted 3 years, 2 months ago
The Answer is correct!
upvoted 46 times

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: C

Azure Cosmos DB is a fully managed NoSQL database for modern app development. Single-digit millisecond response times, and automatic and instant scalability, guarantee speed at any scale.

// Create a new instance of the Cosmos Client
this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)
//ADD THIS PART TO YOUR CODE
await this.CreateDatabaseAsync();

Reference:

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient?view=azure-dotnet
https://docs.microsoft.com/en-us/azure/cosmos-db/sql-api-get-started
upvoted 20 times

dimsok Most Recent 8 months, 3 weeks ago
I dont get this question, Cosmos DB is a No-SQL globally-distributed database, but there are others that .NET can connect.
upvoted 1 times

Nhiendo 10 months ago
Correct Answer: C
upvoted 1 times

AvinashVarma 10 months, 2 weeks ago
Selected Answer: C
CosmosClient(String, String, CosmosClientOptions)

https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient?view=azure-dotnet
upvoted 1 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: C
C. new CosmosClient(EndpointUri, PrimaryKey);
upvoted 1 times

rohitpatil113 1 year, 2 months ago

Selected Answer: C

The Answer is correct!

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: C

C is correct answer.
Global fully managed No-SQL DB = CosmosDB

upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 2 times

Freidrich 1 year, 8 months ago

Selected Answer: C

The correct answer is C.

upvoted 1 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with highly voted answer.

upvoted 5 times

Mev4953 1 year, 9 months ago

Got this in the exam 01/22

upvoted 5 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 5 times

sanjayrawat 2 years ago

Answer : C

upvoted 2 times

debanjan10 2 years ago

Order:
Cosmos Client > Database > Container > Item

upvoted 3 times

SlavMar 2 years, 4 months ago

But why we assume it is Cosmos?
App needs to access some No-SQL database that may already exist - maybe it is Cosmos, maybe not

upvoted 3 times

DParekh 2 years, 5 months ago

There are two ways by which we can created cosmosclient instance.
1. CosmosClient cosmosClient = new CosmosClient(connectionString);
2. CosmosClient cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)

So answer C is correct.

upvoted 2 times

Question #11 *Topic 3*

You have an existing Azure storage account that stores large volumes of data across multiple containers.

You need to copy all data from the existing storage account to a new storage account. The copy process must meet the following requirements:

- Automate data movement.
- Minimize user input required to perform the operation.
- Ensure that the data movement process is recoverable.

What should you use?

A. AzCopy

B. Azure Storage Explorer

C. Azure portal

D. .NET Storage Client Library

Correct Answer: *A*

You can copy blobs, directories, and containers between storage accounts by using the AzCopy v10 command-line utility.

The copy operation is synchronous so when the command returns, that indicates that all files have been copied.

Reference:

https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy

Community vote distribution

A (100%)

kondapaturi Highly Voted 2 years, 4 months ago

Answer – AzCopy, The Azcopy tool can be used to copy data from one storage account to another. You can use the tool within automation scripts to ensure the data can be copied automatically.

upvoted 37 times

mlantonis Highly Voted 2 years, 5 months ago

Correct Answer: C

Azure Storage Explorer uses AzCopy to perform all of its data transfer operations. But in this questions, there is a requirement to minimize user interaction which is why AzCopy is more appropriate.

Reference:

https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10

https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy

upvoted 10 times

alperc 1 year, 12 months ago

please stop commenting..you are always misanswering.

upvoted 6 times

altafpatel1984 1 year, 11 months ago

But read very first requirement. It needs to be Automated. I agree with alperc, you are misanswering, even such a silly question.

upvoted 4 times

azurelearner666 2 years, 4 months ago

AzCopy is A... please do not write confusing responses...
response is already correct, A - AzCopy

upvoted 45 times

NombreFalso Most Recent 8 months, 2 weeks ago

Selected Answer: A

Typical AzCopy

upvoted 2 times

AvinashVarma 10 months, 2 weeks ago

Selected Answer: A

AzCopy
Reference: https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: A

A. AzCopy

upvoted 1 times

Unpredictable_m 1 year ago

Correct answer is A
AzCopy is a command-line utility that you can use to copy blobs or files to or from a storage account

upvoted 1 times

Chaitya 1 year, 2 months ago

AzCopy

upvoted 1 times

ghuna 1 year, 2 months ago

Selected Answer: A

AzCopy is best option to move data around in blob storage

upvoted 1 times

xRiot007 1 year, 4 months ago

You will use AzCopy. It can be automated in a script, reducing user interaction and you can also recover

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: A

A is the correct answer.
Everytime.

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose the A. AzCopy

upvoted 2 times

idroj 1 year, 8 months ago

Selected Answer: A

AzCopy is the correct answer

upvoted 3 times

Exocsinstr 1 year, 10 months ago

https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-configure
Ans: should be AzCopy. This link show you can recover the AzCopy from where it failed.

azcopy jobs resume <job-id> --source-sas="<sas-token>" --destination-sas="<sas-token>"
Azure Data Factory in the Azure portal may need much configuration by user to setup the pipeline to perform copy.

upvoted 1 times

francis6170 2 years, 4 months ago

got this in the exam :)

upvoted 6 times

jeet1985 2 years, 4 months ago

Then ? please suggest the correct answer?

upvoted 4 times

debanjan10 2 years ago

A lot of these questions are common in the exam it seems.

upvoted 1 times

UnknowMan 2 years, 5 months ago

The Answer is correct!

upvoted 4 times

glam 2 years, 5 months ago

correct

upvoted 3 times

AzureAz204Fan 2 years, 5 months ago

Please note that Azure Storage Explorer uses AzCopy to perform all of its data transfer operations. But in this questions, there is a requirement to minimize user interaction which is why AzCopy is more appropriate.
https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10

upvoted 4 times

Question #12 *Topic 3*

DRAG DROP -

You are developing a web service that will run on Azure virtual machines that use Azure Storage. You configure all virtual machines to use managed identities.

You have the following requirements:

☞ Secret-based authentication mechanisms are not permitted for accessing an Azure Storage account.

☞ Must use only Azure Instance Metadata Service endpoints.

You need to write code to retrieve an access token to access Azure Storage. To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

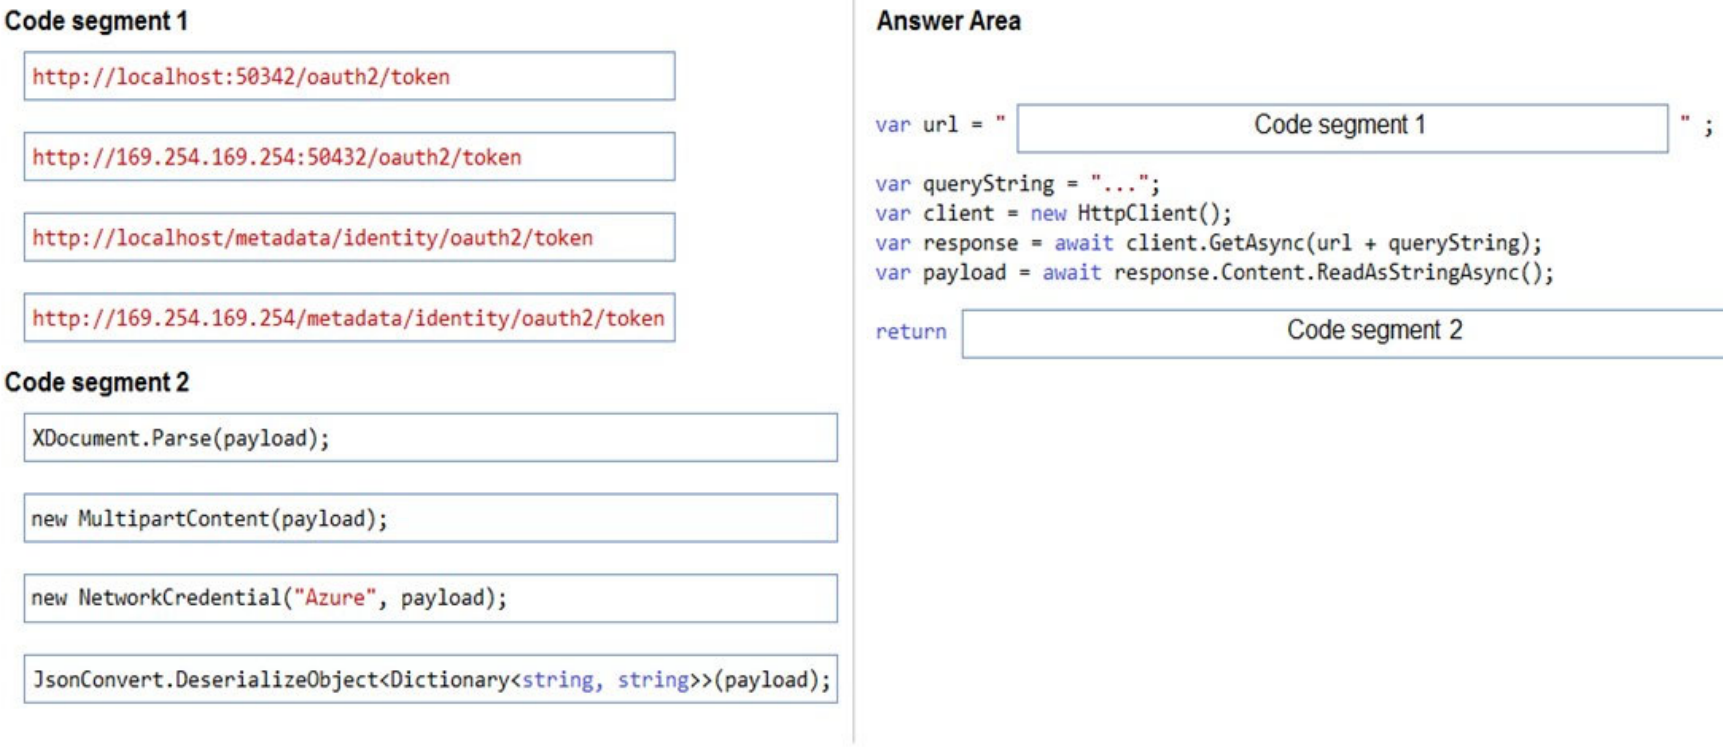

Correct Answer:

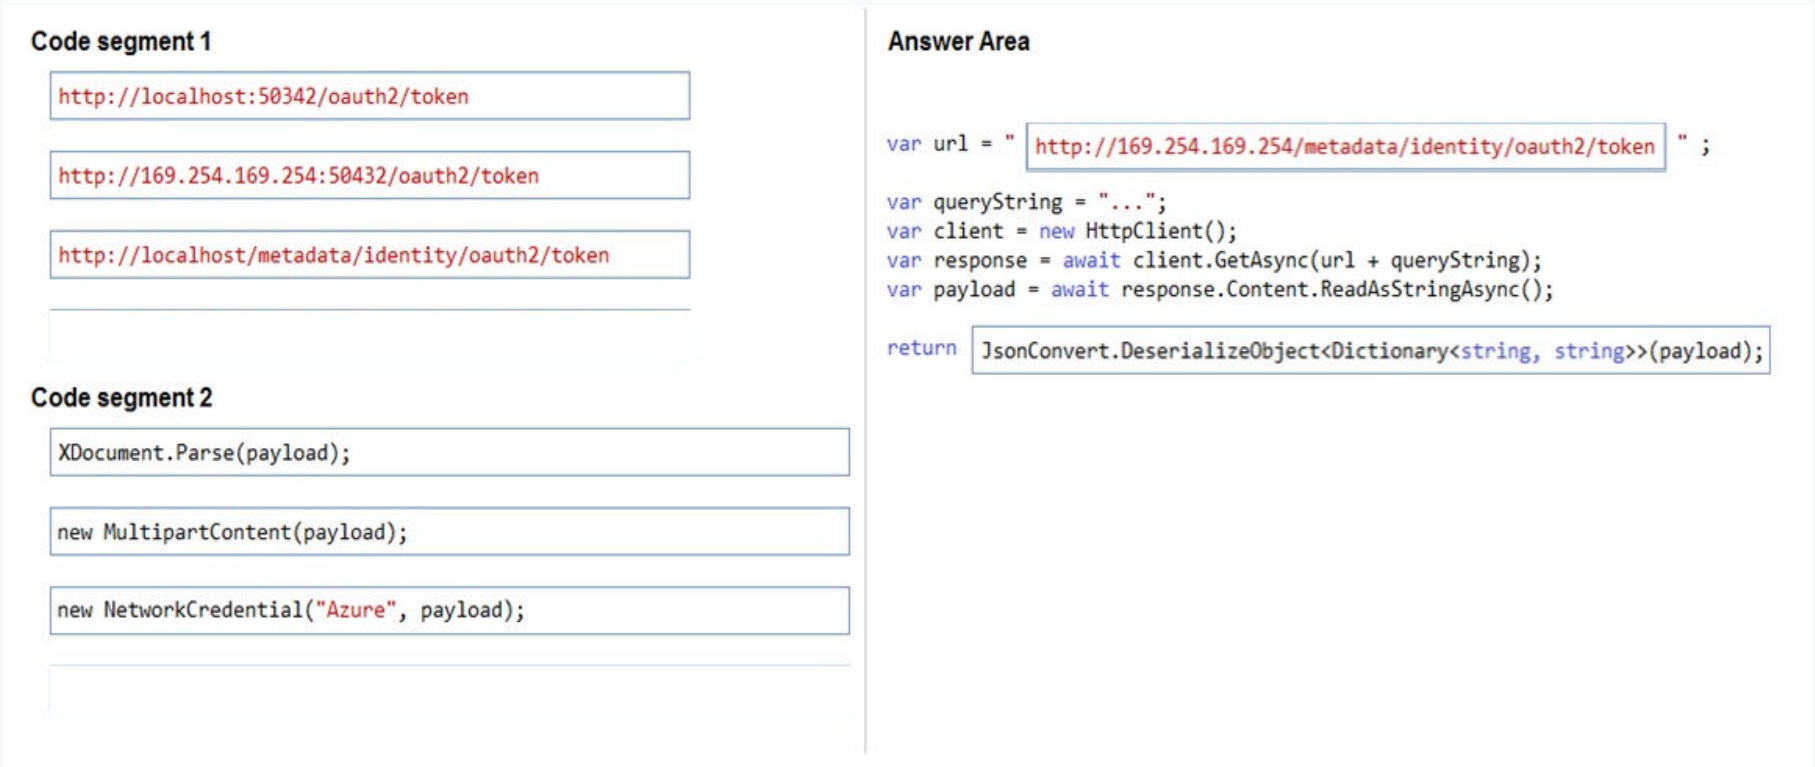

Azure Instance Metadata Service endpoints "/oauth2/token"

Box 1: http://169.254.169.254/metadata/identity/oauth2/token

Sample request using the Azure Instance Metadata Service (IMDS) endpoint (recommended):

GET 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' HTTP/1.1

Metadata: true

Box 2: JsonConvert.DeserializeObject<Dictionary<string,string>>(payload);

Deserialized token response; returning access code.

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token

https://docs.microsoft.com/en-us/azure/service-fabric/how-to-managed-identity-service-fabric-app-code

mlantonis Highly Voted 2 years, 5 months ago

Box 1: http://169.254.169.254/metadata/identity/oauth2/token
Sample request using the Azure Instance Metadata Service (IMDS) endpoint (recommended):
GET 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' HTTP/1.1
Metadata: true

Box 2: JsonConvert.DeserializeObject<Dictionary<string,string>>(payload);
Deserialized token response; returning access code.

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token

https://docs.microsoft.com/en-us/azure/service-fabric/how-to-managed-identity-service-fabric-app-code

upvoted 58 times

Knightie 1 year, 2 months ago

IMDS is a REST API that's available at a well-known, non-routable IP address (169.254.169.254). You can only access it from within the VM.
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/instance-metadata-service?tabs=windows

upvoted 7 times

warchoon 8 months, 1 week ago

Powerful Words And The Magic Of Abracadabra :)

upvoted 1 times

clarionprogrammer Highly Voted 2 years, 6 months ago

It is correct.
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token

upvoted 12 times

dddddd111 Most Recent 1 month, 1 week ago

I got this same question. Provided answers are correct. (Note: I failed the exam 20/9/23. I only scored 644 and I felt bad. I think because many questions here in Examtopics are not accurate. I suggest following the most voted answers and don't just not rely on Examtopics answers. At the beginning of the exam, you will be asked which programming languages you want to use. C#/Python. I chose C#. Also, I just want to add that some questions here are really in the actual exams, but the choices are written and formatted differently. Please be aware of that. Goodluck. I feel bad for failing it, but I want to retake next month. I will try Python. T_T

upvoted 5 times

Kanasan 2 months ago

Correct!
https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token

upvoted 2 times

mihailos 2 months ago

Got it in exam 28/08/23. Went with proposed answer. Scored 912

upvoted 1 times

macobuzi 2 months, 1 week ago

Answer is correct!

upvoted 1 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 3 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 1 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 4 times

Esward 8 months, 2 weeks ago

It was there in 13 Feb 2023 exam

upvoted 6 times

martijn44 1 year ago

I was taught to use $IDENTITY_ENDPOINT but indeed that always seems to refer to the 169.254.169.254 address

upvoted 1 times

mijkol 1 year, 4 months ago

I get confused because of the IP, I get we shouldn't go for the local one, but why is this an AIPIPA IP adress?

is there a reason for that?

upvoted 2 times

Rini100 1 year, 6 months ago

Got this on 20 Apr 2022

upvoted 2 times

kondapaturi 2 years, 4 months ago

BOX1 - http://169.254.169.254/metadata/identity/oauth2/token ,To get the metadata from the local service on the machine, the right URL is http://169.254.169.254/metadata/identity/oauth2/token

BOX2 – JsonConvert.DeserializeObject<string,string>(payload)

You can deserialize the response using the JsonConvert.DeserializeObject method. You can then get a dictionary collection and then get the access key from there.

upvoted 5 times

goatlord 2 years, 4 months ago

Correct

upvoted 2 times

UnknowMan 2 years, 5 months ago

Correct ! => https://docs.microsoft.com/fr-fr/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token

upvoted 2 times

glam 2 years, 5 months ago

correct

upvoted 2 times

DRAG DROP -

You are developing a new page for a website that uses Azure Cosmos DB for data storage. The feature uses documents that have the following format:

```
{
      "name": "John",
      "city" : "Seattle"
}
```

You must display data for the new page in a specific order. You create the following query for the page:

```
SELECT*
FROM People p
ORDER BY p.name, p.city DESC
```

You need to configure a Cosmos DB policy to support the query.

How should you configure the policy? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

JSON segments

- orderBy
- sortOrder
- ascending
- descending
- compositeIndexes

Answer Area

```
{
   "automatic": true,
   "ngMode": "Consistent",
   "includedPaths": [
     {
        "path": "/*"
     }
   ], "excludedPaths": [],
   "  [blank]  ": [
     [
            {                                                    }
             "path": "/name", "order": "descending"

            },
            {
               "path": "/city", "order": "  [blank]  "

            }
        ]
   ]
```

Correct Answer:

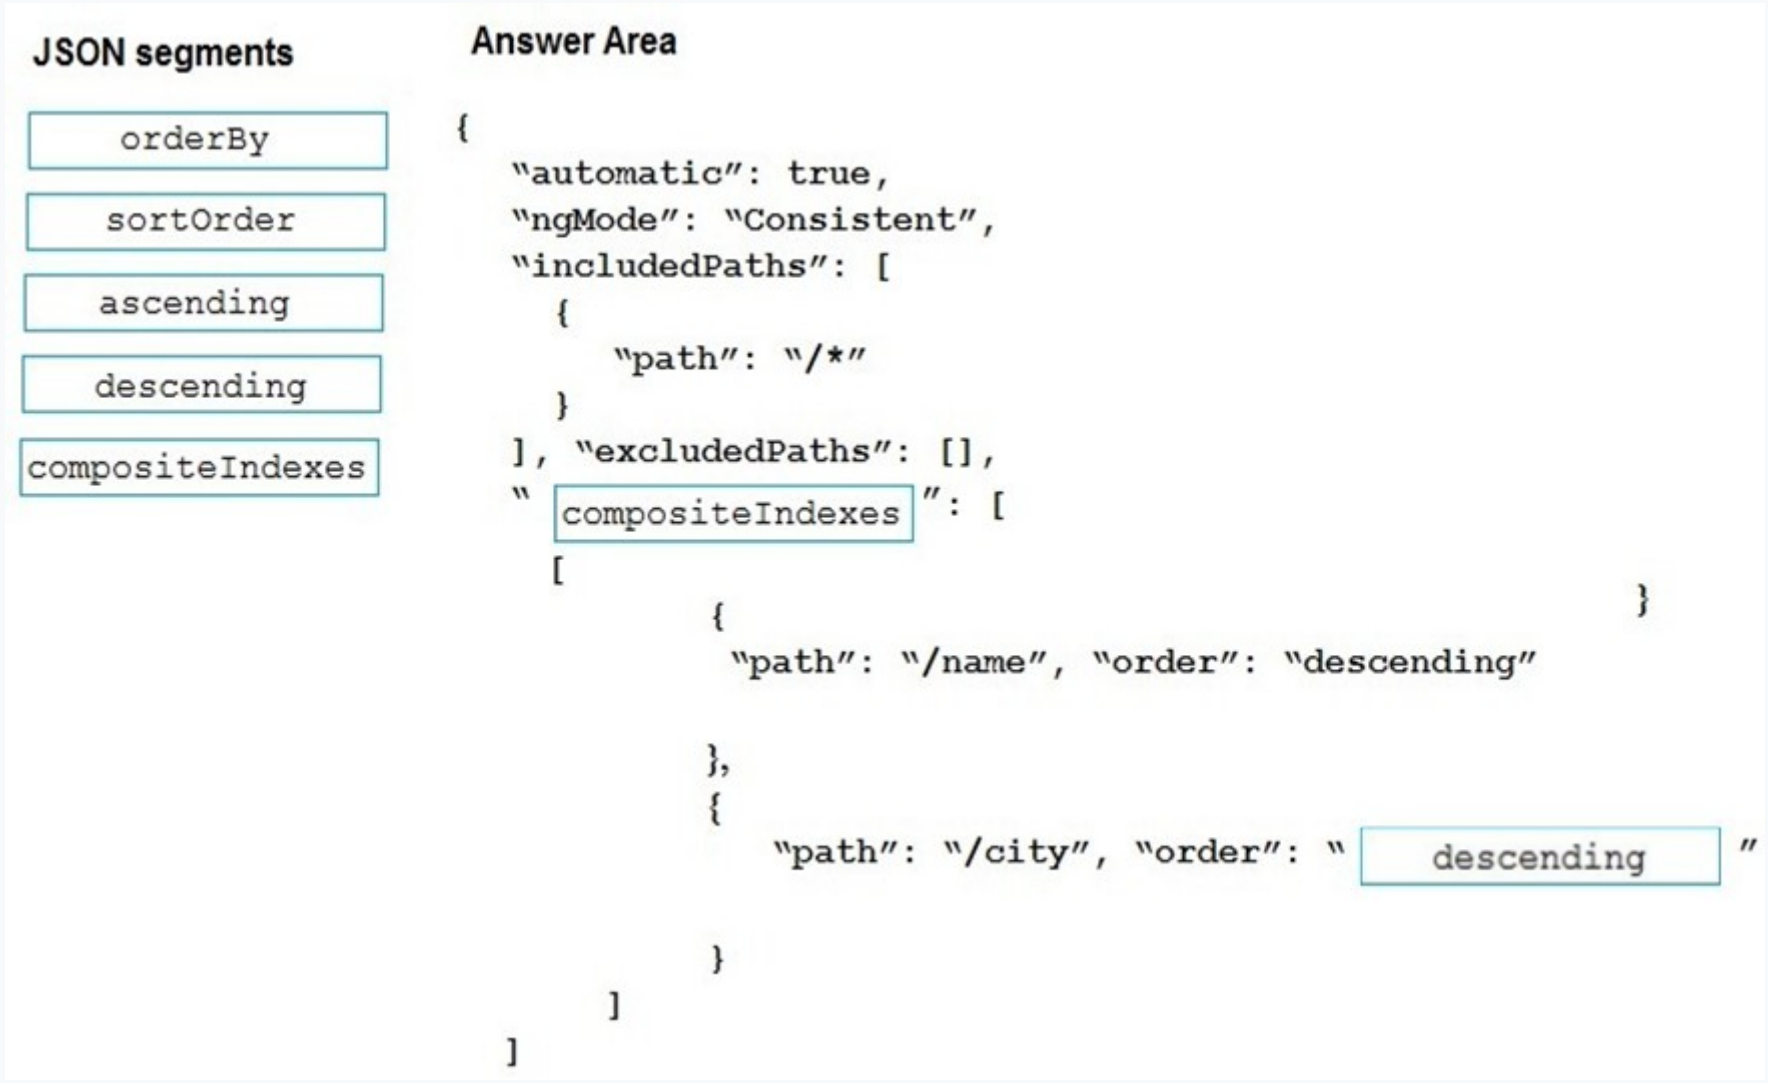

Box 1: compositeIndexes -
You can order by multiple properties. A query that orders by multiple properties requires a composite index.

Box 2: descending -
Example: Composite index defined for (name ASC, age ASC):
It is optional to specify the order. If not specified, the order is ascending.

```
{
"automatic":true,
"indexingMode":"Consistent",
"includedPaths":[
{
"path":"/*"
}
],
"excludedPaths":[],
"compositeIndexes":[
[
{
"path":"/name",
},
{
"path":"/age",
}
]
]
}
```

kayleena93 Highly Voted 2 years, 11 months ago

ORDER BY queries on multiple properties:
The composite index also supports an ORDER BY clause with the opposite order on all paths.

So I think it's about reversed index to the query. Answer should be 'ascending'. You cannot support ASC (default), DESC query with DESC, DESC index.

upvoted 85 times

[Removed] 3 months, 2 weeks ago

No. It supports, but it is not required. You can have (ASC, ASC), (DESC, DESC), (ASC, DESC), (DESC, ASC). So the answer is Descending

upvoted 1 times

john4p 1 year, 10 months ago

The problem here is the SQL that makes many people think that
"ORDER BY p.name, p.city DESC"
means it's ordered by name and city both descending.
But the DESC only applies to city. name is ASC - this would be less confusing:
"ORDER BY p.name ASC, p.city DESC"

Thus in the JSON you can only state ascending+descending or the opposite: descending+ascending.
Since descending for name is already set the answer is "ascending".

At first I had misread the SQL wrong myself and didn't understand kayleena's comment right away.

upvoted 28 times

GCMan Highly Voted 2 years, 11 months ago

"name" field should be marked ascending (default if not specified). It's mislabeled

upvoted 34 times

cloud_exam1 2 years, 11 months ago

I think so. The answer is correct, but the name field should be marked ascending(default).
We can find an example in the following link.
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-indexing-policy?tabs=dotnetv2%2Cpythonv3
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-indexing-policy?tabs=dotnetv2%2Cpythonv3

upvoted 15 times

warchoon 8 months, 1 week ago

The client wants to walk with his back to the front. Who are we to tell him what to do?

upvoted 1 times

scottmct 2 years, 8 months ago

NO. Box 2 is "ascending"
See explanation here:
https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy#order-by-queries-on-multiple-properties
"The composite index also supports an ORDER BY clause with the __opposite order on all paths__."
The table in the section also shows an example similar to this question.

upvoted 30 times

[Removed] 3 months, 2 weeks ago

No. Box 2 is Descending

upvoted 1 times

ArturKon 1 year, 1 month ago

I agree with you. For example index on (A asc, B asc) works for queries with ORDER BY (A asc, B asc) and (A desc, B desc).
Not working with ORDER BY (A asc, B desc), (A desc, B asc) or even (B asc, A asc).

upvoted 1 times

Rockm0uld 1 year, 7 months ago

the table also seems to suggest DESC followed by ASC is not supported by composite index? Confused

upvoted 2 times

Rockm0uld 1 year, 7 months ago

Looking again i think it refers to " What it will not support is non-matching clauses. if ASC, ASC or DESC, DESC will not match our question." as mentioned by edengoforit.
Has to be opposites to make any sense.

upvoted 1 times

p2006 Most Recent 1 month ago

Got on 9/25/2023
compositeindexes
ascending

upvoted 2 times

leviatas 1 month ago

Got this quesiton in examn - 2023.09.25. Got Case Study Contoso

upvoted 1 times

mihailos 2 months ago

Got it in exam 28/08/23. Went with proposed answer. Scored 912

upvoted 1 times

basquiat 2 months, 3 weeks ago

Got this one on 2023-08-08

upvoted 1 times

applepie 3 months ago

got this question today, answer compositeIndex, descending - 7/30/2023, score 895/1000

upvoted 3 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 1 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed (city should be "asc"). Scored 917

upvoted 3 times

JH81 4 months ago

Got this on 6/28/2023 and passed with 850. Went with composite and ascending.

upvoted 2 times

kiro_kocha 4 months, 1 week ago

According to Documentation: https://learn.microsoft.com/en-us/azure/cosmos-db/index-policy#composite-indexes
ORDER BY queries on multiple properties:
Composite Index Sample ORDER BY Query Supported by Composite Index?
(name ASC, age ASC) SELECT * FROM c ORDER BY c.name ASC, c.age asc Yes
(name ASC, age ASC) SELECT * FROM c ORDER BY c.age ASC, c.name asc No
(name ASC, age ASC) SELECT * FROM c ORDER BY c.name DESC, c.age DESC Yes
(name ASC, age ASC) SELECT * FROM c ORDER BY c.name ASC, c.age DESC No
(name ASC, age ASC, timestamp ASC) SELECT * FROM c ORDER BY c.name ASC, c.age ASC, timestamp ASC Yes
(name ASC, age ASC, timestamp ASC) SELECT * FROM c ORDER BY c.name ASC, c.age ASC No
So correct answer is compositeIndexes and descending
Look the table:
Composite Index: (name ASC, age ASC)
Sample ORDER BY Query: SELECT * FROM c ORDER BY c.name DESC, c.age DESC
Supported by Composite Index?: Yes

upvoted 2 times

RikinPatel 4 weeks, 1 day ago

This perfect explaination... Thanks...

upvoted 1 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam at 10-June-2023

upvoted 1 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

My case:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 2 times

sarmaria 7 months, 2 weeks ago

Got this question in the exam on 16/03/2023. Chosen composite and ascending. Scored 846. Make sure to prepare for case studies. I got city and lights case study.

upvoted 3 times

dimsok 8 months, 3 weeks ago

Example is wrong, name should be Ascending and City Descending

upvoted 1 times

rgullini 6 months, 3 weeks ago

Wrong. Read the information about Composite Index in Cosmos DB

upvoted 1 times

Esward 9 months ago

Composite index. p.name ascending, p.city descending
https://learn.microsoft.com/en-us/azure/cosmos-db/index-policy#composite-indexes

upvoted 3 times

Lucifer14 10 months ago

In 28-12-2022 exam

upvoted 1 times

Question #14 *Topic 3*

HOTSPOT -

You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway.

Data from traffic sensors are stored in Azure Event Hub.

Traffic data is consumed by four departments. Each department has an Azure Web App that displays the time series-based reports and contains a WebJob that processes the incoming data from Event Hub. All Web Apps run on App Service Plans with three instances.

Data throughput must be maximized. Latency must be minimized.

You need to implement the Azure Event Hub.

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

NOTE: Each correct selection is worth one point.

Hot Area:

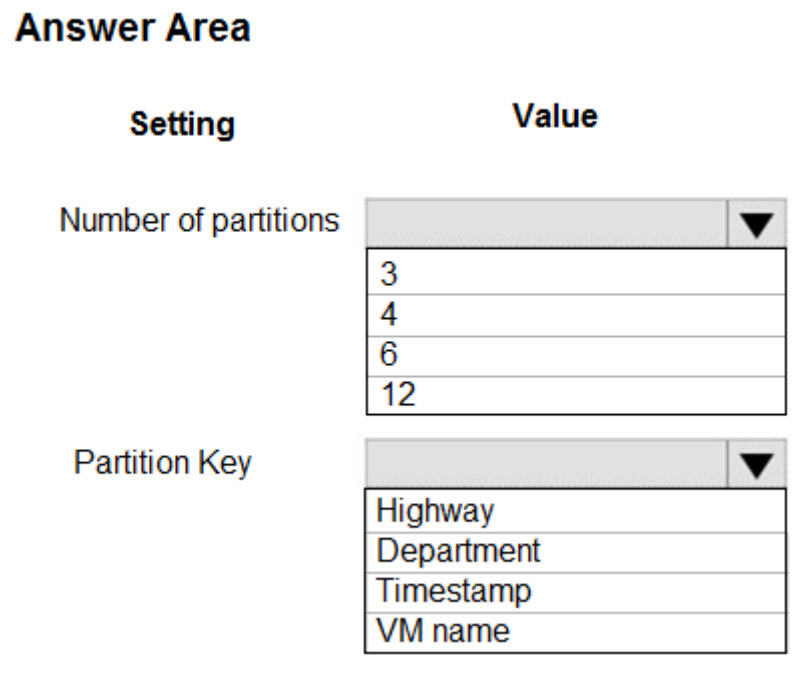

Correct Answer:

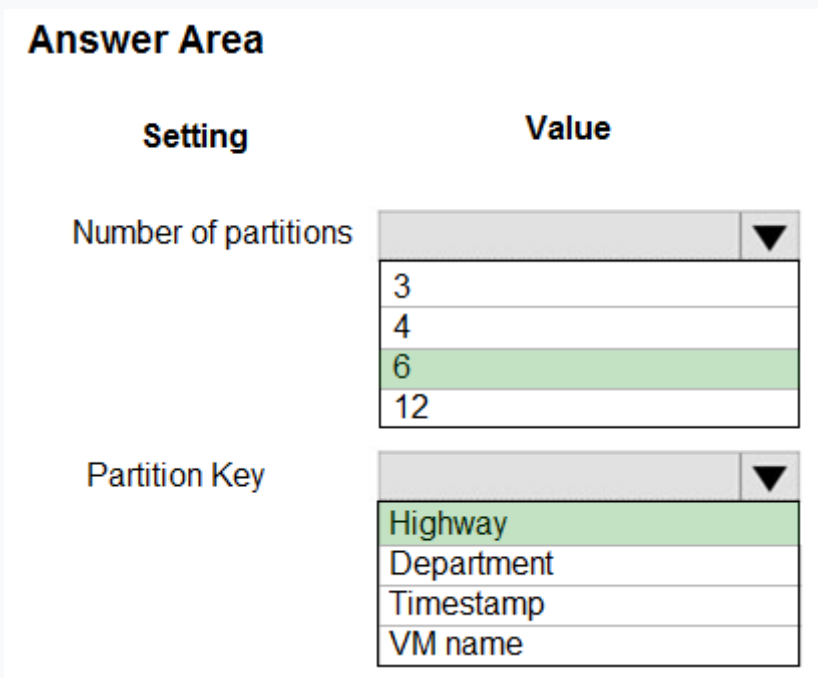

Box 1: 6 -

The number of partitions is specified at creation and must be between 2 and 32.

There are 6 highways.

Box 2: Highway -

Reference:

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features

dantest Highly Voted 2 years, 11 months ago

Partitions relate to producers - and the logical way to partition the incoming data is by the only value you have at that point, the highway name/id. So the selected answer is correct (6 Partitions, by Highway).

People are getting confused by the departments which would actually each be an event consumer with an associated Consumer Group which would have it's own isolated view of each of the highway partitions.

upvoted 129 times

Robert12345Robert 2 years, 4 months ago

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-faq
The number of partitions in an event hub directly relates to the number of concurrent readers you expect to have

upvoted 9 times

PhilLI 1 year, 11 months ago

If you read the features page you'll understand that this is not the simple answer to the question. https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features
Event receivers organised in consumer groups read all partitions. Via a leasing mechanism the receiver within a consumer group can make sure they don't read duplicate data.
Partitions are more related to Event producers (can be 1-on-1)

upvoted 1 times

Dani_ac7 1 year, 2 months ago

If we have 4 groups of consumers and each consumer group, according to the theory, reads independently and isolated from the rest, highways and 6 partitions makes a lot of sense

upvoted 2 times

MiraA 2 years ago

One thought...
The assignment mentions "Partition Key" (not "Partition Id").

"Producers can provide a value for the event key. When they do, a hashing-based partitioner determines a hash value from the key. The event then goes to the partition associated with that hash value."
https://docs.microsoft.com/en-us/azure/architecture/reference-architectures/event-hubs/partitioning-in-event-hubs-and-kafka#distribute-events-to-partitions

So it is possible for 2 or more highways to be - by a chance - hashed to a single partition leaving 1 or more partitions idle at all. If really unlucky then all 6 highways would be hashed to one partition.

The very same situation is with Highway, Department and VM name used as the "Partition Key" as they are discrete values (6 highways, 4 departments, N virtual machines). The Timestamp could do a better job - but "spraying" data from one highway across all partitions...

The best job could do "Partition Id".
"Producers can specify a partition ID with an event. The event then goes to the partition with that ID." So each highway could get hardcoded its own partition preserving order of the data.

upvoted 2 times

PhilLI 1 year, 11 months ago

I guess you're right, although losing maximum availability.

upvoted 1 times

minaritochuck 1 year, 8 months ago

it's a best practice for publishers(producers) to remain unaware of the specific partitioning model chosen for an event hub and to only specify a partition key that is used to consistently assign related events to the same partition.

upvoted 1 times

minaritochuck 1 year, 8 months ago

oh nevermind, I guess it would be most effective when partitions num matches the producers num.

upvoted 1 times

sasisang Highly Voted 2 years, 11 months ago

The answer should be 4 and Highway.
Exam Topics - Please provide correct answers. What is the use of buying questions on your site...if you are not sure of the answer yourself

upvoted 40 times

kwaazaar 2 years, 6 months ago

But theres 6 highways, so why not 6 partitions?

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

There are 4 consumers.
Partitions are a data organization mechanism that relates to the downstream parallelism required in consuming applications. The number of partitions in an event hub directly relates to the number of concurrent readers you expect to have.
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-faq

upvoted 2 times

warchoon 8 months, 1 week ago

"The event publisher is only aware of its partition key, not the partition to which the events are published. This decoupling of key and partition insulates the sender from needing to know too much about the downstream processing. A per-device or user unique identity makes a good partition key, but other attributes such as geography can also be used to group related events into a single partition.

Specifying a partition key enables keeping related events together in the same partition and in the exact order in which they arrived. The partition key is some string that is derived from your application context and identifies the interrelationship of the events. A sequence of events identified by a partition key is a stream. A partition is a multiplexed log store for many such streams."

upvoted 1 times

guardna 2 years, 6 months ago

But each department has 3 instances of the Web App/Job. So isn't there 12 consumers?

upvoted 6 times

jessicazheng 1 year, 9 months ago

https://www.linkedin.com/pulse/azure-event-hub-understanding-designing-partitions-unit-kamal-pathak

upvoted 1 times

BrettusMaximus 2 years, 6 months ago

There are 6 highways and 6 reports. Each department only needs to read one partition to produce their report on that one highway. If you had 4 partitions you would have to duplicate all the data 4 times (BAD) and then each department would need to read all the data and filter the data for one report (VERY BAD).

upvoted 11 times

BrettusMaximus 2 years, 6 months ago

When you create the Hub it does not know either the Departments or the VMs

upvoted 3 times

adilkhan Most Recent 7 months ago

6 and highway chatGPT

upvoted 1 times

macobuzi 2 months, 1 week ago

Lol , i asked ChatGPT and it said 12. 4 departments * 3 instances = 12 partitionkey needed

upvoted 1 times

RealRaymond 8 months, 3 weeks ago

12 (4*3) physical partitions; timestamp as partition key

upvoted 1 times

dimsok 8 months, 3 weeks ago

The number of partitions to use in Azure Event Hub depends on the desired level of parallelism and the expected ingress rate. For maximum data throughput, you would want to use as many partitions as possible. Each partition allows for a separate stream of events to be processed in parallel, increasing the ingress rate. On the other hand, having too many partitions can lead to increased latency, as the events must be divided among the partitions

Answer is 12.

The partition key is a value that determines which partition an event should be sent to. The events with the same partition key are guaranteed to be ordered and processed by the same partition. In a traffic monitoring system, the partition key could be related to the highway being monitored.

Answer is highway.

upvoted 1 times

michaelfastov 11 months ago

https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features
"Your own applications must be able to keep up with processing the volume of events that are being sent into an event hub"
In our case events are highway reports (6). 4 applications are consumers.
6 is correct.

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

selected answer is correct

upvoted 1 times

Alluru 1 year ago

Number of partitions – 6 – selected because there are 6 highways.
Partition Key – Highway – map the incoming events into specific partition. Partition-Key is sender supplied value passed into event-hub.
https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions

upvoted 1 times

Ranzzzan 1 year, 6 months ago

got this on exam

upvoted 3 times

Rini100 1 year, 6 months ago

Got this on 20 Apr 2022

upvoted 3 times

henry1985 1 year, 6 months ago

Thinking this is about how to choose partions when there's an ordered delivery requirement. The question states 'The system produces time series analysis-based reports for each highway' and there's six highways.

Event Hubs ensures that all events sharing a partition key value are stored together and delivered in order of arrival. So if 6 partitions are selected then the consumers can read the events in order and create the time series reports.

Ref. https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#publishing-an-event

upvoted 1 times

xRiot007 1 year, 7 months ago

Partitions - 4, because there are 4 departments. Partitioning is about slicing data for concurrent reading, it has nothing to do with duplication of data, like some silly "explanations" below.

The partition key should be the Highway because that will be your main filter when creating your reports and is also a good identifier for the event source.

upvoted 2 times

MohmmadFayez 1 year, 10 months ago

To partition we will depends on the producer which is highways,
for consumers(Applications) we will use 4 "consumer groups" ,one consumer group for each application ,so each App will have his specific view for the data,each consumer group will include 3 instances of the app, no more than one instance from the same consumer group can share the same portion at the same time.
So since we will have 6 consumer groups ,then we can assign 2 portions for each instance

upvoted 1 times

dho 2 years ago

I would rather choose 12 partitions (4 dep with 3 instances = 12 readers)
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-faq#partitions
The number of partitions in an event hub directly relates to the number of concurrent readers you expect to have.

and Highway
You can use a partition key to map incoming event data into specific partitions for the purpose of data organization.

upvoted 3 times

tackchen 2 years ago

4*3=12 partitions is absolutely wrong. partition is used to organize data source and is used by consumers separately. 6 partitions, 4 consumer groups, in each consumer group there are 3 receivers.

upvoted 2 times

mlantonis 2 years, 5 months ago

Box 1: 6

Box 2: Highway

Reference:

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features

upvoted 17 times

kishe 2 years, 5 months ago

The answer should be 6 and Highway according to the architecture of event hubs: https://docs.microsoft.com/en-us/azure/event-hubs/media/event-hubs-about/event_hubs_architecture.svg

On the consumer side, there should be 4 consumer group, and each group has 3 readers.

upvoted 6 times

finnishr 2 years, 2 months ago

This is 100% correct

upvoted 3 times

Question #15 *Topic 3*

DRAG DROP -

You are developing a microservices solution. You plan to deploy the solution to a multinode Azure Kubernetes Service (AKS) cluster.

You need to deploy a solution that includes the following features:

- reverse proxy capabilities
- configurable traffic routing
- TLS termination with a custom certificate

Which components should you use? To answer, drag the appropriate components to the correct requirements. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

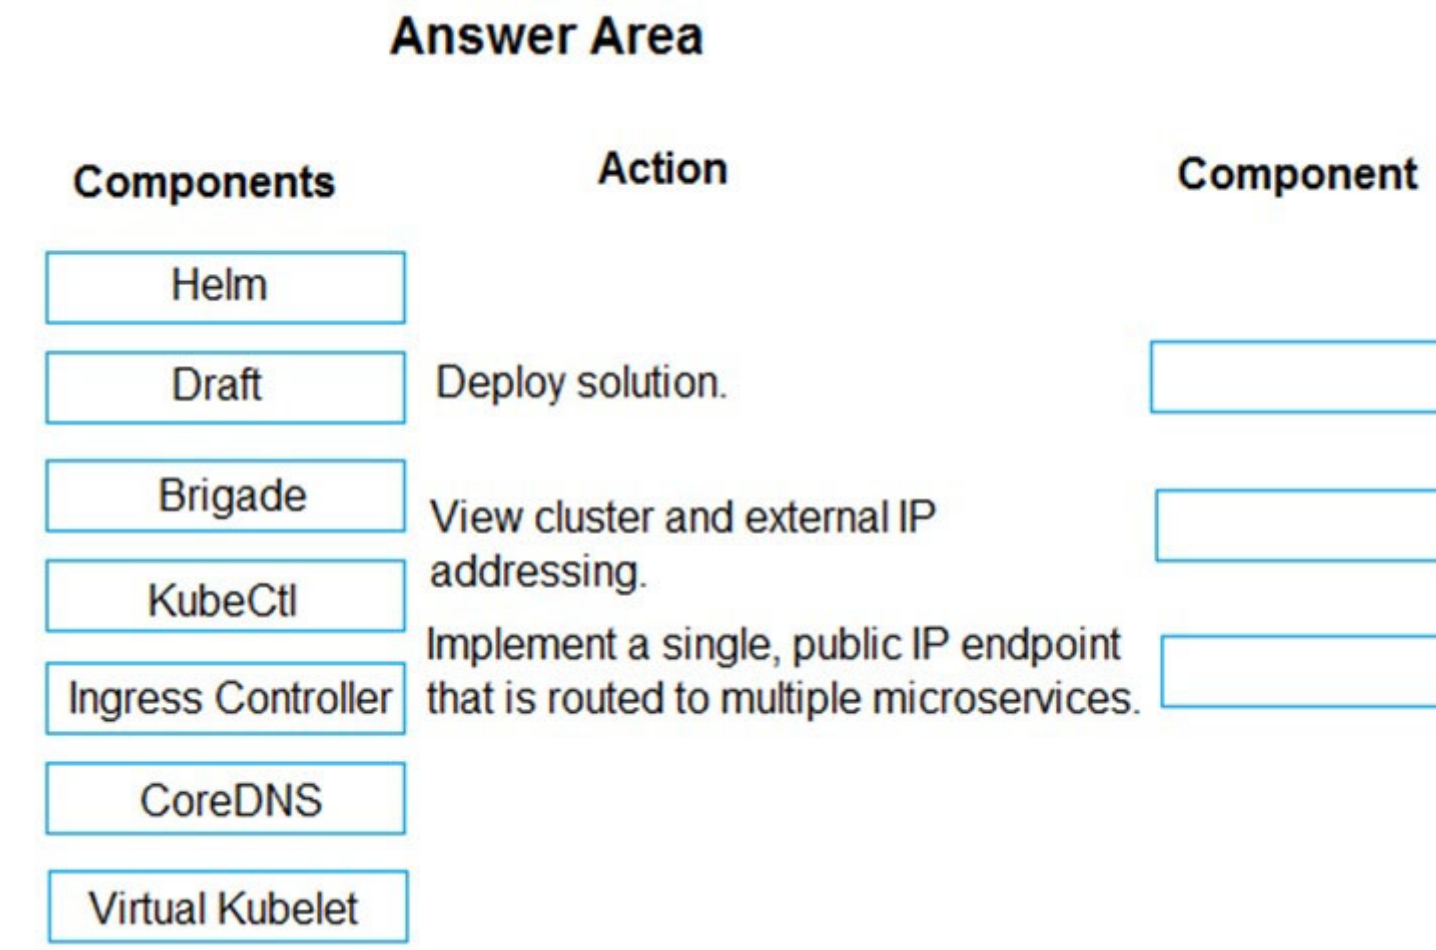

Correct Answer:

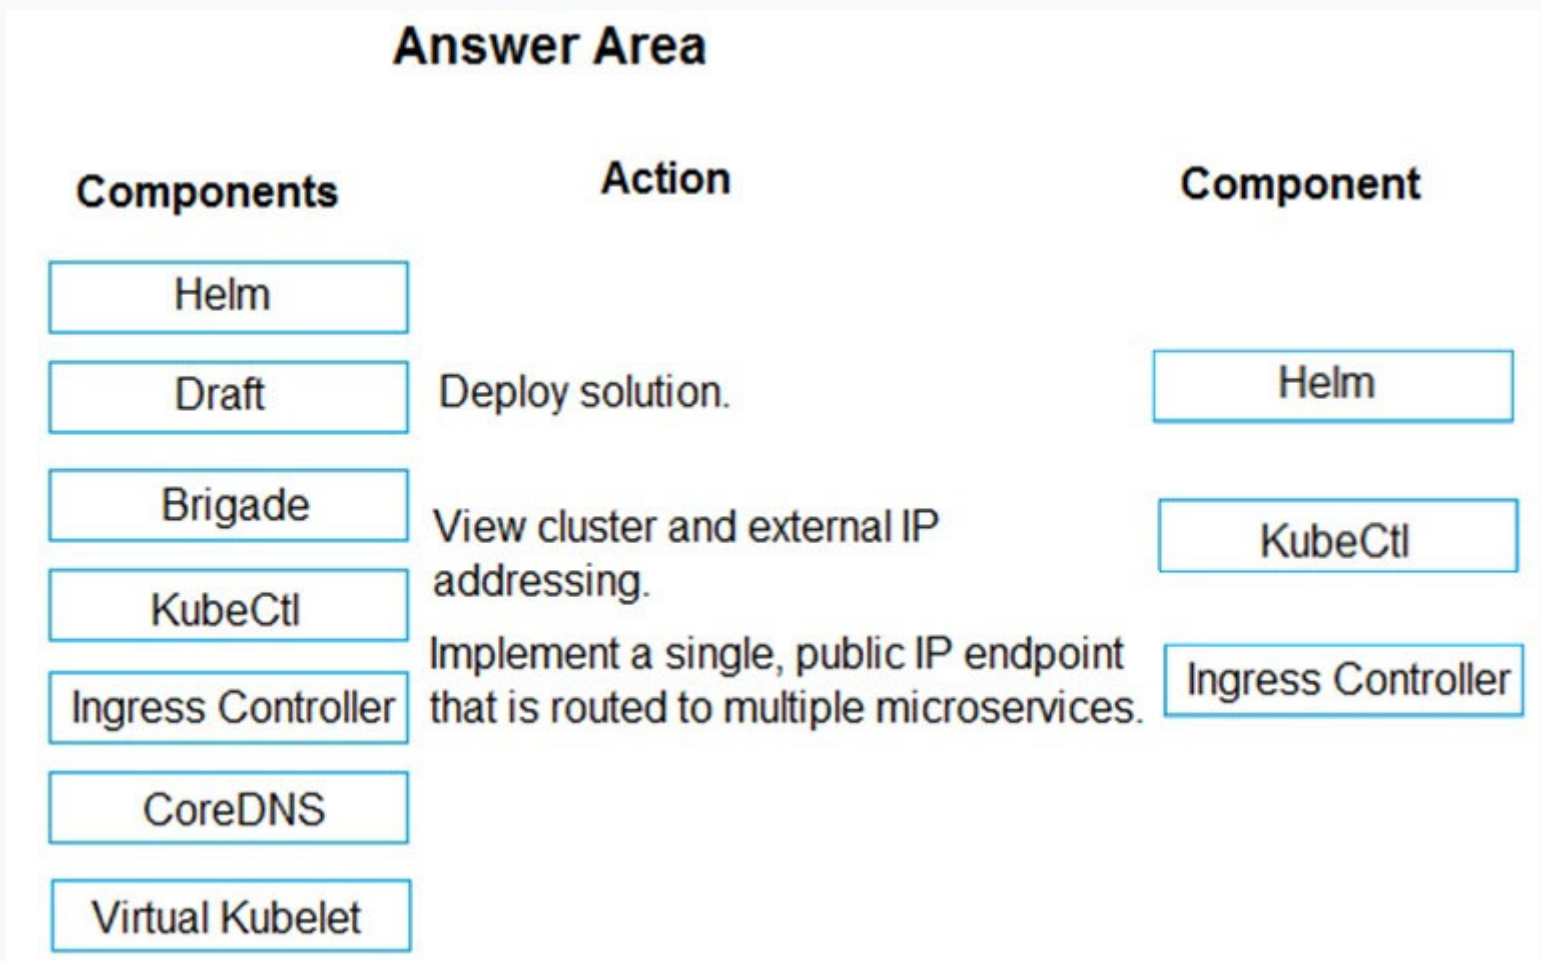

Box 1: Helm -

To create the ingress controller, use Helm to install nginx-ingress.

Box 2: kubectl -

To find the cluster IP address of a Kubernetes pod, use the kubectl get pod command on your local machine, with the option -o wide .

Box 3: Ingress Controller -
An ingress controller is a piece of software that provides reverse proxy, configurable traffic routing, and TLS termination for Kubernetes services. Kubernetes ingress resources are used to configure the ingress rules and routes for individual Kubernetes services.
Incorrect Answers:
Virtual Kubelet: Virtual Kubelet is an open-source Kubernetes kubelet implementation that masquerades as a kubelet. This allows Kubernetes nodes to be backed by Virtual Kubelet providers such as serverless cloud container platforms.
CoreDNS: CoreDNS is a flexible, extensible DNS server that can serve as the Kubernetes cluster DNS. Like Kubernetes, the CoreDNS project is hosted by the
CNCF.
Reference:
https://docs.microsoft.com/bs-cyrl-ba/azure/aks/ingress-basic https://www.digitalocean.com/community/tutorials/how-to-inspect-kubernetes-networking

mlantonis Highly Voted 2 years, 5 months ago

Box 1: Helm
Helm helps you manage Kubernetes applications — Helm Charts help you define, install, and upgrade even the most complex Kubernetes application. To create the ingress controller, use Helm to install nginx-ingress.

Box 2: Kubectl
The Kubernetes command-line tool, kubectl, allows you to run commands against Kubernetes clusters. To find the cluster IP address of a Kubernetes pod, use the kubectl get pod command on your local machine, with the option -o wide .

Box 3: Ingress Controller
An ingress controller is a piece of software that provides reverse proxy, configurable traffic routing, and TLS termination for Kubernetes services. Kubernetes ingress resources are used to configure the ingress rules and routes for individual Kubernetes services. Using an ingress controller and ingress rules, a single IP address can be used to route traffic to multiple services in a Kubernetes cluster.
upvoted 68 times

mlantonis 2 years, 5 months ago
Reference:

https://helm.sh

https://kubernetes.io/docs/tasks/tools

https://kubernetes.io/docs/concepts/services-networking/ingress-controllers

https://docs.microsoft.com/bs-cyrl-ba/azure/aks/ingress-basic

https://www.digitalocean.com/community/tutorials/how-to-inspect-kubernetes-networking
upvoted 9 times

robin1337 1 year, 7 months ago
Correct, I am a certified CKA und CKAD. The question isn't that well written though as you could also deploy a solution with kubectl...
upvoted 10 times

Leandromellor Highly Voted 2 years, 11 months ago
I believe there is no AKS question in the exam:
https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B
upvoted 15 times

matejka 2 years, 11 months ago
Haven't done the exam yet, but found many comments (eg. on Udemy tutorials) that Kuberentes actually IS part of the exam.
upvoted 6 times

luppittegui 2 years, 10 months ago
It's not: https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B
"Azure Kubernetes Service (AKS) is out of scope"
upvoted 4 times

daveonplanetearth 2 years, 8 months ago
I also thought AKS was excluded but people say they have had AKS questions in the exam. I don't understand why they would want to exclude it.
upvoted 1 times

Cornholioz 2 years, 10 months ago
And yet, here we are with this question. I'm guessing it has appeared in the exam and hence here. I am wondering what's actually "excluded" from Azure in this weird exam. Can't say it for many, but this one is a poorly constructed exam.
upvoted 7 times

cherylm 2 years, 8 months ago
I've also seen a lot of people mentioned in facebook groups they got kubernetes questions on the real exam az-204

upvoted 4 times

Rockm0uld 1 year, 7 months ago

I'm pretty sure I got this question on the exam, they be messin' with us

upvoted 1 times

wrak Most Recent 8 months, 1 week ago

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

upvoted 6 times

macobuzi 2 months, 1 week ago

Microsoft doesn't give a f about the scope. Better safe than sorry!

upvoted 2 times

nvtienanh 1 year, 1 month ago

Box 1: Kubectl or Helm

upvoted 2 times

azurelearner666 2 years, 4 months ago

It is correct.
Helm
KubeCtl
Ingress
simple as that.

upvoted 5 times

mlantonis 2 years, 5 months ago

Answer seems legit, but it's weird getting Kubernetes questions, when AKS is out of the scope of the exam.

upvoted 7 times

glam 2 years, 5 months ago

correct.

upvoted 1 times

Frakandel 2 years, 5 months ago

Answer:
1. KubeCtl or Helm
2. KubeCtl
3. Ingress Controller

upvoted 4 times

Annamarie0408 2 years, 6 months ago

Just took this test last week (last week of March) AKS -- WAS IN FACT -- In the test.

upvoted 7 times

i_ra 2 years, 6 months ago

Is these questions still valid? I'm taking the exam this week can I depend on these?? please answer me

upvoted 2 times

kwaazaar 2 years, 7 months ago

helm and kubectl can both be used to deploy.

upvoted 4 times

svaza 2 years, 8 months ago

Answer is right

upvoted 2 times

Question #16 *Topic 3*

DRAG DROP -

You are implementing an order processing system. A point of sale application publishes orders to topics in an Azure Service Bus queue. The Label property for the topic includes the following data:

Property	Description
ShipLocation	the country/region where the order will be shipped
CorrelationId	a priority value for the order
Quantity	a user-defined field that stores the quantity of items in an order
AuditedAt	a user-defined field that records the date an order is audited

The system has the following requirements for subscriptions:

Subscription type	Comments
FutureOrders	This subscription is reserved for future use and must not receive any orders
HighPriorityOrders	Handle all high priority orders and international orders
InternationalOrders	Handle orders where the country/region is not United States
HighQuantityOrders	Handle only orders with quantities greater than 100 units
AllOrders	This subscription is used for auditing purposes. This subscription must receive every single order. AllOrders has an Action defined that updates the AuditedAt property to include the date and time it was received by the subscription.

You need to implement filtering and maximize throughput while evaluating filters.

Which filter types should you implement? To answer, drag the appropriate filter types to the correct subscriptions. Each filter type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

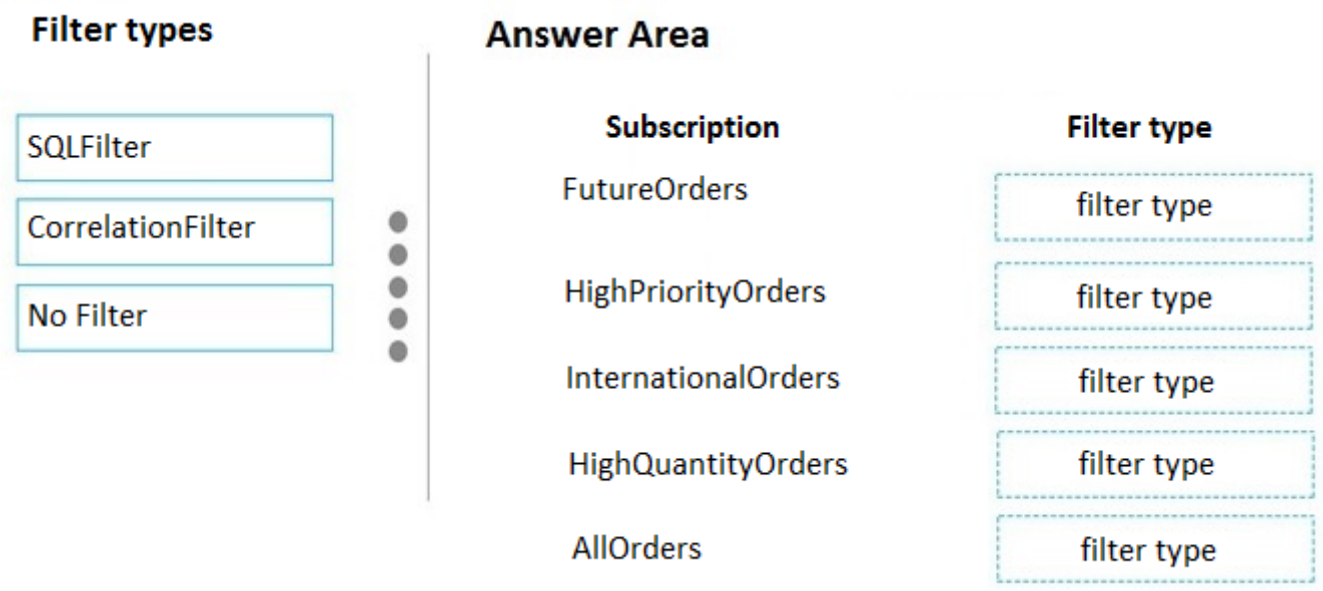

Correct Answer:

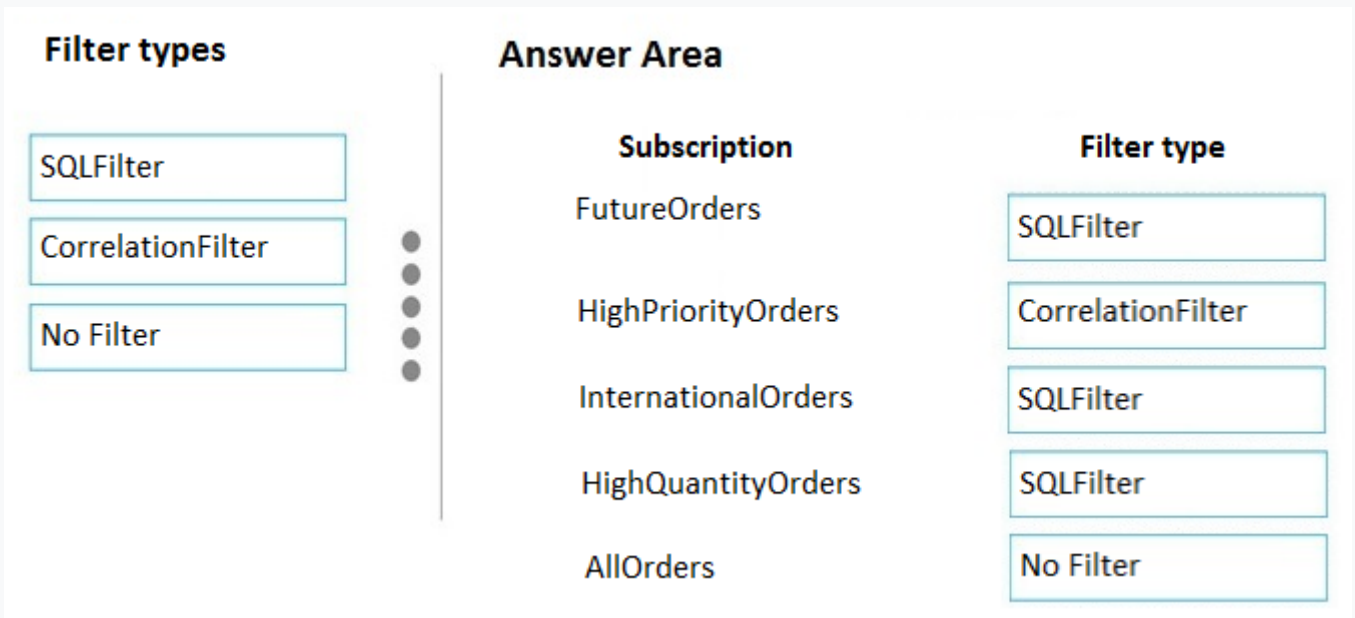

FutureOrders: SQLFilter -

HighPriortyOrders: CorrelationFilter

CorrelationID only -

InternationalOrders: SQLFilter -

Country NOT USA requires an SQL Filter

HighQuantityOrders: SQLFilter -

Need to use relational operators so an SQL Filter is needed.

AllOrders: No Filter -
SQL Filter: SQL Filters - A SqlFilter holds a SQL-like conditional expression that is evaluated in the broker against the arriving messages' user-defined properties and system properties. All system properties must be prefixed with sys. in the conditional expression. The SQL-language subset for filter conditions tests for the existence of properties (EXISTS), as well as for null-values (IS NULL), logical NOT/AND/OR, relational operators, simple numeric arithmetic, and simple text pattern matching with LIKE.
Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties. A common use is to match against the CorrelationId property, but the application can also choose to match against ContentType, Label, MessageId, ReplyTo,
ReplyToSessionId, SessionId, To, and any user-defined properties. A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter. For string expressions, the comparison is case-sensitive. When specifying multiple match properties, the filter combines them as a logical
AND condition, meaning for the filter to match, all conditions must match.
Boolean filters - The TrueFilter and FalseFilter either cause all arriving messages (true) or none of the arriving messages (false) to be selected for the subscription.
Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters

stylebc Highly Voted 2 years, 11 months ago

I think that it should be
-Correlation Filter (with the not existing value of any filed to avoid getting any message)
-SQL filter (as we need to get all high priority AND international orders, but for Correlation filter: A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter and we need not equal)
-SQL filter
-SQL filter
-No Filter

upvoted 129 times

MiraA 2 years, 1 month ago

Just note...
I think the FutureOrders and AllOrders should be accomplished with "TrueFilter" and "FalseFilter" but these options are not available in the assignment. May be in the future.

https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters

upvoted 4 times

mandynotmandy 1 year, 7 months ago

boolean filters are sql filters, true would be 1=1, false would be 1=0

upvoted 3 times

Necky 1 year, 1 month ago

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. So i guess it's
- Future Order > Sql filter
- High Prio > Co-relation
- International & High Quantity > Sql
- All > No filter

upvoted 12 times

[Removed] 3 months, 2 weeks ago

Incorrect

upvoted 1 times

rgullini 6 months, 3 weeks ago

Correlation just check for equality. if the values don´t match a specific criteria they will be ignored.

upvoted 1 times

Esward 9 months ago

Agreed

upvoted 1 times

rgullini 6 months, 3 weeks ago

Wrong:
- SQL Filter: You can exclude all orders by any of the other values but not by any of the lables (SQL Logic)
- Correlation ID: You can filter the high priority orders using the labels "CorrelationId" and "ShipLocation"
- SQL Filter
- SQL Filter
- No Filter: All orders will be retrieved

upvoted 3 times

rgullini 6 months, 3 weeks ago

Disregard this comment please. I was wrong.

upvoted 6 times

chettir01 9 months, 3 weeks ago

I don't agree with you, the second one must be correlationId :
example : // Create a correlation filter with color set to Red and priority set to High
await adminClient.CreateSubscriptionAsync(
new CreateSubscriptionOptions(topicName, "HighPriorityRedOrders"),
new CreateRuleOptions("HighPriorityRedOrdersRule", new CorrelationRuleFilter()
source : https://learn.microsoft.com/en-us/azure/service-bus-messaging/topic-filters

upvoted 1 times

rgullini 6 months, 3 weeks ago

How do you filter the countries with Correlation filters?

upvoted 2 times

TEMPKAKAM Highly Voted 2 years, 11 months ago

The Correct answers are:
No Filter
Correleation Filter
SQL filter
SQL filter
SQL filter

upvoted 40 times

cbn 2 years, 8 months ago

How do you justify Correlation Filter for HighPriorityOrders, when it involves Region <> US ?
I think it should be SQL filter as well.

upvoted 13 times

JustHereToLearn 10 months, 2 weeks ago

I somewhat agree with this answer after going through lot of resources. The only thing that bothers me is answer for the first and last question.
When you create a subscription by default a SQL filter is added which is 1=1 i.e. always evaluates to true.
So my source of confusion is how "No Filter" is treated. I have two interpretations either "No Filter" is we don't add any filter or the other is absence of a filter.
Mostly I believe it is treated as absence of a filter in which case the this answer looks correct.

upvoted 2 times

rashjan 2 years, 11 months ago

Why a filter for all orders?

upvoted 3 times

MrNair007 2 years, 11 months ago

See the basic is subscription does not have filter then does not receive any data. And since all orders needs all the orders it should have SQL filter with 1=1 so that all orders are passed to it.

upvoted 2 times

TakumaK 2 years, 5 months ago

if no filter explicitly specified, the true filter will be assigned which enables ALL messages. So your comment is wrong and makes confused.

upvoted 11 times

brtx 2 years, 11 months ago

https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters
"If you don't explicitly specify a filter condition for the rule, the applied filter is the true filter that enables all messages to be selected into the subscription."

upvoted 23 times

danielcr 2 years, 8 months ago

So.. The answer is correct.

upvoted 3 times

leviatas Most Recent 1 month ago

Got this quesiton in examn, went with answer. - 2023.09.25. Got Case Study Contoso

upvoted 2 times

longnguyendh 2 months, 4 weeks ago

It should be SQL, Correlation, SQL, SQL and NO Filter

upvoted 2 times

applepie 3 months ago

got this question today, go with answer -
Correlation
Correlation
SQL
SQL

No filter
7/30/2023, score 895/1000
upvoted 2 times

NightshadeRC 3 months, 1 week ago
Had this question today: 2023-07-26
I went:
- Correlation Filter
- SQL Filter
- SQL Filter
- SQL Filter
- No Filter
Every question on the exam was on ExamTopics. I entered every recommended answer and got 940
upvoted 6 times

DavidDev 7 months, 1 week ago
CorrelationFilter
SqlFilter
SqlFilter
SqlFilter
filterType
upvoted 1 times

Videira 8 months ago
On my exam 2023-02-25
upvoted 3 times

kandrew5 8 months, 1 week ago
Priotity should have a Correlation filter according to the example in Microsoft documentation: https://learn.microsoft.com/en-us/azure/service-bus-messaging/topic-filters
upvoted 1 times

narenazure 9 months ago
A SqlFilter holds a SQL-like conditional expression that is "evaluated" in the broker against the arriving messages' user-defined properties and system properties.
A CorrelationFilter holds a set of conditions that are "matched" against one or more of an arriving message's user and system properties.
So, for evaluation cases go for SQL Filter and "match exists" case go for Correlation filter
upvoted 2 times

dgcc97 10 months, 3 weeks ago
Based on the documentation:
https://learn.microsoft.com/en-us/azure/service-bus-messaging/topic-filters

My understanding is that Boolean Filters are a subset of SQL Filters -> as such the first is SQL Filter

Correlation Filter is used when you can check a property for a determined value (equality only - meaning that != is not included) and if you want to check more than one property, it is through a logical AND operation (both equalities must be present) -> as such, it can't be applied any.

SQL Filters are applied for all other validations types -> as such, is applied to the 2º,3º and 4º - and permit to use actions that annotate the message by adding, removing, or replacing properties and their values -> and as such the AllOrders is also a SQL Filter.
upvoted 3 times

BennyJai 10 months, 3 weeks ago
I got this in 5 dec test. Confused on no filter .
upvoted 1 times

mdg3501 10 months, 3 weeks ago
got this on 2022-12-7
upvoted 2 times

TheExamMaster2020 11 months, 2 weeks ago
Did my exam on 15th November 2022. This question was on it.
upvoted 3 times

coffecold 1 year ago
A correlation filter checks for equality and equality only. So priority>min priority is doubtful to realize.
I think it must be:
SQL (just false)
SQL (priority>min priority)
SQL (country != us)
SQL (size>100)
No filter
upvoted 6 times

gmishra88 1 year ago
First and second have multiple options. Impossible to answer with surety what the question maker wants as answers
- Correlation filter can be used with an undefined value with equality. SQL filter can also be used with an inequality
- Not clear how it determines something is international. International is when we know what is national and international maybe not national. For

that the country is required, I guess (unless Microsoft has an innovative programming logic using quantum computing). So to check the inequality you can use sqlfilter. Whatever can be done with correlation filter can be done with sqlfilter. I do not know why microsoft adds so many illogical features

upvoted 3 times

NK203 8 months ago

Complex filter rules require processing capacity. In particular, the use of SQL filter rules cause lower overall message throughput at the subscription, topic, and namespace level. Whenever possible, applications should choose correlation filters over SQL-like filters because they're much more efficient in processing and have less impact on throughput.
https://learn.microsoft.com/en-us/azure/service-bus-messaging/topic-filters

upvoted 1 times

P123123 11 months, 1 week ago

Completely agree

upvoted 1 times

Ishan_63 1 year, 1 month ago

For first one, sql filter is correct because as per the requirement the Future Order won't receive any message so the filter should be (1=0) and this comes under SQL filter category. By default all subscriptions receive all messages sent to Topic because the default filter condition is set as True(1=1) by default for all subscriptions. So for the last one No Filter is also correct

upvoted 2 times

Question #17 *Topic 3*

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:

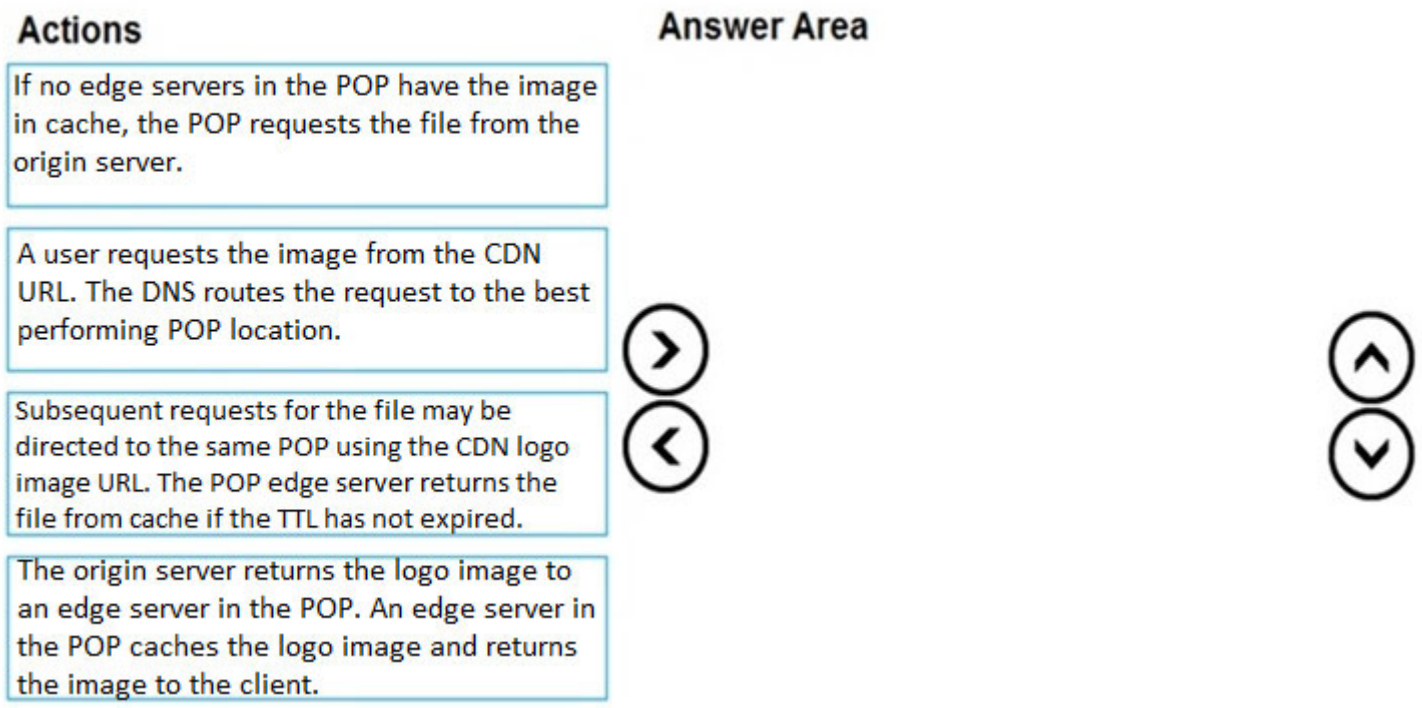

Correct Answer:

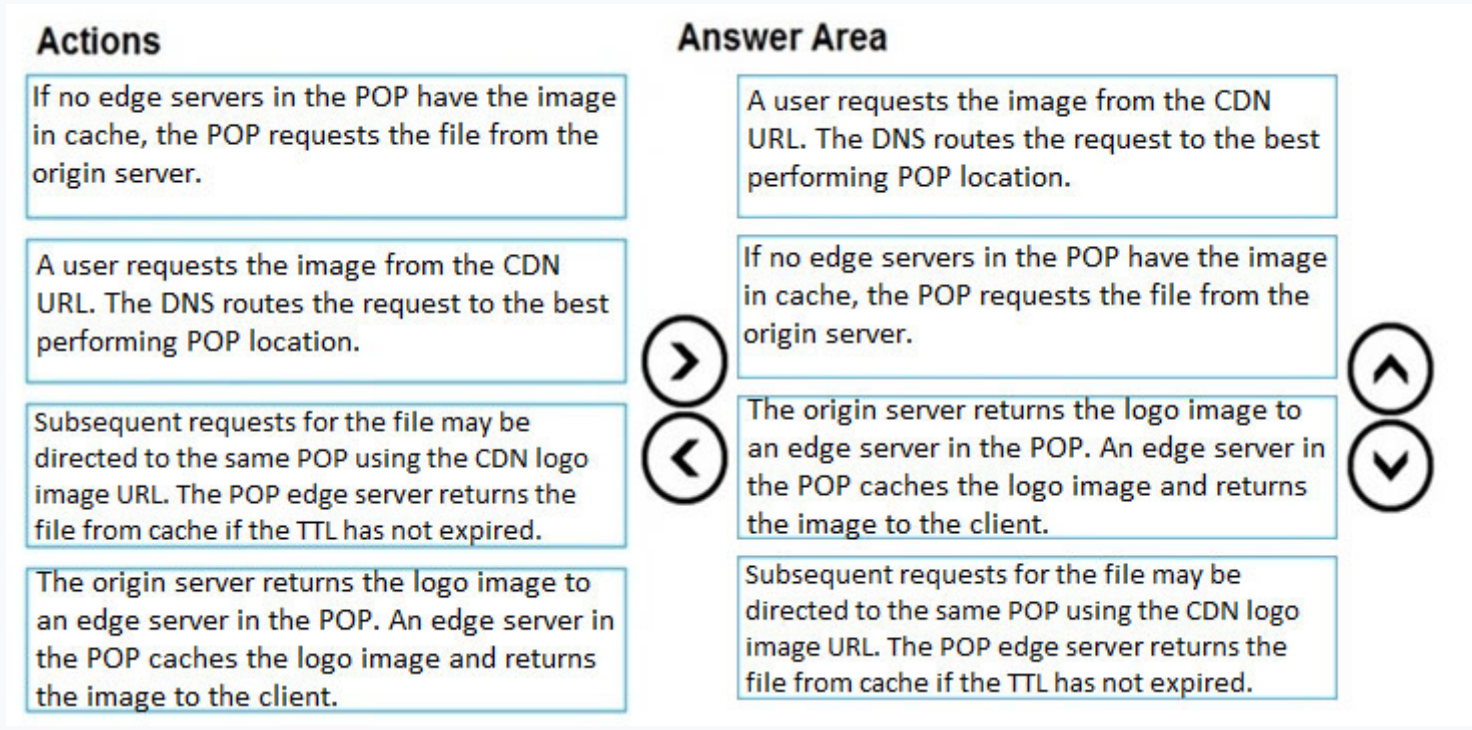

Step 1: A user requests the image..

A user requests a file (also called an asset) by using a URL with a special domain name, such as <endpoint name>.azureedge.net. This name can be an endpoint hostname or a custom domain. The DNS routes the request to the best performing POP location, which is usually the POP that is geographically closest to the user.

Step 2: If no edge servers in the POP have the..

If no edge servers in the POP have the file in their cache, the POP requests the file from the origin server. The origin server can be an Azure Web App, Azure

Cloud Service, Azure Storage account, or any publicly accessible web server.

Step 3: The origin server returns the..

The origin server returns the file to an edge server in the POP.

An edge server in the POP caches the file and returns the file to the original requestor (Alice). The file remains cached on the edge server in the POP until the time-to-live (TTL) specified by its HTTP headers expires. If the origin server didn't specify a TTL, the default TTL is seven days.

Step 4: Subsequent requests for..

Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.

If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.

Reference:

https://docs.microsoft.com/en-us/azure/cdn/cdn-overview

Ash111 Highly Voted 2 years, 11 months ago

Given ans is correct

upvoted 87 times

KhabibcandefeatGSP 2 years, 4 months ago

It feels so good when there are no debates on what is the correct answer.

upvoted 54 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000

upvoted 5 times

Tealon Highly Voted 2 years, 10 months ago

The given answer is correct.

upvoted 10 times

Tarajee Most Recent 1 month, 2 weeks ago

Got on my exam 2023sept

upvoted 2 times

warchoon 8 months, 1 week ago

"... returns the file directly from the cache. This process results in a faster, more responsive user experience" - last in order. Hmmm. Fine

upvoted 1 times

kyawtlearning 1 year, 2 months ago

Given answer is correct.

upvoted 2 times

Eltooth 1 year, 4 months ago

Answer is correct.

upvoted 2 times

gursimran_s 1 year, 6 months ago

Got it in April 2023

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

time travel!!

upvoted 9 times

meoukg 1 year, 7 months ago

Got it on 03/2022, chose the same as the given answer.

upvoted 2 times

silva_831 2 years, 4 months ago

The given answer is correct

upvoted 2 times

mlantonis 2 years, 5 months ago

The Answer is correct

Reference:

https://docs.microsoft.com/en-us/azure/cdn/cdn-overview

upvoted 5 times

warchoon 8 months, 1 week ago

"How it works" does not exactly represent the order.
"returns the file from cache if" should be before "if no ... image in cache"

upvoted 1 times

warchoon 8 months, 1 week ago

I know that "Subsequent" is the bad word for this choice and the author could be stupid enough to use wrong order.

upvoted 1 times

prabhjot 2 years, 5 months ago

correct ans

upvoted 1 times

UnknowMan 2 years, 5 months ago

The given answer is correct.

upvoted 2 times

glam 2 years, 5 months ago

correct.
upvoted 1 times

Fr3ddy 2 years, 9 months ago
The az-204 exam voucher costs 100 euros.
This value includes Study material + dump
upvoted 5 times

azurelearner666 2 years, 4 months ago
And? this comment is silly and does not help.
It's like saying: the sky is blue, except when it is clouded (or dark).
Admin: can you delete this useless message above? (and mine)
upvoted 11 times

coffecold 1 year ago
Admin, can you delete all "this answer is correct" minus one, and not delete the comment of Fr3ddy?
upvoted 2 times

mmdex 8 months, 2 weeks ago
Admin, can you delete this whole discussion? I don't see any comment that would have any value whatsoever.
upvoted 1 times

Saurabh_Kulkarni 2 years, 1 month ago
Lol...
upvoted 4 times

pac1311 2 years, 9 months ago
correct!
upvoted 2 times

sndp 2 years, 10 months ago
Answer is correct.
upvoted 5 times

Question #18 *Topic 3*

You are developing an Azure Cosmos DB solution by using the Azure Cosmos DB SQL API. The data includes millions of documents. Each document may contain hundreds of properties.
The properties of the documents do not contain distinct values for partitioning. Azure Cosmos DB must scale individual containers in the database to meet the performance needs of the application by spreading the workload evenly across all partitions over time.
You need to select a partition key.
Which two partition keys can you use? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. a single property value that does not appear frequently in the documents

B. a value containing the collection name

C. a single property value that appears frequently in the documents

D. a concatenation of multiple property values with a random suffix appended

E. a hash suffix appended to a property value

Correct Answer: *DE*
You can form a partition key by concatenating multiple property values into a single artificial partitionKey property. These keys are referred to as synthetic keys.
Another possible strategy to distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in this way, you can perform parallel write operations across partitions.
Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys

Community vote distribution
DE (100%)

TEMPKAKAM Highly Voted 2 years, 11 months ago
The given answer is correct
upvoted 73 times

mlantonis Highly Voted 2 years, 5 months ago
D and E

Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys
upvoted 17 times

[Removed] Most Recent 3 months, 2 weeks ago
Given answer is incorrect. Correct answer: AE
upvoted 1 times

Net_IT 1 week, 6 days ago
That is not the given answer? It is D and E.
upvoted 1 times

narenazure 9 months ago
Selected Answer: DE
It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key
upvoted 1 times

mdg3501 10 months, 3 weeks ago
got this on 2022-12-7
upvoted 5 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: DE

D. a concatenation of multiple property values with a random suffix appended Most Voted
E. a hash suffix appended to a property value Most Voted

upvoted 2 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 3 times

coffecold 1 year ago

at option E : a hash value will always deliver the same result on the same data. It is not a random value as stated in the proposed solution.
Since A,B and C fall off, and two answers must be chosen still I would go for D,E.

upvoted 2 times

xRiot007 1 year, 4 months ago

Selected Answer: DE

If no property in the document data will have unique values, you need to make one.
This is called a synthetic partition key.
These sorts of keys are made by adding a unique suffix at the end of some property.
One other way is to create a property that will have the hashed data + a random suffix.
The objective is to have a property that is random enough so that you can rely on it to be your key.

upvoted 5 times

Eltooth 1 year, 4 months ago

Selected Answer: DE

D & E are the correct answers.

upvoted 1 times

Rini100 1 year, 6 months ago

Got this on 20 Apr 2022

upvoted 2 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with highly voted answer.

upvoted 6 times

Mev4953 1 year, 9 months ago

Got this in the exam 01/22

upvoted 6 times

tramlong888 1 year, 10 months ago

the given answer is correct.

upvoted 1 times

sanjayrawat 2 years ago

seems, the given answer is correct.

upvoted 1 times

ning 2 years, 2 months ago

"The properties of the documents do not contain distinct values for partitioning." based on this statement, no single property in document can be used for partition, so only possibility is D / E, though in reality, I think you might be able to use something else

upvoted 4 times

francis6170 2 years, 4 months ago

got this in the exam :)

upvoted 5 times

carlos0808 2 years, 2 months ago

About how many questions that are present here, appeared in the exam?

upvoted 1 times

cool_tool 2 years, 3 months ago

Go Fun yourself

upvoted 2 times

Question #19 *Topic 3*

HOTSPOT -

You are developing an Azure-hosted e-commerce web application. The application will use Azure Cosmos DB to store sales orders. You are using the latest SDK to manage the sales orders in the database.

You create a new Azure Cosmos DB instance. You include a valid endpoint and valid authorization key to an appSettings.json file in the code project.

You are evaluating the following application code: (Line number are included for reference only.)

```
01 using System;
02 using System.Threading.Tasks;
03 using Microsoft.Azure.Cosmos;
04 using Microsoft.Extensions.Configuration;
05 using Newtonsoft.Json;
06 namespace SalesOrders
07 {
08   public class SalesOrder
09   {
10     . . .
11   }
12   internal class ManageSalesOrders
13     {
14       private static async Task GenerateSalesOrders()
15         {
16           IConfigurationRoot configuration = new ConfigurationBuilder().AddJsonFile("appSettings.json").Build();
17           string endpoint = configuration["EndPointUrl"];
18           string authKey = configuration["AuthorizationKey"];
19           using CosmosClient client = new CosmosClient(endpoint, authKey);
20           Database database = null;
21           using (await client.GetDatabase("SalesOrders").DeleteStreamAsync()) { }
22           database = await client.CreateDatabaseIfNotExistsAsync("SalesOrders");
23           Container container1 = await database.CreateContainerAsync(id: "Container1", partitionKeyPath: "/AccountNumber");
24           Container container2 = await database.CreateContainerAsync(id: "Container2", partitionKeyPath: "/AccountNumber");
25           SalesOrder salesOrder1 = new SalesOrder() { AccountNumber = "123456" };
26           await container1.CreateItemAsync(salesOrder1, new PartitionKey(salesOrder1.AccountNumber));
27           SalesOrder salesOrder2 = new SalesOrder() { AccountNumber = "654321" };
28           await container1.CreateItemAsync(salesOrder2, new PartitionKey(salesOrder2.AccountNumber));
29           SalesOrder salesOrder3 = new SalesOrder() { AccountNumber = "109876" };
30           await container2.CreateItemAsync(salesOrder3, new PartitionKey(salesOrder3.AccountNumber));
31           _ = await database.CreateUserAsync("User1");
32           User user1 = database.GetUser("User1");
33           _ = await user1.ReadAsync();
34     }
35   }
36 }
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
A database named SalesOrders is created. The database will include two containers.	○	○
Container1 will contain two items.	○	○
Container2 will contain one item.	○	○

Correct Answer:

Answer Area

Statements	Yes	No
A database named SalesOrders is created. The database will include two containers.	○	○
Container1 will contain two items.	○	○
Container2 will contain one item.	○	○

Box 1: Yes -

The createDatabaseIfNotExistsAsync method checks if a database exists, and if it doesn't, create it.

The Database.CreateContainerAsync method creates a container as an asynchronous operation in the Azure Cosmos service.

Box 2: Yes -

The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service.

Box 3: Yes -

Reference:

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient.createdatabaseifnotexistsasync

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.database.createcontainerasync https://docs.microsoft.com/en-us/dotnet/api/azure.cosmos.cosmoscontainer.createitemasync

mlantonis Highly Voted 2 years, 5 months ago

Box 1: Yes
The createDatabaseIfNotExistsAsync method checks if a database exists, and if it doesn't, create it. (Line 22)
The Database.CreateContainerAsync method creates a container as an asynchronous operation in the Azure Cosmos service. (Line 23 and 24)

Box 2: Yes
The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service. (Line 26 and 28)

Box 3: Yes
The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service. (Line 30)

upvoted 82 times

Esward 9 months ago

Agreed

upvoted 1 times

chettir01 10 months, 1 week ago

The first box is not correct, the name of database is database and not salesOrder !

upvoted 1 times

Vladimir_Gajinov 1 month, 2 weeks ago

So, this answer is correct one.

upvoted 1 times

Vladimir_Gajinov 1 month, 2 weeks ago

You are right but in Azure Cosmos DB, the database name is typically the same as the database ID, but spaces in the ID are replaced with hyphens ("-").

upvoted 1 times

chettir01 10 months, 1 week ago

SalesOrders is the databaseID, not the name !

upvoted 1 times

chettir01 10 months, 1 week ago

```
// New instance of Database class referencing the server-side database
// The name of instance is database2, and we need an ID to create it !
Database database2 = await client.CreateDatabaseIfNotExistsAsync(
id: "adventureworks-2"
);
```

upvoted 1 times

mlantonis 2 years, 5 months ago

Reference:

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient.createdatabaseifnotexistsasync

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.database.createcontainerasync

https://docs.microsoft.com/en-us/dotnet/api/azure.cosmos.cosmoscontainer.createitemasync

upvoted 7 times

AOE Highly Voted 2 years, 2 months ago

Line 21 is tricky, it assumes the database is already created as it calls DeleteStreamAsync. I'm confused.

upvoted 7 times

warchoon 8 months, 1 week ago

It returns ResponseMessage which is IDisposable

upvoted 1 times

coffecold 1 year ago

"using" is equal to a "try-finaly" with noting in finaly so the null exeption is already been taken care of .
https://stackoverflow.com/questions/2522822/will-dispose-be-called-in-a-using-statement-with-a-null-object

upvoted 2 times

Knightie 1 year, 2 months ago

// Delete a Database resource where database_id is the ID property of the Database resource you wish to delete.
Database database = this.cosmosClient.GetDatabase(database_id);
await database.DeleteStreamAsync();

That is deleted so the create if not exist will surely create the DB

upvoted 1 times

ReyPirata Most Recent 2 months, 1 week ago

On my exam 2023-08-20. Scored 925
Yes
Yes
Yes

upvoted 2 times

ks321 3 months, 1 week ago

Box 1: Yes
Box 2: Yes
Box 3: No - container2.Create is called only once

upvoted 1 times

macobuzi 2 months, 1 week ago

I think you are confused! It said container2 has 1 item. So box3 - Yes is correct

upvoted 2 times

Videira 8 months ago

On my exam 2023-02-25

upvoted 1 times

Lucifer14 10 months ago

In 28-12-2022 exam

upvoted 3 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 2 times

sam5678 1 year ago

yes yes yes

upvoted 1 times

chettir01 10 months, 1 week ago

No, it's not correct, The first box is not correct, the name of database on the code is database and not salesOrder !

upvoted 1 times

chettir01 10 months, 1 week ago

SalesOrders is the databaseID, not the name !

upvoted 1 times

Eltooth 1 year, 4 months ago

Answer is correct:
Yes
Yes
Yes

upvoted 1 times

chettir01 10 months, 1 week ago

No, it's not correct, The first box is not correct, the name of database on the code is database and not salesOrders !
SalesOrders is the ID of the database not it's name !

upvoted 1 times

AymanAkk 10 months, 1 week ago

there is no name of database, that s a variable called database of type Databse and it will receive the newly created database with name salesOrder !

upvoted 4 times

nargzul 2 years, 4 months ago

I'm just curious, how do you know that Container 1 contains 2 items and container2 contains 1 item and not the opposite? I see that we are partitionning on the account number, but not sure to understand how the partition is made?

upvoted 2 times

Knightie 1 year, 2 months ago

container 1 will have 2 partitions, container 2 only 1 partition because partition key is the account number. But container 1 will still have 2 items, 1 item per partition. Hope this help.

upvoted 1 times

BroGood 2 years, 4 months ago

see Lines 26 & 28 contain items for Container 1
Only Line 30 containers an item for Container 2

upvoted 5 times

ariel_dev 2 years, 4 months ago

container1 is called two times adding the items, container2 just one.

upvoted 6 times

AlokSingh 2 years, 4 months ago

Correct Answer

upvoted 1 times

Question #20 *Topic 3*

DRAG DROP -

You develop an Azure solution that uses Cosmos DB.

The current Cosmos DB container must be replicated and must use a partition key that is optimized for queries.

You need to implement a change feed processor solution.

Which change feed processor components should you use? To answer, drag the appropriate components to the correct requirements. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view the content.

NOTE: Each correct selection is worth one point.

Select and Place:

Components

- Host
- Delegate
- Lease container
- Monitored container

Answer Area

Requirement	Component
Store the data from which the change feed is generated.	Component
Coordinate processing of the change feed across multiple workers.	Component
Use the change feed processor to listen for changes.	Component
Handle each batch of changes.	Component

Correct Answer:

Components

Answer Area

Requirement	Component
Store the data from which the change feed is generated.	Monitored container
Coordinate processing of the change feed across multiple workers.	Lease container
Use the change feed processor to listen for changes.	Host
Handle each batch of changes.	Delegate

Box 1: The monitored container -

The monitored container has the data from which the change feed is generated. Any inserts and updates to the monitored container are reflected in the change feed of the container.

Box 2: The lease container -

The lease container acts as a state storage and coordinates processing the change feed across multiple workers. The lease container can be stored in the same account as the monitored container or in a separate account.

Box 3: The host: A host is an application instance that uses the change feed processor to listen for changes. Multiple instances with the same lease configuration can run in parallel, but each instance should have a different instance name.

Box 4: The delegate -

The delegate is the code that defines what you, the developer, want to do with each batch of changes that the change feed processor reads.

Reference:

https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-processor

MattXu Highly Voted 2 years, 4 months ago

The given answer is correct.

upvoted 53 times

TakumaK 2 years, 3 months ago

You are right. The given answer is correct.

upvoted 9 times

john4p 1 year, 10 months ago

And your given answer, that MattXu is right, is also correct.

upvoted 16 times

dbobspurfpoo 1 year, 10 months ago

that observation is correct, john

upvoted 14 times

john4p 1 year, 10 months ago

Thank you for the confirmation.

upvoted 12 times

riyachavan 1 week, 6 days ago

You are welcome.

upvoted 1 times

gmishra88 Highly Voted 1 year ago

Microsoft has obscrure names in documentation and has nothing to do with the azure component itself. "Delegate", "Host component", "Compute instance". How's remembering this from one page they wrote mean anything. We do not call any of them these when we implement the change feed processor

upvoted 9 times

Dixavado Most Recent 1 month ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 2 times

shekhar11 1 month, 1 week ago

Question in my exam 22sept 2023

upvoted 1 times

nikipediaa 1 month ago

me too today

upvoted 1 times

Tarajee 1 month, 2 weeks ago

On my exam 2023sept

upvoted 2 times

kayvg 2 months, 2 weeks ago

1. Monitored container
> You want to track changes of the data you store

2. Lease container
> A lease container is responsible for maintaining information about which workers are processing which data from the Monitored Container

3. Host
> They represent the worker instances responsible for processing changes from the change feed

4. Delegates
> Methods you define to handle actual changes detected by the change feed (so when the change feed receives changes, it invokes delegate methods)

The difference between a monitored container and a host is that a host actually distributes work to delegates, while a monitored container tracks changes in data and gives that info to the change feed processor.

upvoted 2 times

KrishTeam 3 months ago

please provide correct answers

upvoted 1 times

unraval 3 months, 1 week ago

The discussions are meant to help prepare for the exam so please try to be helpful

upvoted 2 times

unraval 3 months, 1 week ago

People are commenting here to be funny, but can they care to provide an explanation?

upvoted 1 times

Chris2349 4 months, 2 weeks ago

Received this on 15th of June 2023. Went with the given answer.

upvoted 2 times

Videira 8 months ago

On my exam 2023-02-25

upvoted 3 times

mdg3501 10 months, 3 weeks ago

got this on 2022-12-7

upvoted 3 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 5 times

BennyJai 10 months, 3 weeks ago

I also Got same in 5 dec exam.

upvoted 2 times

sam5678 1 year ago

dzzdvz

upvoted 1 times

herrmutig 1 year, 6 months ago

Answer is correct.

*Note that the "Host" Component should be called Compute Instance instead.

https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-processor

upvoted 4 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 2 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with the given answer

upvoted 5 times

Question #21 *Topic 3*

HOTSPOT -

You are developing a web application that will use Azure Storage. Older data will be less frequently used than more recent data.

You need to configure data storage for the application. You have the following requirements:

- Retain copies of data for five years.
- Minimize costs associated with storing data that is over one year old.
- Implement Zone Redundant Storage for application data.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Requirement	Solution
Configure an Azure Storage account	Implement Blob Storage Implement Azure Cosmos DB Implement Storage (general purpose v1) Implement StorageV2 (general purpose v2)
Configure data retention	Snapshot blobs and move them to the archive tier Set a lifecycle management policy to move blobs to the cool tier Use AzCopy to copy the data to an on-premises device for backup Set a lifecycle management policy to move blobs to the archive tier

Correct Answer:

Answer Area

Requirement	Solution
Configure an Azure Storage account	Implement Blob Storage Implement Azure Cosmos DB Implement Storage (general purpose v1) Implement StorageV2 (general purpose v2)
Configure data retention	Snapshot blobs and move them to the archive tier Set a lifecycle management policy to move blobs to the cool tier Use AzCopy to copy the data to an on-premises device for backup Set a lifecycle management policy to move blobs to the archive tier

Reference:

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy?toc=/azure/storage/blobs/toc.json

gmishra88 Highly Voted 1 year ago

So, because it is ZRS, and that does not support arhive tier, it cannot be moved to archive tier even though the questions mention the red-herring key-word "infrequently accessed" (which triggers feelings for archive tier). For no logically apparent reason Microsoft decided not to support archive tier in ZRS and unfortunately I have to remember that Microsoft "feature"?

upvoted 27 times

OPT_001122 11 months, 2 weeks ago

ZRS, and that does not support arhive tier - this is key point, Thanks

upvoted 10 times

jakobste 2 months, 2 weeks ago

I assume they wanted to charge people more money for ZRS and thus no Archive tier.

upvoted 1 times

surprise0011 6 months, 2 weeks ago

also GZRS and RA-GZRS does not support archieve tier
https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview#archive-access-tier

upvoted 2 times

18Marks Highly Voted 10 months, 3 weeks ago

Got in exam. go with given answer

upvoted 8 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000

upvoted 2 times

adilkhan 7 months ago

were all the questions from the exam topic?

upvoted 3 times

Dixavado Most Recent 1 month ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 1 times

applepie 3 months ago

got this question today, go with the provided answer - 7/30/2023, score 895/1000

upvoted 1 times

Priya0703 8 months, 1 week ago

On exam 20-02-2023

upvoted 3 times

Esward 9 months, 2 weeks ago

1. Implement storage V2
2. Set lifecycle management policy to move to cool tier

which means given answers are correct

upvoted 2 times

Hasti123 1 month, 1 week ago

Why not Blob storage ?

upvoted 1 times

Alluru 1 year ago

Given answer is correct. Data retention policy relay on cool tier

upvoted 4 times

gmishra88 1 year ago

Microsoft plays on the unfortunate choice of words in their documentation: "Rarely used", "Infrequently used". The difference I cannot find

upvoted 4 times

OPT_001122 1 year, 1 month ago

Given answer is correct or not?

upvoted 1 times

micro9000 11 months, 3 weeks ago

Yup, the answer is correct
Based on this: https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview?source=recommendations
Only storage accounts that are configured for LRS, GRS, or RA-GRS support moving blobs to the archive tier. The archive tier isn't supported for ZRS, GZRS, or RA-GZRS accounts. For more information about redundancy configurations for Azure Storage, see Azure Storage redundancy.

upvoted 4 times

Enigma___ 1 year, 1 month ago

Only storage accounts that are configured for LRS, GRS, or RA-GRS support moving blobs to the Archive tier. The Archive tier isn't supported for ZRS, GZRS, or RA-GZRS accounts.
https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview?source=recommendations

upvoted 3 times

finnishr 1 year, 1 month ago

Correct. Can't pick the archive option since ZRS needs to be used.

upvoted 2 times

RochaG2 1 year, 1 month ago

Data in all tiers, including the Archive tier, is always copied from the primary to the secondary during geo-replication. The Archive tier for Blob Storage is currently supported for LRS, GRS, and RA-GRS accounts, but not for ZRS, GZRS, or RA-GZRS accounts.

https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy

upvoted 2 times

sghaha 1 year, 6 months ago

https://docs.microsoft.com/ko-kr/azure/storage/common/storage-account-overview

upvoted 1 times

CellCS 10 months ago

here is en-us link : https://learn.microsoft.com/en-us/azure/storage/common/storage-account-overview

upvoted 1 times

Question #22 *Topic 3*

HOTSPOT -

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

You have the following requirements:

- Code must be scalable and allow for growth.
- Each record must consist of a playerId, gameId, score, and time played.
- When users reach a new high score, the system will save the new score using the SaveScore function below.

Each game is assigned an Id based on the series title.

▪

You plan to store customer information in Azure Cosmos DB. The following data already exists in the database:

PartitionKey	RowKey	Email
Harp	Walter	wharp@contoso.com
Smith	Steve	ssmith@contoso.com
Smith	Jeff	jsmith@contoso.com

You develop the following code to save scores in the data store. (Line numbers are included for reference only.)

```
01 public void SaveScore(string gameId, string playerId, int score, long timePlayed)
02 {
03 CloudStorageAccount storageAccount = CloudStorageAccount.Parse(connectionString);
04 CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
05 CloudTable table = tableClient.GetTableReference("scoreTable");
06 table.CreateIfNotExists();
07 var scoreRecord = new PlayerScore(gameId, playerId, score, timePlayed);
08 TableOperation insertOperation = TableOperation.Insert(scoreRecord);
09 table.Execute(insertOperation);
10 }
```

You develop the following code to query the database. (Line numbers are included for reference only.)

```
01 CloudTableClient tableClient = account.CreateCloudTableClient();
02 CloudTable table = tableClient.GetTableReference ("people");
03 TableQuery < CustomerEntity> query = new TableQuery <CustomerEntity >( )
04 .Where( TableQuery.CombineFilters (
05 TableQuery.GenerateFilterCondition (PartitionKey, QueryComparisons.Equal , "Smith"),
06 TableOperators.And , TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal, "ssmith@contoso.com")
07 ));
08 await table.ExecuteQuerySegmentedAsync< CustomerEntity>(query, null);
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
SaveScore will work with Cosmos DB.	◯	◯
SaveScore will update and replace a record if one already exists with the same playerId and gameId.	◯	◯
Leader board data for the game will be automatically partitioned using gameId.	◯	◯
SaveScore will store the values for the gameId and playerId parameters in the database.	◯	◯

Correct Answer:

Answer Area

Statements	Yes	No
SaveScore will work with Cosmos DB.	◯	◯
SaveScore will update and replace a record if one already exists with the same playerId and gameId.	◯	◯
Leader board data for the game will be automatically partitioned using gameId.	◯	◯
SaveScore will store the values for the gameId and playerId parameters in the database.	◯	◯

Box 1: Yes -

Create a table.

A CloudTableClient object lets you get reference objects for tables and entities. The following code creates a CloudTableClient object and uses it to create a new

CloudTable object, which represents a table

```
// Retrieve storage account from connection-string.
CloudStorageAccount storageAccount =
CloudStorageAccount.parse(storageConnectionString);
// Create the table client.
CloudTableClient tableClient = storageAccount.createCloudTableClient();
// Create the table if it doesn't exist.
String tableName = "people";
CloudTable cloudTable = tableClient.getTableReference(tableName); cloudTable.createIfNotExists();
```

Box 2: No -

New records are inserted with TableOperation.insert. Old records are not updated.

To update old records TableOperation.insertOrReplace should be used instead.

Box 3: No -

Box 4: Yes -

Reference:

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

yukkki Highly Voted 9 months, 3 weeks ago

given answer is correct.

upvoted 8 times

[Removed] 3 months, 2 weeks ago

Incorrect. First one is "No"

upvoted 2 times

Esward 9 months, 2 weeks ago

yes, you are correct! given answers are correct

upvoted 2 times

OPT_001122 Highly Voted 12 months ago

the given answer seems correct

upvoted 6 times

Azr0112 Most Recent 1 month ago

Yes, SaveScore will work with CosmosDb.
This is an example of Azure Cosmos DB for Table.
https://learn.microsoft.com/en-us/azure/cosmos-db/table/quickstart-dotnet?tabs=azure-cli%2Cwindows

upvoted 1 times

[Removed] 3 months, 2 weeks ago

No, No, No, Yes

upvoted 1 times

[Removed] 9 months, 3 weeks ago

Inserting will fail, this would require InsertOrReplace to work

upvoted 3 times

bbq598 10 months, 1 week ago

For box 2 there will be an exception if the same partition key and row key already exist.

upvoted 2 times

warchoon 8 months, 1 week ago

So you need to answer box 3 first ;)

upvoted 1 times

Yumi21 11 months, 1 week ago

It seems that in the code no partition key for scoreTable is specified, which means that here we are using a single-partition collection. I would say in this case 3rd option is No, so no automatical partitioning will happen.

upvoted 2 times

TheExamMaster2020 1 year ago

Y,N,Y,Y

upvoted 1 times

coffecold 1 year ago

What code should be there if doing automatic partitioning (C) remains a total secret to me. Should be somewhere on the table level...

upvoted 3 times

DivyaRajkumar 1 year, 1 month ago

The given answer looks right to me,
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.table?view=azure-dotnet

upvoted 2 times

azurepaul 1 year, 1 month ago

CloudTableClient is for Table storage - CosmosClient is for Cosmos DB I don't think it will work with Cosmos

upvoted 4 times

PrepX 4 months, 1 week ago

A = No
SaveScore() method uses a storage account connectionString:
CloudStorageAccount storageAccount = CloudStorageAccount.Parse(connectionString);
and uses storageAccount to create CloudTableClient.

This is not a CosmosDB connection string, CosmosDB requires a Uri + credential (not given in given code!) So given code wil not work 100% to connect to CosmosDB even though it uses a compatible API.

upvoted 1 times

ArturKon 1 year, 1 month ago

https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.table.cloudtableclient?view=azure-dotnet

upvoted 1 times

ArturKon 1 year, 1 month ago

Cosmos DB offers Table API, you can use CosmosClient to connect to this.

upvoted 2 times

Enigma___ 1 year, 1 month ago

Please do a simple bing search and you'll see it is Cosmos Db
https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.table.cloudtableclient?view=azure-dotnet

upvoted 1 times

finnishr 1 year, 1 month ago

Correct 100%!

upvoted 2 times

BogdanG 1 year, 6 months ago

Answer is Correct.

upvoted 3 times

Question #23 *Topic 3*

You develop and deploy a web application to Azure App Service. The application accesses data stored in an Azure Storage account. The account contains several containers with several blobs with large amounts of data. You deploy all Azure resources to a single region.
You need to move the Azure Storage account to the new region. You must copy all data to the new region.
What should you do first?

A. Export the Azure Storage account Azure Resource Manager template

B. Initiate a storage account failover

C. Configure object replication for all blobs

D. Use the AzCopy command line tool

E. Create a new Azure Storage account in the current region

F. Create a new subscription in the current region

Correct Answer: *A*
To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice and finally, delete the resources in the source region.
To get started, export, and then modify a Resource Manager template.
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal

Community vote distribution
A (100%)

BogdanG Highly Voted 1 year, 6 months ago
Answer is correct.
upvoted 10 times

BhavikaSNN Highly Voted 1 year, 6 months ago
Selected Answer: A
We can create a new Storage account in the new region, using the existing storage account ARM template. All we need to do is change the region name after exporting the ARM of existing account..
upvoted 8 times

gmishra88 1 year ago
And the name of the storage account which needs to be unique. I think all these questions are about some sentence in some azure documentation.
upvoted 4 times

Nhiendo Most Recent 10 months ago
Step One is always "Export"
upvoted 6 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: A
A. Export the Azure Storage account Azure Resource Manager template
upvoted 3 times

TheExamMaster2020 11 months, 2 weeks ago
Did my exam on 15th November 2022. This question was on it.
upvoted 5 times

Question #24 *Topic 3*

HOTSPOT -

You are developing an application to collect the following telemetry data for delivery drivers: first name, last name, package count, item id, and current location coordinates. The app will store the data in Azure Cosmos DB.

You need to configure Azure Cosmos DB to query the data.

Which values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

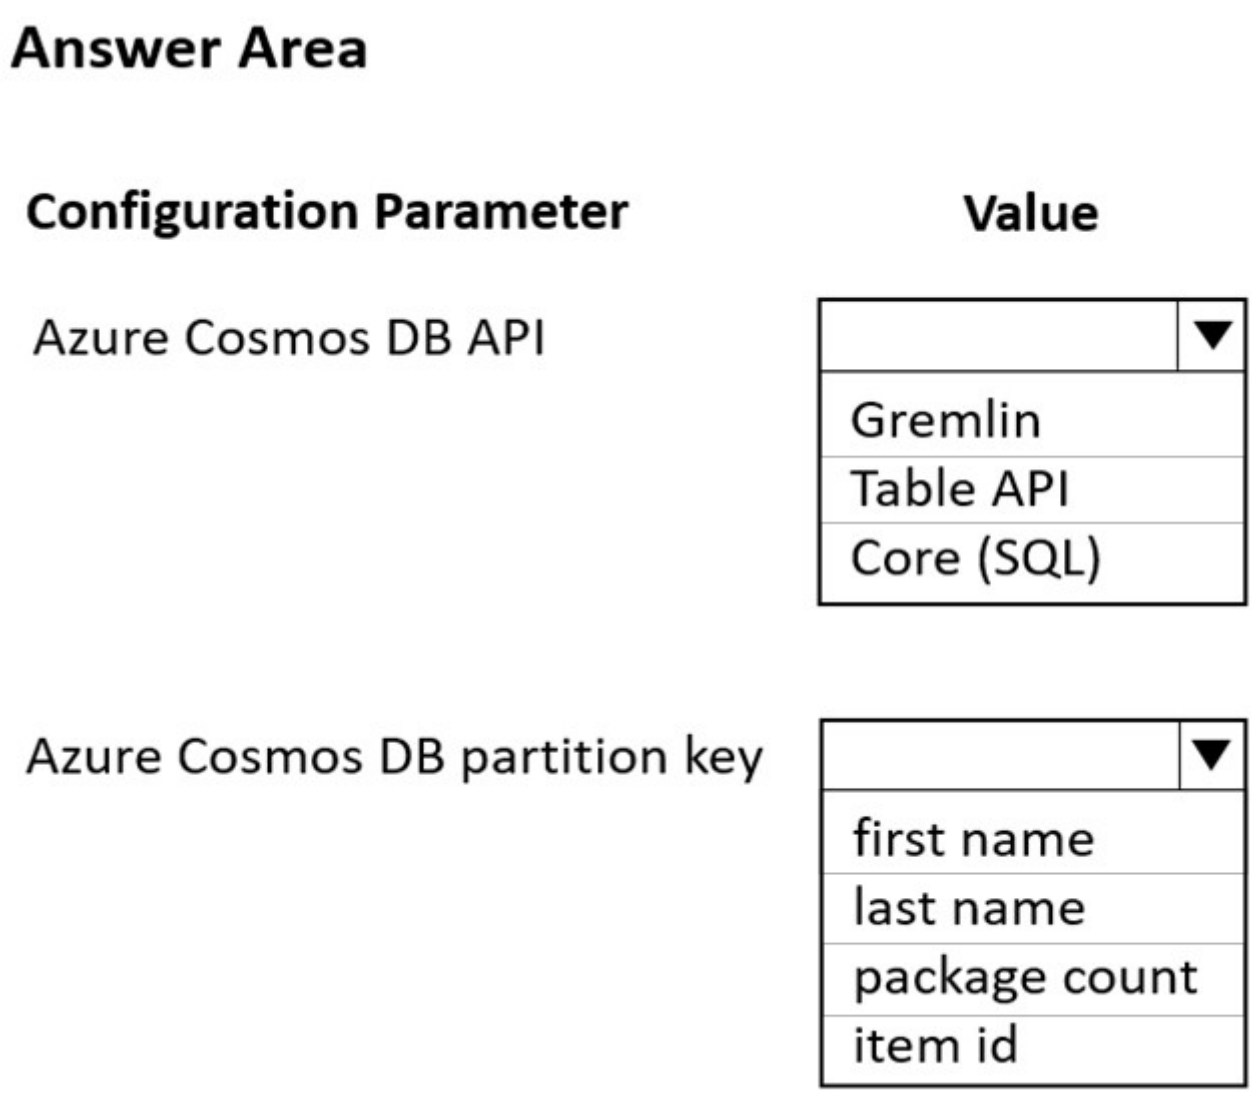

Correct Answer:

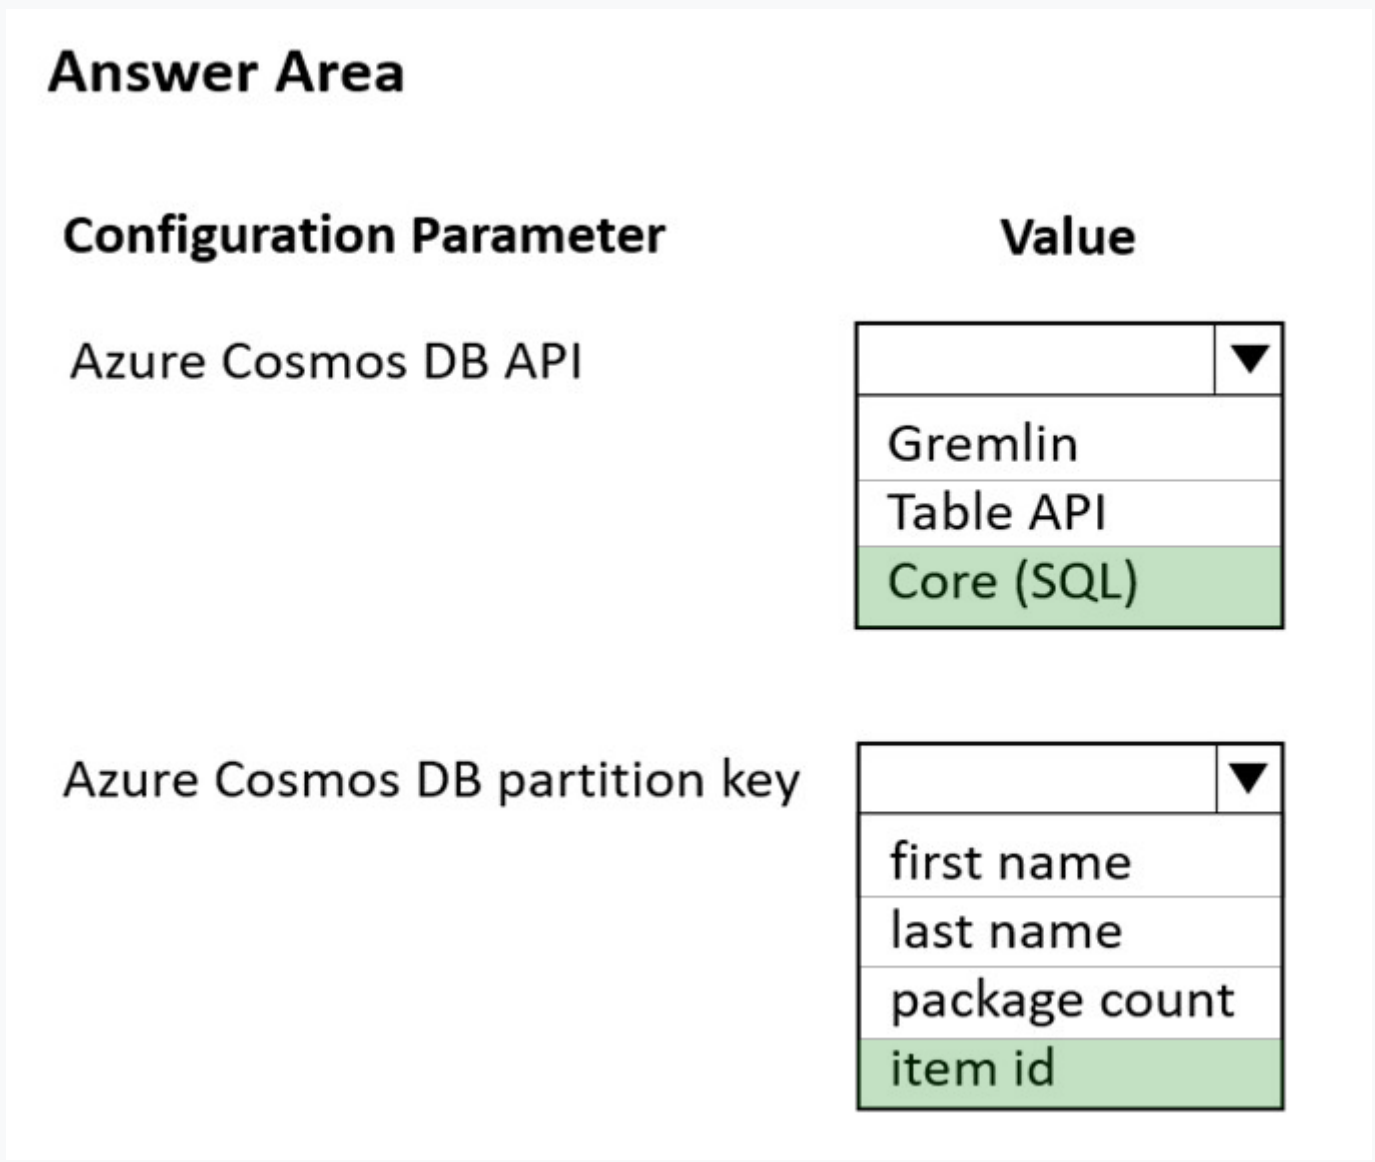

Box 1: Core (SQL)

Core(SQL) API stores data in document format. It offers the best end-to-end experience as we have full control over the interface, service, and

the SDK client libraries. SQL API supports analytics and offers performance isolation between operational and analytical workloads.

Box 2: item id -
item id is a unique identifier and is suitable for the partition key.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api
https://docs.microsoft.com/en-us/azure/cosmos-db/partitioning-overview

[Removed] Highly Voted 1 year, 1 month ago

Without knowing the functionality or the usage pattern or what it is for. Good lord, Microsoft

upvoted 24 times

serpevi Highly Voted 1 year, 1 month ago

Got this in 09/22 , went with SQL and Item Id, score 927.

upvoted 19 times

Tarajee Most Recent 1 month, 2 weeks ago

On my exam 2023sept

upvoted 5 times

NightshadeRC 3 months, 1 week ago

Had this question in today's exam: 2023-07-26

upvoted 5 times

[Removed] 3 months, 2 weeks ago

Answer seems correct, but the question is very bad. It doesn't even tell about the usage, so it could be Table API as well

upvoted 6 times

sarmaria 7 months, 2 weeks ago

Got this question in the exam on 16/03/2023. Went with SQL and Item Id . Make sure to prepare for case studies. I got city and lights case study.

upvoted 5 times

MarcoStewart 4 months, 3 weeks ago

where can one find these case studies?

upvoted 3 times

warchoon 8 months, 1 week ago

They changed names again. "Core (SQL)" is "Api for NoSQL" now. Its wonderful to choose "SQL" for "NoSQL".
https://www.c-sharpcorner.com/article/road-to-az-2044/#:~:text=Core%20SQL%20API%2C%20default%20API%20for%20using%20Azure%20Cosmos%20DB%20enables%20querying%20your%20data%20with%20a%20language%20very%20close%20to%20SQL%3B
https://learn.microsoft.com/en-us/azure/cosmos-db/choose-api#coresql-api:~:text=API%20for%20NoSQL%20is%20native%20to%20Azure%20Cosmos%20DB.

upvoted 8 times

Esward 9 months, 2 weeks ago

given answers are correct

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

SQL and Item Id

upvoted 2 times

rol204 1 year ago

correct

upvoted 2 times

kampatra 1 year, 1 month ago

Answer is correct:
https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api

upvoted 4 times

dtctx 1 year, 6 months ago

Received this in test on 4/29 and passed the test.

Went with Table API and Item ID. I do not know if Table API is correct, but I am confident that Item ID is.

upvoted 4 times

sghaha 1 year, 6 months ago

https://docs.microsoft.com/ko-kr/azure/cosmos-db/choose-api

upvoted 1 times

Question #25 *Topic 3*

DRAG DROP -

You are implementing an Azure solution that uses Azure Cosmos DB and the latest Azure Cosmos DB SDK. You add a change feed processor to a new container instance.

You attempt to read a batch of 100 documents. The process fails when reading one of the documents. The solution must monitor the progress of the change feed processor instance on the new container as the change feed is read. You must prevent the change feed processor from retrying the entire batch when one document cannot be read.

You need to implement the change feed processor to read the documents.

Which features should you use? To answer, drag the appropriate features to the cored requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each cored selection is worth one point.

Select and Place:

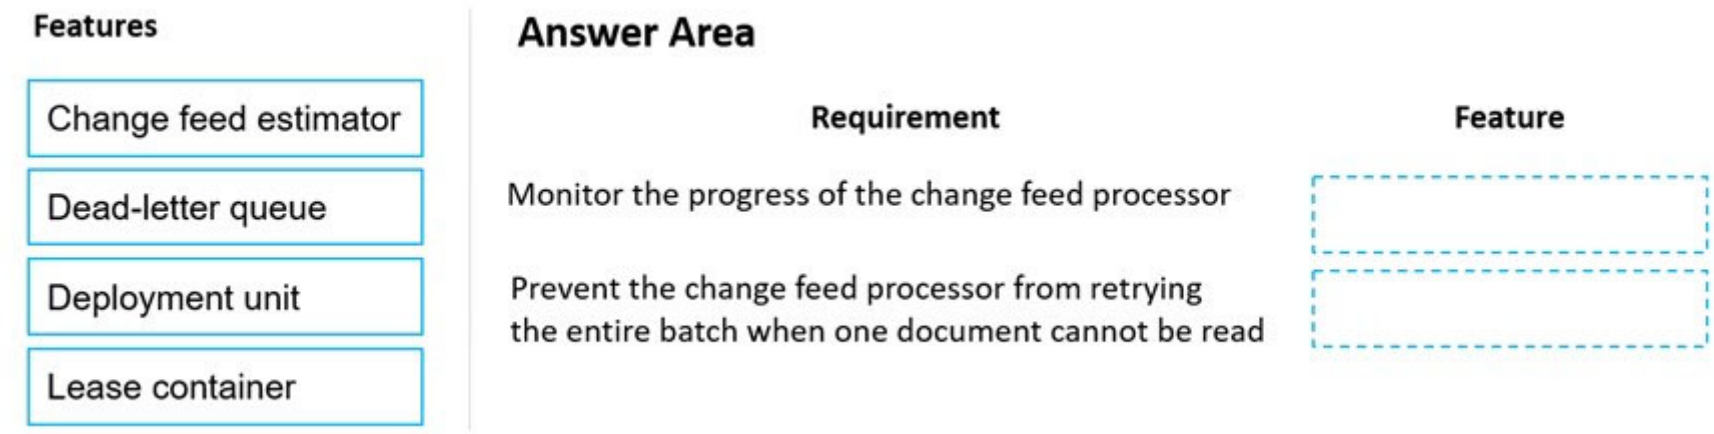

Correct Answer:

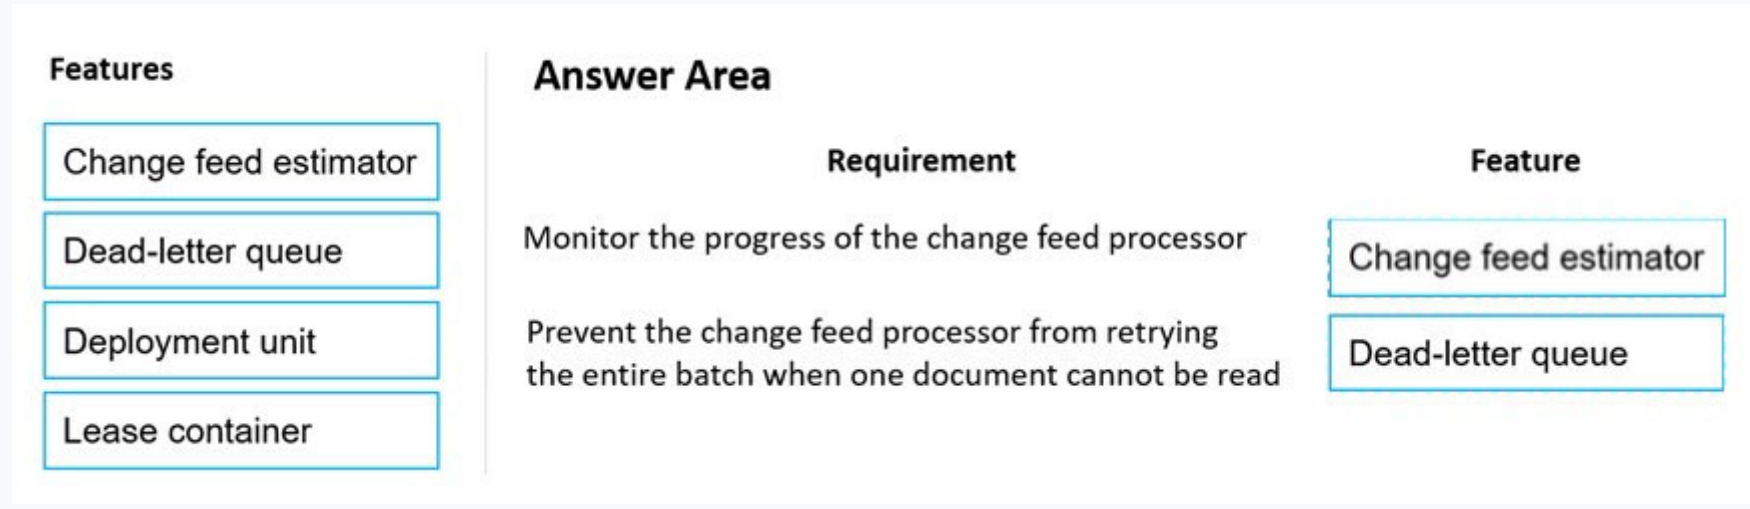

Box 1: Change feed estimator -

You can use the change feed estimator to monitor the progress of your change feed processor instances as they read the change feed or use the life cycle notifications to detect underlying failures.

Box 2: Dead-letter queue -

To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to a dead-letter queue. This design ensures that you can keep track of unprocessed changes while still being able to continue to process future changes. The dead-letter queue might be another Cosmos container. The exact data store does not matter, simply that the unprocessed changes are persisted.

Reference:

https://docs.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-processor

finnishr Highly Voted 1 year, 1 month ago

The answer is correct.

upvoted 11 times

sarmaria Highly Voted 7 months, 2 weeks ago

Got this on 16/03/23. Went with proposed solution. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 8 times

dddddd111 Most Recent 1 month, 1 week ago

I got this same question. Provided answers are correct. (Note: I failed the exam 20/9/23. I only scored 644 and I felt bad. I think because many questions here in Examtopics are not accurate. I suggest following the most voted answers and don't just not rely on Examtopics answers. At the beginning of the exam, you will be asked which programming languages you want to use. C#/Python. I chose C#. Also, I just want to add that some

questions here are really in the actual exams, but the choices are written and formatted differently. Please be aware of that. Goodluck. I feel bad for failing it, but I want to retake next month. I will try Python. T_T

upvoted 5 times

Nitin23 1 month, 1 week ago

this is a bad site

upvoted 1 times

kayvg 2 months, 2 weeks ago

The answer is correct

A change feed estimator is used to monitor teh progress of your change feed processor instances as they read the change feed

A dead-letter queue is holding a queue for messages that cannot bbe delivered to their destination

A deployment unit is to provide a container for an application or service, which is not relevant

The lease container is used to coordinate processing the change feed, which is not relevant aswell

upvoted 1 times

kvtborad 2 months, 3 weeks ago

I got this question on 6th August 2023. chose highly voted. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.

upvoted 1 times

JH81 4 months ago

Got this on 6/28/2023 and passed with 850. Went with answer.

upvoted 1 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

My case:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

narenazure 9 months ago

The answer is correct.
To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to an errored-message queue. This design ensures that you can keep track of unprocessed changes while still being able to continue to process future changes. The errored-message queue might be another Azure Cosmos DB container. The exact data store does not matter, simply that the unprocessed changes are persisted.
Error handling section - https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet

upvoted 1 times

Esward 9 months, 2 weeks ago

Given answers are correct!

upvoted 1 times

Esward 9 months, 2 weeks ago

To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to an errored-message queue. This design ensures that you can keep track of unprocessed changes while still being able to continue to process future changes. The errored-message queue might be another Azure Cosmos DB container. The exact data store does not matter, simply that the unprocessed changes are persisted.

In addition, you can use the change feed estimator to monitor the progress of your change feed processor instances as they read the change feed or use the life cycle notifications to detect underlying failures.

upvoted 2 times

Sreedharc 11 months, 2 weeks ago

got in 11/11/2022

upvoted 3 times

vcfvct 1 year ago

Got this on 10/21/2022

upvoted 3 times

OPT_001122 12 months ago

thanks for mentioning the date

upvoted 1 times

gmishra88 1 year ago

Sure, this is also taken from a certification. But where does a dead letter queue come in change feed processor. If this comes, yes, this is the correct answer. But it all doesn't make any sense

upvoted 1 times

coffecold 1 year ago

It is do-it-your self stuff:
"To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to a dead-letter queue "
https://learn.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-processor?tabs=dotnet

upvoted 3 times

Pentagon 12 months ago

If just for saving the last position, I think it should be lease container.
"When the delegate finishes processing the changes successfully, update the lease store with the latest processed point in time and go to #1"
Dead-letter queue seems to be prevent re-trying too much instead of resuming at the last done position.

upvoted 1 times

Question #26 *Topic 3*

HOTSPOT -

You are developing an application that uses a premium block blob storage account. The application will process a large volume of transactions daily. You enable

Blob storage versioning.

You are optimizing costs by automating Azure Blob Storage access tiers. You apply the following policy rules to the storage account. (Line numbers are included for reference only.)

```
01 {
02   "rules" : [
03     {
04       "name" : "versionRule",
05       "enabled" : true,
06       "type" : "Lifecycle",
07       "definition" : {
08         "actions" : {
09           "version" : {
10             "tierToCool" : {
11               "daysAfterCreationGreaterThan" : 60
12             },
13             "delete" : {
14               "daysAfterCreationGraterThan" : 365
15             }
16           }
17         },
18         "filters" : {
19           "blobTypes" : [ "blockBlob" ], "prefixMatch" : [ "transactions" ]
20         }
21       }
22     }
23   ]
24 }
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Block blobs prefixed with **transactions** will transition blobs that have not been modified in over 60 days to cool storage, and delete blobs not modified in 365 days	◯	◯
Blobs are moved to cool storage if they have not been accessed for 60 days	◯	◯
The policy rule tiers previous versions within a container named transactions that are 60 days or older to the cool tier and deletes previous versions that are 365 days or older	◯	◯
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool	◯	◯

Correct Answer:

Answer Area

Statements	Yes	No
Block blobs prefixed with **transactions** will transition blobs that have not been modified in over 60 days to cool storage, and delete blobs not modified in 365 days	○	○
Blobs are moved to cool storage if they have not been accessed for 60 days	○	○
The policy rule tiers previous versions within a container named transactions that are 60 days or older to the cool tier and deletes previous versions that are 365 days or older	○	○
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool	○	○

Box 1: No -

Would be true if daysAfterModificationGreaterThan was used, but here daysAfterCreationGreaterThan

Box 2: No -

Would need to use the daysAfterLastAccessTimeGreaterThan predicate.

Box 3: Yes -

Box 4: Yes -

With the lifecycle management policy, you can:

Transition blobs from cool to hot immediately when they are accessed, to optimize for performance.

Reference:

https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview

Dani_ac7 Highly Voted 1 year, 1 month ago

With this image, all answers are NO:

- Container named transaction is not in code
- is no present line "enableAutoTierToHotFromCool": true

upvoted 19 times

Tenk14 1 year, 1 month ago

I guess, third statement (The policy rule tiers..) result is Yes.
Container name "transactions" is in prefixMatch.

https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview#archive-data-after-ingest

Solution is:
- No
- No
- Yes
- No

upvoted 39 times

ChiragShah4885 1 year, 1 month ago

Container name "transactions" is in prefixMatch means its name must start from "transactions". Its name is not "transactions". Dani_ac7's answer looks correct

upvoted 2 times

ArturKon 1 year, 1 month ago

Looks like 3rd is NO, according to filter's guide:

"Filter blobs by name or first letters. To find items in a specific container, enter the name of the container followed by a forward slash, then the blob name or first letters. For example, to show all blobs starting with "a", type: "mycontainer/a"."

So it looks like we are searching blobs with "transactions" prefix in all containers in 3rd.

upvoted 1 times

cwn53066 1 year, 1 month ago

But name "transactions" satisfies rule `name must start from "transactions"`, why it doesn't fit?

upvoted 4 times

Esward 9 months, 2 weeks ago

Correct: if you want to match the blobs within a specific container, you should mention the conatiner name/blob name

{
"rules": [
{
"name": "agingRule",
"enabled": true,
"type": "Lifecycle",
"definition": {
"filters": {
"blobTypes": ["blockBlob"],
"prefixMatch": ["sample-container/blob1", "container2/blob2"]
},
"actions": {
"baseBlob": {
"tierToCool": { "daysAfterModificationGreaterThan": 30 },
"tierToArchive": { "daysAfterModificationGreaterThan": 90 }
}
}
}
}
]
}

upvoted 1 times

baroo1 Highly Voted 8 months, 2 weeks ago

No - Not modified, created.
No - Not accessed, created.
Yes - Rules are matching the statement. The prefix "transactions" can be applicable for containers as well. "container" / "container/blob" or "blob" can be used under this context.
Source: https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview#archive-data-after-ingest
No - "enableAutoTierToHotFromCool": "true" should be enabled.

upvoted 14 times

nikipediaa Most Recent 1 month ago

Question was on exam 2023-09-26

upvoted 1 times

nikipediaa 1 month ago

sorry, not this one, but similar with containers and numbers :)

upvoted 1 times

aragones 5 months, 2 weeks ago

Got this 2023-05-12.
Without 4th question.

my cases also:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 2 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

My case:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went N,N,Y, score 926
last box was not there only first three

upvoted 6 times

surprise0011 6 months, 4 weeks ago

premium block blob storage does not support access tiers. This is confusing

upvoted 3 times

Saluk_DE 7 months ago

Question was in Exam 2023-03-30

upvoted 1 times

sarmaria 7 months, 2 weeks ago

Got this on 16/03/23. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 4 times

rasojol523 7 months, 3 weeks ago

Box 3: No ?
'DayAfterCreationGraterThan 60' does not include 'just 60 days'. Therefore, if something has been made for 60 days, it would not be included in '60 days or later', so box 3 would be NO ?

upvoted 1 times

Ayman99 8 months ago

Was on the exam 27.02.2023
Went with NNYN.
Score 870

upvoted 3 times

warchoon 8 months, 1 week ago

Box 1, 2: No - this is about past versions, not about blobs
Box 3: Yes - exactly policy description
Box 4: No - first this is about past version not blob, second even blob moving back needs enableAutoTierToHotFromCool

upvoted 2 times

Adiu 8 months, 1 week ago

One important thing - the rule is for version and according to MS:
The baseBlob element in a lifecycle management policy refers to the current version of a blob. The version element refers to a previous version.

upvoted 2 times

Michael2023 9 months ago

Dani_ac7 is right I guess, but I think because we have here an Premium Block Blob Account/Storage which doesnt support Access Tier (Hot, Cool, Archive)... Access Tier is only supported in Standard general-purpose v2 Accounts

https://learn.microsoft.com/en-us/azure/storage/blobs/storage-feature-support-in-storage-accounts#premium-block-blob-accounts

upvoted 4 times

adilkhan 9 months, 2 weeks ago

Do I need to buy the contributor access to pass this exam?

upvoted 1 times

Jhilphis 9 months, 1 week ago

yes you need.

upvoted 2 times

dgcc97 10 months, 2 weeks ago

Answer is:
No
No
Yes
No

The third is Yes due to PrefixMatch being used to specifying containers.

Documentation where this question appears as example and as explanation of what it does:
https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview#manage-previous-versions

The fourth is No, due to no present line "enableAutoTierToHotFromCool": true

upvoted 4 times

Trimack93 11 months ago

It is worth mentioning that we have an application that uses a premium block blob storage account and as of now, tiering is not supported for these (see https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview#rule-actions). So it looks like 3rd answer (moving blobs to cold tier) should also be NO.

upvoted 5 times

Question #27 *Topic 3*

An organization deploys Azure Cosmos DB.
You need to ensure that the index is updated as items are created, updated, or deleted.
What should you do?

A. Set the indexing mode to Lazy.

B. Set the value of the automatic property of the indexing policy to False.

C. Set the value of the EnableScanInQuery option to True.

D. Set the indexing mode to Consistent.

Correct Answer: *D*
Azure Cosmos DB supports two indexing modes:
Consistent: The index is updated synchronously as you create, update or delete items. This means that the consistency of your read queries will be the consistency configured for the account.
None: Indexing is disabled on the container.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy

Community vote distribution
D (100%)

finnishr Highly Voted 1 year, 1 month ago
Correct !!!
upvoted 13 times

kayvg Most Recent 2 months, 2 weeks ago
Answer: D (consistent)
Setting indexing mode to consistent ensures that index is updated in sync with data

Setting it to lazy means index is updated async, meaning you may not get latest version of data

Setting the value of the automatic property to false means you have to manually manage indexing of data

Setting EnableScanInQuery to true allows query operations to perform scans across all documents, this is not really relevant + it provides poor performance
upvoted 3 times

lyggwtwtczxnhznebw 3 months, 3 weeks ago
Selected Answer: D
Setting the indexing mode to Consistent ensures that the index is always kept up to date with the changes made to the items in Azure Cosmos DB. With this mode, whenever an item is created, updated, or deleted, the index is automatically updated in real-time to reflect those changes.
upvoted 1 times

[Removed] 7 months, 3 weeks ago
Selected Answer: D
D indexing mode ensures that the index is always up-to-date with the data stored in the container and changes are immediately reflected in the index.
upvoted 1 times

Esward 9 months, 2 weeks ago
Given answer is correct.

Consistent: The index is updated synchronously as you create, update or delete items. This means that the consistency of your read queries will be the consistency configured for the account.
None: Indexing is disabled on the container. This mode is commonly used when a container is used as a pure key-value store without the need for secondary indexes. It can also be used to improve the performance of bulk operations. After the bulk operations are complete, the index mode can be set to Consistent and then monitored using the IndexTransformationProgress until complete.
upvoted 3 times

OPT_001122 11 months ago
Selected Answer: D
D. Set the indexing mode to Consistent.
upvoted 3 times

Question #28 *Topic 3*

You are developing a .Net web application that stores data in Azure Cosmos DB. The application must use the Core API and allow millions of reads and writes.
The Azure Cosmos DB account has been created with multiple write regions enabled. The application has been deployed to the East US2 and Central US regions.
You need to update the application to support multi-region writes.
What are two possible ways to achieve this goal? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Update the ConnectionPolicy class for the Cosmos client and populate the PreferredLocations property based on the geo-proximity of the application.

B. Update Azure Cosmos DB to use the Strong consistency level. Add indexed properties to the container to indicate region.

C. Update the ConnectionPolicy class for the Cosmos client and set the UseMultipleWriteLocations property to true.

D. Create and deploy a custom conflict resolution policy.

E. Update Azure Cosmos DB to use the Session consistency level. Send the SessionToken property value from the FeedResponse object of the write action to the end-user by using a cookie.

Correct Answer: *CD*
C: The UseMultipleWriteLocations of the ConnectionPolicy class gets or sets the flag to enable writes on any locations (regions) for geo-replicated database accounts in the Azure Cosmos DB service.
Note: Once an account has been created with multiple write regions enabled, you must make two changes in your application to the ConnectionPolicy for the
Cosmos client to enable the multi-region writes in Azure Cosmos DB. Within the ConnectionPolicy, set UseMultipleWriteLocations to true and pass the name of the region where the application is deployed to ApplicationRegion. This will populate the PreferredLocations property based on the geo-proximity from location passed in. If a new region is later added to the account, the application does not have to be updated or redeployed, it will automatically detect the closer region and will auto-home on to it should a regional event occur.
Azure core API application " ConnectionPolicy class" cosmos db multiple write regions enabled
D: With multi-region writes, when multiple clients write to the same item, conflicts may occur. When a conflict occurs, you can resolve the conflict by using different conflict resolution policies.
Note: Conflict resolution policy can only be specified at container creation time and cannot be modified after container creation.
Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.documents.client.connectionpolicy https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-multi-master https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-manage-conflicts

Community vote distribution
AC (85%)

coffecold Highly Voted 1 year ago
The goal is

"You need to update the application to support multi-region writes",

that is enable multi-region writes (bool, option C) and add the regions (option A)
Then you have to apply the Conflict resolution policies.This can be LLW(default, not mentioned) or custom (option D).

Hence : there is only ONE way to to support multi-region writes (both apply C AND A) and there are subsequently TWO ways to apply the Conflict resolution policies (@ SQL) to solve write, update and delete conflicts of which one is mentioned in the question (D).
To support multi-region writes I would answer A and C , but they have to be set both, not one or the other.
See https://learn.microsoft.com/en-us/azure/cosmos-db/sql/how-to-multi-master?tabs=api-async and https://learn.microsoft.com/en-us/azure/cosmos-db/conflict-resolution-policies
upvoted 16 times

dgcc97 10 months, 1 week ago
From the documentation you provided:
Within the ConnectionPolicy, set UseMultipleWriteLocations to true and pass the name of the region where the application is deployed to ApplicationRegion. This will populate the PreferredLocations property based on the geo-proximity from location passed in. If a new region is later added to the account, the application does not have to be updated or redeployed, it will automatically detect the closer region and will auto-home on to it should a regional event occur.

I take that paragraph to mean that it automatically updates the PreferredLocations property, when you set the UseMultipleWriteLocations property with true and pass the region name.

That means you actually only need to update UseMultipleWriteLocations, thus C.

And as you mentioned, apply the Conflict resolution policies (D).

upvoted 2 times

warchoon 8 months ago

It's a trick. "you must make two changes in your application to the ConnectionPolicy" "pass the name of the region where the application is deployed to ApplicationRegion. This will populate the PreferredLocations property". I think this is an "update"
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async#:~:text=you%20must%20make%20two%20changes%20in%20your%20application%20to%20the%20ConnectionPolicy

upvoted 1 times

baroo1 Highly Voted 8 months, 2 weeks ago

Selected Answer: AC

A + C =
Clearly stated in documentation that both are required: https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async.
And according to the question "Each correct answer presents part of the solution". "Part"

Not D -> "Custom conflict resolution policy is available only for API for NoSQL accounts and can be set only at creation time." The question states the account was already created.
Source: https://learn.microsoft.com/en-us/azure/cosmos-db/conflict-resolution-policies

upvoted 10 times

applepie Most Recent 3 months ago

got this question today, answer A,C - 7/30/2023, score 895/1000

upvoted 1 times

JH81 4 months ago

Selected Answer: AC

Got this on 6/28/2023 and passed with 850. Went with A, C.

upvoted 2 times

deathRac3 5 months, 1 week ago

Selected Answer: AC

A&C to support multi-region writes

upvoted 2 times

nekkilodeon 6 months, 1 week ago

Selected Answer: AC

A & C
https://learn.microsoft.com/en-us/azure/cosmos-db/sql/how-to-multi-master?tabs=api-async

upvoted 2 times

adilkhan 6 months, 3 weeks ago

A,C chat gpt

upvoted 2 times

vikas4239 8 months, 2 weeks ago

A&C

ConnectionPolicy, set UseMultipleWriteLocations to true and pass the name of the region where the application is deployed to ApplicationRegion. This will populate the PreferredLocations property based on the geo-proximity from location passed in

https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async

upvoted 1 times

AkhiRk 8 months, 3 weeks ago

Answer should be CD
The questions mentions .NET web application and if you check below link you can see for .net the option C goes
And for multi write scenarios conflict might arise and hence option D
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async

upvoted 2 times

Esward 9 months, 2 weeks ago

AC are the correct answers as per microsoft docs
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async

upvoted 1 times

imanonion 9 months, 3 weeks ago

Selected Answer: AC

Option C: Update the ConnectionPolicy class for the Cosmos client and set the UseMultipleWriteLocations property to true. This will enable the application to use multiple write regions when storing data in Azure Cosmos DB.
Option A: Update the ConnectionPolicy class for the Cosmos client and populate the PreferredLocations property based on the geo-proximity of

the application. This will allow the application to use the closest write region to the application's location, improving performance and reducing latency.
Other options, such as updating the consistency level or creating a custom conflict resolution policy, are not directly related to enabling multi-region writes and are not necessary for this goal.

upvoted 4 times

rotimislaw 10 months ago

Selected Answer: AC

As stated by others

upvoted 2 times

Siepie 10 months, 1 week ago

Correct answers should be A and C

```
ConnectionPolicy policy = new ConnectionPolicy
{
ConnectionMode = ConnectionMode.Direct,
ConnectionProtocol = Protocol.Tcp,
UseMultipleWriteLocations = true ----> C
};
policy.SetCurrentLocation("West US 2"); -----> A
```

upvoted 3 times

Siepie 10 months, 1 week ago

Sorry SetCurrentLocation is only the region in which the application is being deployed and where Azure Cosmos DB is replicated. But the ConnectionPolicy class does also have a property PreferredLocations: Gets and sets the preferred locations (regions) for geo-replicated database accounts in the Azure Cosmos DB service. For example, "East US" as the preferred location.

upvoted 1 times

OPT_001122 1 year ago

C and D are the correct answers?

upvoted 1 times

POOOJAAAAAAAAAA 10 months, 1 week ago

IMO C and D

upvoted 1 times

[Removed] 1 year, 1 month ago

The documentation says: " To enable multi-region writes in your application, set ApplicationRegion to the region in which the application is being deployed and where Cosmos DB is replicated"
Now the option to use that property is in .Net SDK V2. Microsoft says I do not need to know .Net for this exam and they ask questions from V2? So, I should try this in both v2 and v3 to be able to pass this exam.

upvoted 2 times

finnishr 1 year, 1 month ago

Correct!!!

upvoted 1 times

Dani_ac7 1 year, 1 month ago

Selected Answer: CD

Discarded questions:
B. Update Azure Cosmos DB to use the Strong consistency level. Add indexed properties to the container to indicate region. --> MAKES NO SENSE
E. Update Azure Cosmos DB to use the Session consistency level. Send the SessionToken property value from the FeedResponse object of the write action to the end-user by using a cookie. --> MAKES NO SENSE

upvoted 4 times

HOTSPOT -

You are developing a solution to store documents in Azure Blob storage. Customers upload documents to multiple containers. Documents consist of PDF, CSV,

Microsoft Office format and plain text files.

The solution must process millions of documents across hundreds of containers. The solution must meet the following requirements:

- Documents must be categorized by a customer identifier as they are uploaded to the storage account.
- Allow filtering by the customer identifier.
- Allow searching of information contained within a document
- Minimize costs.

You create and configure a standard general-purpose v2 storage account to support the solution.

You need to implement the solution.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Requirement	Solution
Search and filter by customer identifier	Azure Cognitive Search Azure Blob index tags Azure Blob inventory policy Azure Blob metadata
Search information inside documents	Azure Cognitive Search Azure Blob index tags Azure Blob inventory policy Azure Blob metadata

Correct Answer:

Answer Area

Requirement	Solution
Search and filter by customer identifier	Azure Blob index tags
Search information inside documents	Azure Cognitive Search

Box 1: Azure Blob index tags -
As datasets get larger, finding a specific object in a sea of data can be difficult. Blob index tags provide data management and discovery capabilities by using key- value index tag attributes. You can categorize and find objects within a single container or across all containers in your storage account. As data requirements change, objects can be dynamically categorized by updating their index tags. Objects can remain in-place with their current container organization.

Box 2: Azure Cognitive Search -
Only index tags are automatically indexed and made searchable by the native Blob Storage service. Metadata can't be natively indexed or searched. You must use a separate service such as Azure Search.
Azure Cognitive Search is the only cloud search service with built-in AI capabilities that enrich all types of information to help you identify and explore relevant content at scale. Use cognitive skills for vision, language, and speech, or use custom machine learning models to uncover insights from all types of content.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-manage-find-blobs https://azure.microsoft.com/en-us/services/search/

jeanfmc Highly Voted 1 year, 1 month ago
right
Azure Blob Index tags: https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-index-how-to?tabs=azure-portal
Azure Cognitive Search: Search inside documents
upvoted 17 times

TheExamMaster2020 Highly Voted 11 months, 2 weeks ago
Did my exam on 15th November 2022. This question was on it.
upvoted 7 times

RikinPatel Most Recent 4 weeks, 1 day ago
As per Goodle Bard this is correct ans.
Azure Blob index tags
Azure Cognitive Search
upvoted 1 times

[Removed] 3 months, 2 weeks ago
1. metadata
2. cognitive search
upvoted 2 times

MysticalSam 4 months, 3 weeks ago
This question was in today's exam on 10-June-2023
upvoted 3 times

Videira 8 months ago
On my exam 2023-02-25
upvoted 3 times

trini99 8 months ago

This question was on the 24/02/2023

upvoted 2 times

tmuras 8 months, 2 weeks ago

This question was on exam on 17/02/23.

upvoted 3 times

Esward 9 months, 2 weeks ago

Given answers are correct
Azure blob index tags
Cognitive serach
https://learn.microsoft.com/en-us/azure/search/search-blob-storage-integration

upvoted 4 times

ks1586 10 months, 1 week ago

Did my exam on 17th December 2022. This question was on it.

upvoted 3 times

uditrm12 10 months, 1 week ago

examtopic questions are enough to pass ?

upvoted 2 times

Ciupaz 1 month ago

Also follow some course on Udemy (or Pluralsight), just to see actions, and read all the discussions here on ExamTopics. Then, if you have time/energy, study Microsoft docs.

upvoted 1 times

HellBaby 10 months ago

No - you need a lot of other resources in order to get a better chance. The Microsoft exams are no joke.

upvoted 4 times

OPT_001122 10 months, 3 weeks ago

1. index tag
2. cognitive search

upvoted 2 times

Question #30 *Topic 3*

HOTSPOT -
You are developing a web application by using the Azure SDK. The web application accesses data in a zone-redundant BlockBlobStorage storage account.
The application must determine whether the data has changed since the application last read the data. Update operations must use the latest data changes when writing data to the storage account.
You need to implement the update operations.
Which values should you use? To answer, select the appropriate option in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

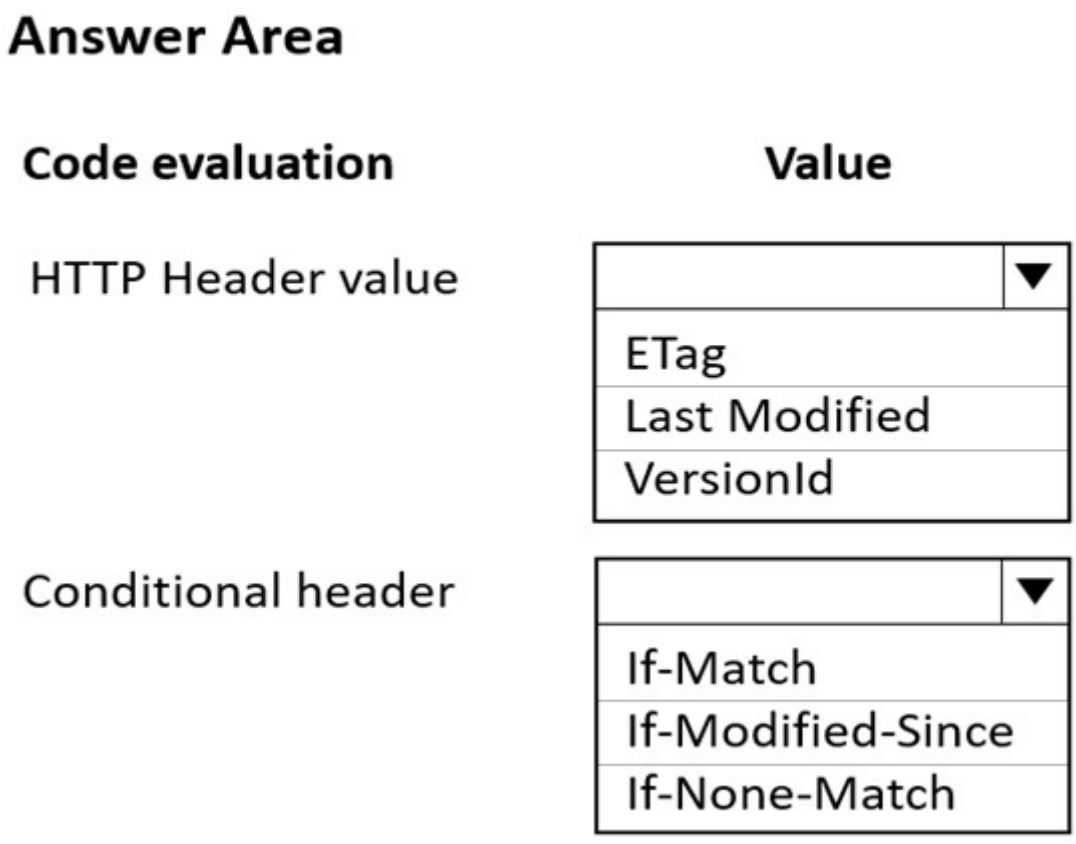

Correct Answer:

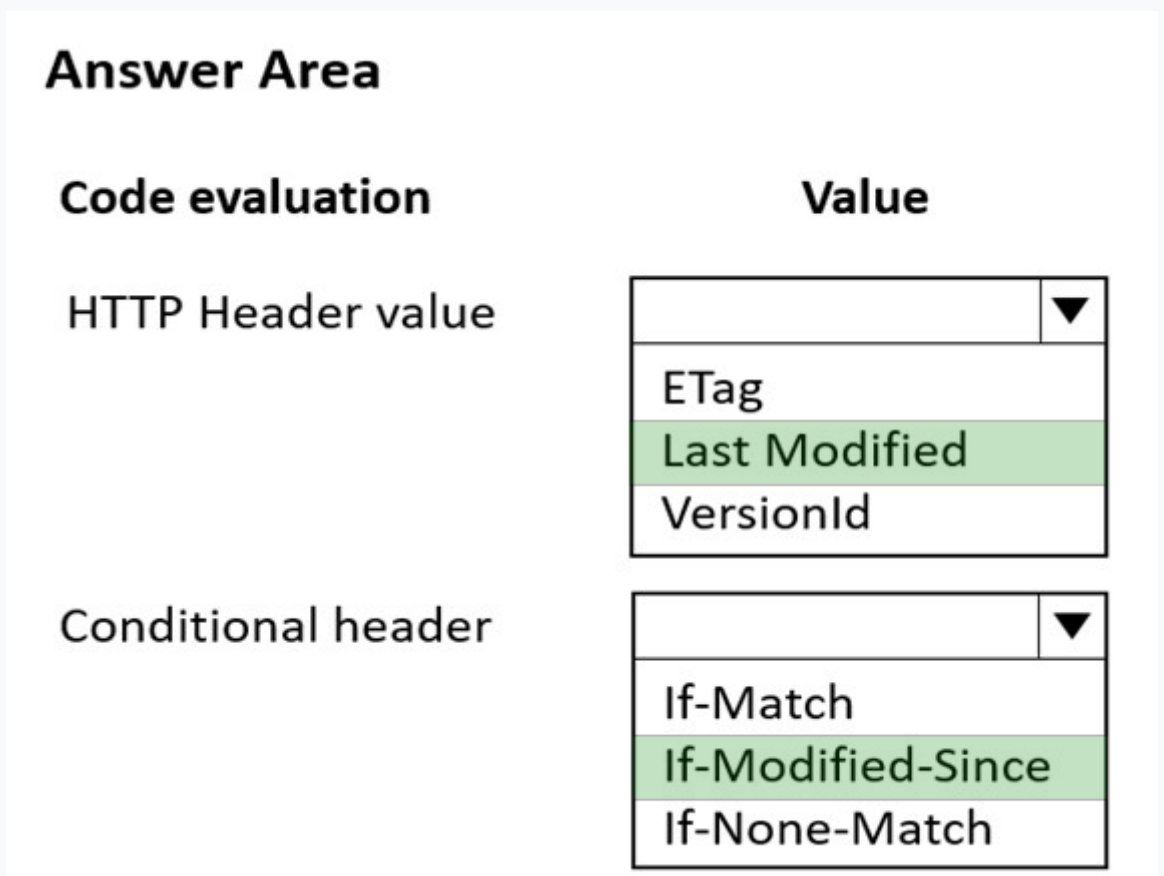

Box 1: Last Modified -
The Last-Modified response HTTP header contains a date and time when the origin server believes the resource was last modified. It is used as a validator to determine if the resource is the same as the previously stored one. Less accurate than an ETag header, it is a fallback mechanism.

Box 2: If-Modified-Since -
Conditional Header If-Modified-Since:
A DateTime value. Specify this header to perform the operation only if the resource has been modified since the specified time.
Incorrect:

Not ETag/If-Match -
Conditional Header If-Match:
An ETag value. Specify this header to perform the operation only if the resource's ETag matches the value specified. For versions 2011-08-18 and newer, the
ETag can be specified in quotes.
Reference:
https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Last-Modified https://docs.microsoft.com/en-us/rest/api/storageservices/specifying-conditional-headers-for-blob-service-operations

tcybu Highly Voted 1 year, 1 month ago

I think it should be:
- ETag - server returns this tag for a resource to ensure we operate on the same version of the resource in subsequent API calls
- If-Match - update is processed by the server only if the ETag provided matches the latest resource version ETag

The reason for that is we want to make sure we update the latest version of a resource:
"Update operations must use the latest data changes when writing"
So, when using Last-Modified with If-Modified-Since, the operation executes only when another client modifies the resource between our READ and WRITE operations.
If we wanted to use Last-Modified instead, we would need If-Unmodified-Since instead.

upvoted 29 times

Tenk14 1 year, 1 month ago

I agree with you. (ETag + If-Match)
https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/If-Match

upvoted 4 times

[Removed] Highly Voted 1 year, 1 month ago

I referred to the documentation also and I cannot convince my mind I care enough to even understand this because there is no chance I will ever have to know these options in details unless I have to use it. Who reads and understands all these information when this is available as documentation and with Microsoft documentation it is all about trial and error and hoping for the best

upvoted 16 times

richardc7 Most Recent 1 month ago

Got examn 28/09/23

upvoted 1 times

Nitin23 1 month, 1 week ago

both are wrong in given solution.
Correct anser - etag and If-match

upvoted 2 times

LxstInNxght 2 months, 2 weeks ago

It's ETag and If-Match.

The ETag (Entity Tag) is an HTTP header that represents a unique identifier for a version of the resource. When sending an update request, you can include the If-Match header with the ETag value of the data you previously read. If the current ETag of the resource matches the provided ETag, it means the data hasn't changed.

upvoted 1 times

[Removed] 3 months, 2 weeks ago

Correct answer

upvoted 1 times

calmriver44 5 months, 3 weeks ago

It is Etag and If-Match

check this page https://learn.microsoft.com/en-us/rest/api/storageservices/specifying-conditional-headers-for-blob-service-operations for conditional header "If-Match" - We can specify ETag value. Specify this header to perform the operation only if the resource's ETag matches the value specified.

upvoted 2 times

adilkhan 7 months ago

Etag and if-match ChatGPT

upvoted 1 times

kandrew5 8 months, 1 week ago

Since we also need the update to work which is a PUT operation the correct answer is ETag/If-Match

upvoted 1 times

mmdex 8 months, 2 weeks ago

ETag
If-Match

There are 2 possible options to validate an update against unwanted changes:
- ETag + If-Match
- Last-Modified + If-Unmodified-Since

The second option is not available here, so it must be the first one. Besides, even if it was available, it would still be better to check ETag, because Last-Modified is weaker (it has only 1-second resolution, so it does not detect 2 changes within 1 second, so you cannot ensure you have the latest version) and I guess the zone-redundancy may also play some role in how Last-Modified is handled.

upvoted 2 times

tmndungu 8 months, 2 weeks ago

For those who have taken AZ-204 test, should I trust ExamTopics answers or the Community's answers?

upvoted 5 times

rgullini 6 months, 3 weeks ago

You should trust the discussions and other users experiencies. The default answers in the questions might be right or wrong.

upvoted 5 times

Esward 9 months, 1 week ago

Etag, If-Match are correct answers as per MS docs
Optimistic concurrency: An application performing an update will, as part of its update, determine whether the data has changed since the application last read that data. For example, if two users viewing a wiki page make an update to that page, then the wiki platform must ensure that the second update does not overwrite the first update. It must also ensure that both users understand whether their update was successful. This strategy is most often used in web applications.
https://learn.microsoft.com/en-us/azure/storage/blobs/concurrency-manage?tabs=dotnet

upvoted 4 times

Dotar 9 months, 2 weeks ago

I think it should be E-tag and if-none-match.
See the following source on how e-tag works: https://en.wikipedia.org/wiki/HTTP_ETag

upvoted 1 times

Skelletcz 9 months, 3 weeks ago

I think the answer is correct. The question is asking explicitly if "data changed since the application last read the data", so we must evaluate against the date, hence If-Modified-Since.

upvoted 2 times

AsuraX 10 months ago

Answer is correct - in my opinion.

- ETag - The ETag (or entity tag) HTTP response header is an identifier for a specific version of a resource. (Version!)
- Last-Modified - The Last-Modified response HTTP header contains a date and time when the origin server believes the resource was last modified. (Date and Time!)
(Source: https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers)

In Question: The application must determine whether the data has changed since the application last read the data.
My understanding: We are working on dates, not versions! So, Last-Modified and If-Modified-Since are better options.

PS: I honestly have no idea why I need to know this for Azure Certification Exam :((

upvoted 7 times

capable 11 months ago

Got this question in the exam 27/11/2022, anyone, please mention the correct answer and explain

upvoted 1 times

OPT_001122 12 months ago

(ETag + If-Match) is correct?

upvoted 1 times

Question #31 *Topic 3*

HOTSPOT -

An organization deploys a blob storage account. Users take multiple snapshots of the blob storage account over time.

You need to delete all snapshots of the blob storage account. You must not delete the blob storage account itself.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
Delete (Azure.Storage.Blobs.Models.DeleteSnapshotsOption
snapshotsOption = Azure.Storage.Blobs.Models.[          ] . [          ] )
```

DeleteIfExists	IncludeSnapshots
DeleteSnapshotsOption	None
WithSnapshot	OnlySnapshots
WithSnapshotCore	

Correct Answer:

Answer Area

```
Delete (Azure.Storage.Blobs.Models.DeleteSnapshotsOption
snapshotsOption = Azure.Storage.Blobs.Models.[DeleteSnapshotsOption] . [OnlySnapshots] )
```

DeleteIfExists	IncludeSnapshots
DeleteSnapshotsOption	None
WithSnapshot	OnlySnapshots
WithSnapshotCore	

Box 1: DeleteSnapshotsOption -

Sample code in powershell:

//dont forget to add the include snapshots :)

await batchClient.DeleteBlobsAsync(listofURIforBlobs,

Azure.Storage.Blobs.Models.DeleteSnapshotsOption.IncludeSnapshots);

Sample code in .Net:

// Create a batch with three deletes

BlobBatchClient batchClient = service.GetBlobBatchClient();

BlobBatch batch = batchClient.CreateBatch();

batch.DeleteBlob(foo.Uri, DeleteSnapshotsOption.IncludeSnapshots); batch.DeleteBlob(bar.Uri, DeleteSnapshotsOption.OnlySnapshots);

batch.DeleteBlob(baz.Uri);

// Submit the batch

batchClient.SubmitBatch(batch);

Box 2: OnlySnapshots -

Reference:

https://docs.microsoft.com/en-us/dotnet/api/overview/azure/storage.blobs.batch-readme

https://stackoverflow.com/questions/39471212/programmatically-delete-azure-blob-storage-objects-in-bulks

coffecold Highly Voted 1 year ago

it is DeleteSnapshotsOption.OnlySnapshots but the explanation in the answer is bogus as usual.
see https://learn.microsoft.com/en-us/dotnet/api/azure.storage.blobs.models.deletesnapshotsoption?view=azure-dotnet

upvoted 15 times

micro9000 11 months, 3 weeks ago

Yeah, we can ignore the explanation on this one,
We can use OnlySnapshots

upvoted 2 times

Esward Highly Voted 8 months, 2 weeks ago

It was there in 13 Feb 2023 exam

upvoted 6 times

Pulk Most Recent 1 week ago

correct.

upvoted 1 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 2 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 4 times

Esward 9 months, 1 week ago

Given options are correct
https://learn.microsoft.com/en-us/dotnet/api/azure.storage.blobs.models.deletesnapshotsoption?view=azure-dotnet

upvoted 1 times

OPT_001122 10 months, 3 weeks ago

DeleteSnapshotsOption.OnlySnapshots

upvoted 4 times

finnishr 1 year, 1 month ago

Correct!!

upvoted 5 times

Question #32 *Topic 3*

HOTSPOT -

An organization deploys a blob storage account. Users take multiple snapshots of the blob storage account over time.

You need to delete all snapshots of the blob storage account. You must not delete the blob storage account itself.

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

NOTE: Each correct selection is worth one point.

Hot Area:

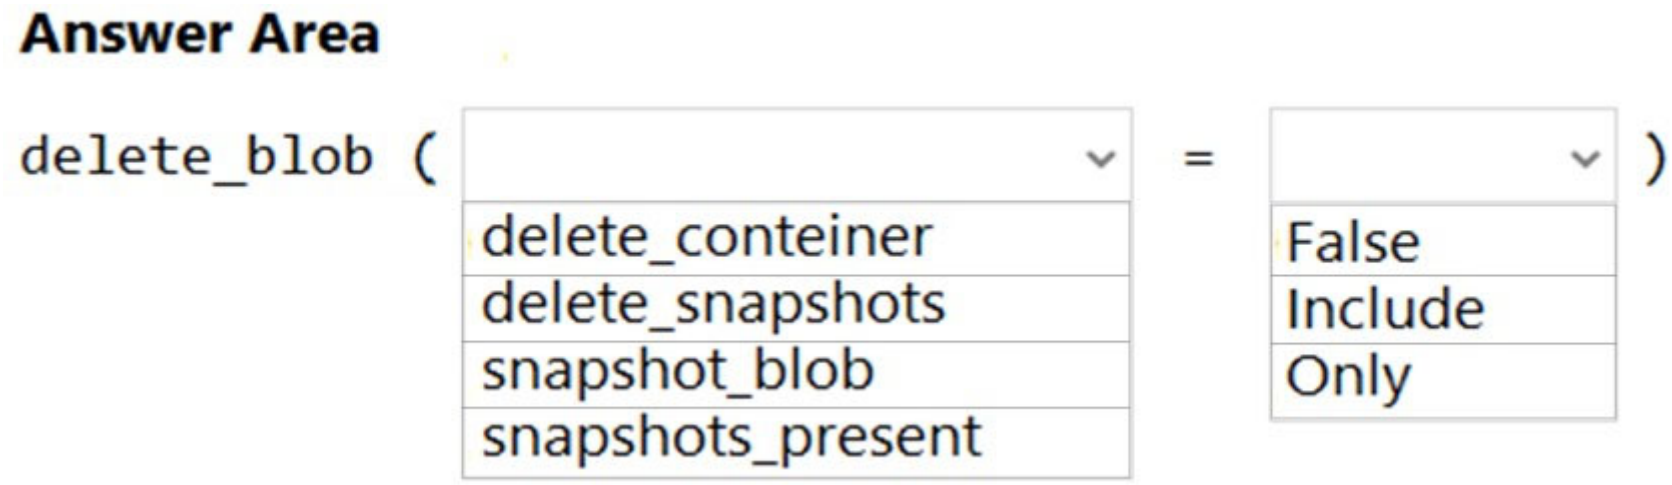

Correct Answer:

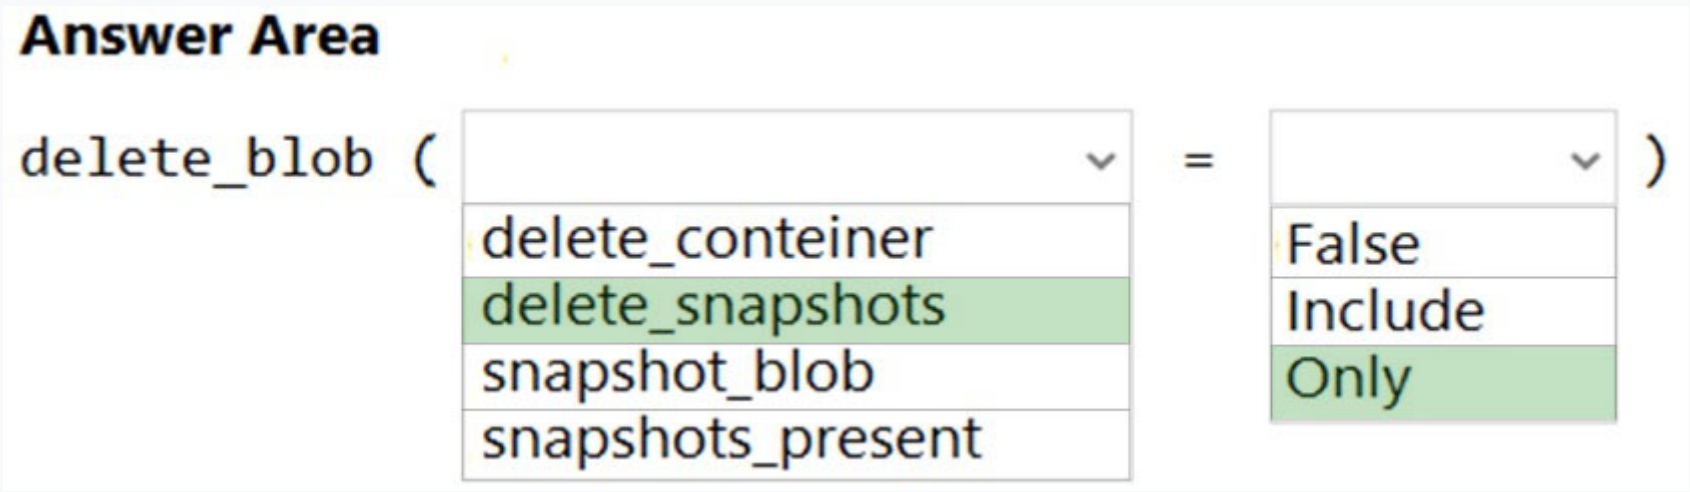

Box 1: delete_snapshots -

Delete only the snapshot (blob itself is retained)

blob_client.delete_blob(delete_snapshots="only")

Box 2: only -

Reference:

https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob_samples_common.py

[Removed] Highly Voted 1 year, 1 month ago

So, according to Microsoft a developer has to remember how the Microsoft guy chose to write in the .Net library and also how the python person thought and the unfortunate inconsistencies they built into these APIs?

upvoted 25 times

toysky731 10 months, 1 week ago

or you can retake the exam?

upvoted 4 times

bbq598 10 months, 1 week ago

haha I think you are right

upvoted 1 times

ElectricPants Highly Voted 1 year, 1 month ago

The answer looks correct and the link proves it:
Link to exact line with comment:
https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob_samples_common.py#L65

upvoted 15 times

Pulk Most Recent 1 week ago

given answers are correct.

upvoted 1 times

Esward 9 months, 1 week ago

Given answers are correct
https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob_samples_common.py

upvoted 1 times

gvasilev 10 months ago

Ah yes my favorite thing as developer - to remember what is something doing instead of google it.. lmao nice "exam" micro$oft

upvoted 3 times

Enigma___ 1 year, 1 month ago

Correct
Required if the blob has associated snapshots. Values include:

"only": Deletes only the blobs snapshots.

"include": Deletes the blob along with all snapshots.
https://learn.microsoft.com/en-us/azure/developer/python/sdk/storage/azure-storage-blob/azure.storage.blob.containerclient?view=storage-py-v12

upvoted 4 times

LucasSomeone1 1 year, 1 month ago

Weird comments. When there is only 1 or 2 comments, everyone is saying examtopics is correct...

upvoted 5 times

kampatra 1 year, 1 month ago

Correct!

upvoted 2 times

Question #33 *Topic 3*

HOTSPOT
-

You are developing an application that monitors data added to an Azure Blob storage account.

You need to process each change made to the storage account.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
var changeFeedClient = new BlobServiceClient("…").GetChangeFeedClient();
var x = default(string);
while (true)
{
    var changeFeed = changeFeedClient. [ ▼ ] ;
        GetChanges()
        GetChangesAsync()
        GetChanges(x).AsPages()
        GetChanges(x).GetEnumerator()

    foreach (var c in changeFeed)
    {
        x = c. [ ▼ ] ;
            ContinuationToken
            GetRawResponse().ReasonPhrase
            Values.Max(x = > x.EventTime).ToString()
            Values.Min(x = > x.EventTime).ToString()

        ProcessChanges(c.Values);

    }
}
```

Correct Answer:

Answer Area

```
var changeFeedClient = new BlobServiceClient("…").GetChangeFeedClient();
var x = default(string);
while (true)
{
    var changeFeed = changeFeedClient. [ ▼ ] ;
        GetChanges()
        GetChangesAsync()
        GetChanges(x).AsPages()
        GetChanges(x).GetEnumerator()

    foreach (var c in changeFeed)
    {
        x = c. [ ▼ ] ;
            ContinuationToken
            GetRawResponse().ReasonPhrase
            Values.Max(x = > x.EventTime).ToString()
            Values.Min(x = > x.EventTime).ToString()

        ProcessChanges(c.Values);

    }
}
```

tcybu Highly Voted 9 months, 3 weeks ago

I've tried the code in VS. Here's some thoughts:
1. box:
- GetChanges() - wrong - var c in the foreach would be BlobChangeFeedEvent which doesn't contain Values property used in ProcessChanges(c.Values) line below
- GetChangesAsync - wrong - code won't compile because it would require await foreach loop instead
- GetChanges(x).AsPages() - correct - it's the only option to make this code even compile
- GetChanges(x).GetEnumerator() - wrong - you cannot use IEnumerator type as foreach source
2. box:
- x = c.ContinuationToken - right - variable x was used as continuationToken parameter in changeFeedClient.GetChanges(x).AsPages() above
- c.GetRawResponse().ReasonPhrase - wrong - that does not make sense to use this value as continuation token

- x = c.Values.Min - wrong - continuation token is a number not date
- x = c.Values.Max - wrong - as above

So to sum up
1. changeFeedClient.GetChanges(x).AsPages()
2. x = c.ContinuationToken;

You can find more about Continuation Token here:
https://jessehouwing.net/azure-devops-accessing-apis-with-large-volumes-of-data/

upvoted 53 times

devex 2 months ago

I tried out the code and agree with the answer.

Just wanted to say that I find it really strange that this question is basically about the syntax of methods in a NuGet package which is not even released (it is currently in preview):
https://www.nuget.org/packages/Azure.Storage.Blobs.ChangeFeed/

Not sure what knowing the exact syntax actually proves...

upvoted 1 times

applepie 3 months ago

got this question today, go with this answer - 7/30/2023, score 895/1000

upvoted 2 times

Elbanna 1 month, 1 week ago

Does all questions came from exam topics ?

upvoted 1 times

Net_IT 1 week, 6 days ago

The majority of the questions

upvoted 1 times

warchoon 8 months ago

var x = default(string);
so, x is string
ContinuationToken is not string

upvoted 3 times

Michael2023 9 months ago

you are right...
1. changeFeedClient.GetChanges(x).AsPages() -> returns an IEnumerable<Page<BlobChangeFeedEvent>> ... when you loop through these pages "Page<BlobChangeFeedEvent>" you will get the options "page.ContinuationToken" and page.Values which are used in this example
100% Correct...
found the code here
https://github.com/Azure/azure-sdk-for-net/tree/main/sdk/storage/Azure.Storage.Blobs.ChangeFeed#resume-with-continuationtoken

upvoted 4 times

leviatas Most Recent 1 month ago

Got this quesiton in examn - 2023.09.25. Got Case Study Contoso

upvoted 2 times

NightshadeRC 3 months, 1 week ago

Had this question in today's exam: 2023-07-26

upvoted 1 times

Videira 8 months ago

On my exam 2023-02-25

upvoted 2 times

tmuras 8 months, 2 weeks ago

This question was on exam on 17/02/23.

upvoted 1 times

red0400 8 months, 1 week ago

what answer do you put ?

upvoted 1 times

Esward 9 months, 1 week ago

As per Microsoft docs:
GetchangesAsync(), contiuationtoken are correct

https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed-how-to

upvoted 2 times

AfPham 8 months, 2 weeks ago

Async option can be ignore, it requires 'await', which is not there in this situation

upvoted 5 times

drjakiell 9 months, 1 week ago

for GetchangesAsync you need the async await approach

upvoted 3 times

frks 9 months, 3 weeks ago

I believe it should be await... getchangesasync.... then getrawresponse, as x is a string. I don't see the connection to cancellationtoken.

upvoted 1 times

chettir01 9 months, 3 weeks ago

i think the fisrt box must be : getChangesAsync()
the second one is correct
source : https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed-how-to

upvoted 3 times

MikeSantos 6 months, 3 weeks ago

I don't agree. There is no await statement before the method

upvoted 2 times

Question #34 *Topic 3*

HOTSPOT
-

You develop an application that sells AI generated images based on user input. You recently started a marketing campaign that displays unique ads every second day.

Sales data is stored in Azure Cosmos DB with the date of each sale being stored in a property named 'whenFinished'.

The marketing department requires a view that shows the number of sales for each unique ad.
You need to implement the query for the view.

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

NOTE: Each correct selection is worth one point.

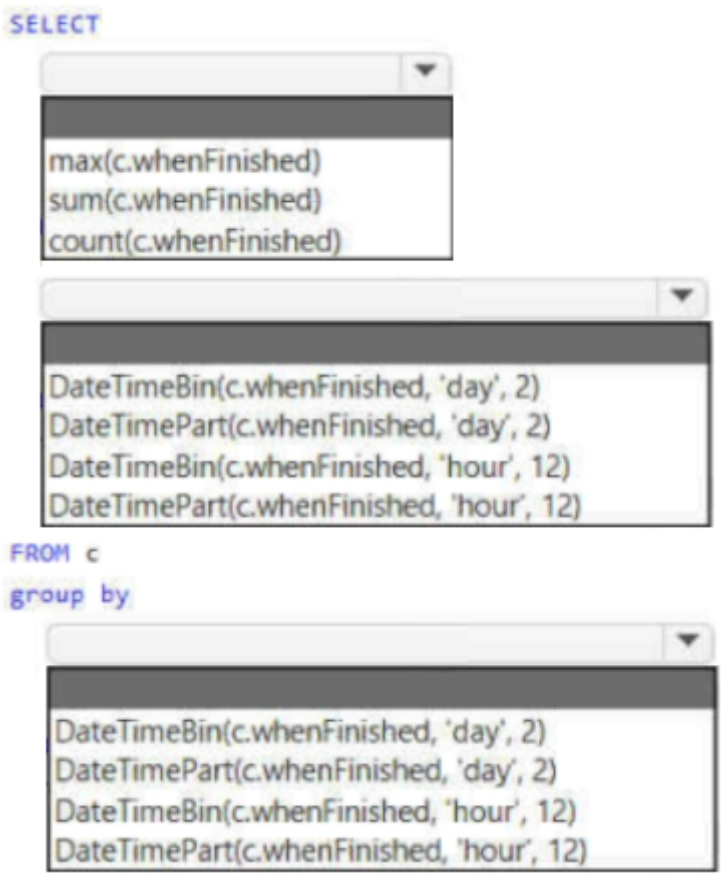

Correct Answer:

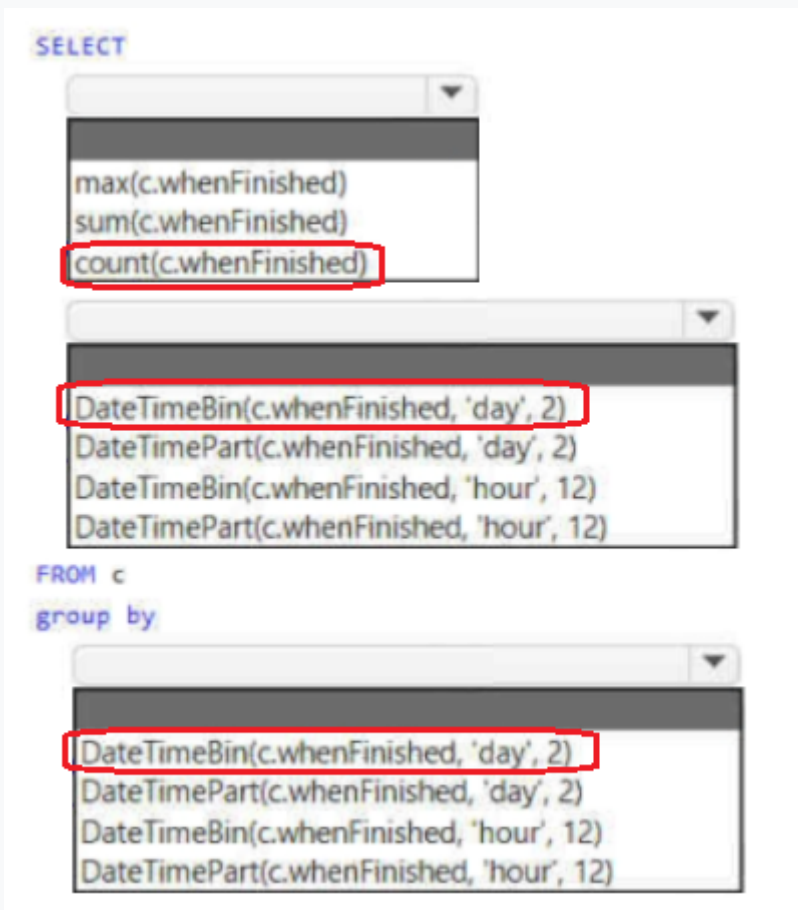

tcybu Highly Voted 9 months, 3 weeks ago

Correct!
Can't be DateTimePart as it takes two args only, see https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/query/date-time-functions

upvoted 18 times

Dotar 9 months, 2 weeks ago

You are right good sir I stand corrected.

upvoted 2 times

MysticalSam Highly Voted 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 5 times

richardc7 Most Recent 1 month ago

got today in 28-09-23

upvoted 1 times

janzak 6 months, 1 week ago

why count and not sum?

upvoted 2 times

MDWPartners 1 month, 1 week ago

the field is a date, not a number, so you cannot sum it.

upvoted 1 times

kgy01 5 months, 1 week ago

[...]The marketing department requires a view that shows the "number of sales"[...]

upvoted 1 times

Esward 9 months, 1 week ago

Given answers Count, DateTimeBin are correct answers as per MS docs
DateTimebin takes 3 and 4 parameters but DateTimepart takes only 2 args

https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/query/date-time-functions

upvoted 4 times

DP_Bhatt 9 months, 1 week ago

Correct!

upvoted 1 times

Dotar 9 months, 2 weeks ago

Anyone have an idea as to why it says day and not hour?

upvoted 2 times

plusJoyed 5 months, 2 weeks ago

It says every second day, so 12 hours would be twice a day. It might work if hours was 48

upvoted 1 times

Dotar 9 months, 3 weeks ago

Pretty sure it has to be DateTimePart(...day) option.
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/query/datetimepart

upvoted 1 times

Question #35 *Topic 3*

HOTSPOT
-

You implement an Azure solution to include Azure Cosmos DB, the latest Azure Cosmos DB SDK, and the Core (SQL) API. You also implement a change feed processor on a new container instance by using the Azure Functions trigger for Azure Cosmos DB.

A large batch of documents continues to fail when reading one of the documents in the batch. The same batch of documents is continuously retried by the triggered function and a new batch of documents must be read.

You need to implement the change feed processor to read the documents.

Which feature should you implement? To answer, select the appropriate features in the answer area.

NOTE: Each correct selection is worth one point.

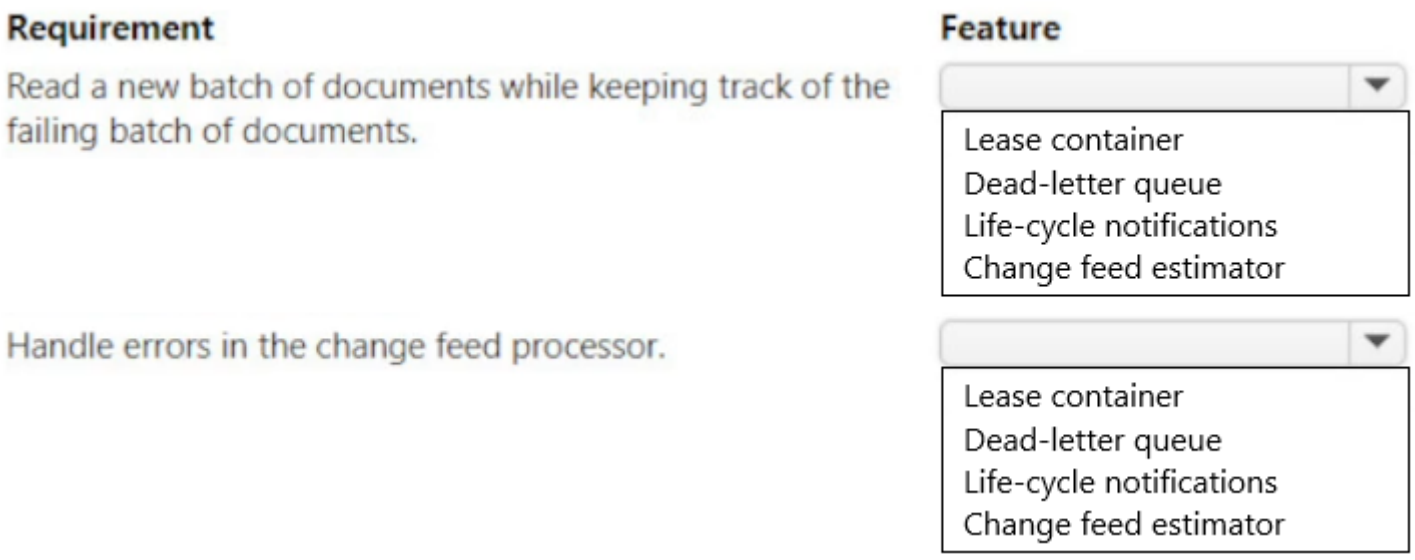

Correct Answer:

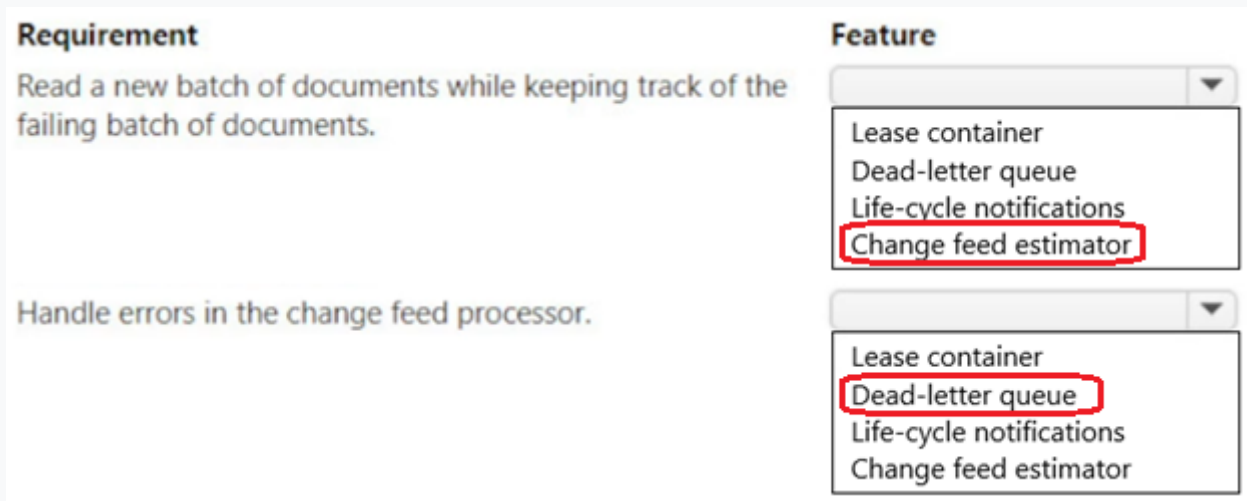

sarmaria Highly Voted 7 months, 2 weeks ago

Got this on 16/03/23. Went with proposed solution. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 10 times

ReyPirata 2 months, 1 week ago

Agreed.
This was on the exam (08/20/2023). Scored 925
Change feed estimator
Dead-letter queue

upvoted 3 times

imanonion Highly Voted 9 months, 3 weeks ago

I agree with the answer!
1. Change feed estimators monitor change feed processors' progress
2. Dead-letter queues handle errors and are able to monitor failed attempts, require failed attempts and even trigger a follow-up action (remediation or response)

upvoted 10 times

dom271219 Most Recent 1 month, 2 weeks ago

1/ "Change feed estimator" solely estimates the rate of changes, number of shards, and distrubution of the feed, but not read or process documents.

There's a error because by confusing with "Change feed processor".
2/ Dead-letter queue
upvoted 1 times

[Removed] 3 months, 2 weeks ago
should be vice versa
upvoted 3 times

juanckar 3 months, 3 weeks ago
This was on the exam (July 2023). Unsure about solution....
upvoted 2 times

aragones 5 months, 2 weeks ago
Got this on 2023may12.

my cases also:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada
upvoted 1 times

Defa 7 months, 3 weeks ago
From: https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet#error-handling

"To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to an errored-message queue. This design ensures that you can keep track of unprocessed changes while still being able to continue to process future changes. The errored-message queue might be another Azure Cosmos DB container. The exact data store does not matter, simply that the unprocessed changes are persisted.

In addition, you can use the change feed estimator to monitor the progress of your change feed processor instances as they read the change feed or use the life cycle notifications to detect underlying failures."

As the questions says 'keeping track' I'd go for feed estimator.
upvoted 2 times

Videira 8 months ago
On my exam 2023-02-25
upvoted 2 times

Ayman99 8 months ago
Was on the exam 27.02.2023
Went for Lifecycle notifications and dead letter queue.
Score 870
upvoted 3 times

mabdo 9 months ago
was on exam 02/2023
upvoted 3 times

alekomar 9 months ago
1. Life-cycle notifications
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet#life-cycle-notifications
Register a handler for WithErrorNotification to be notified when the current host encounters an exception during processing

2. Dead-letter queue
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet#error-handling
To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to an errored-message queue.
upvoted 3 times

warchoon 8 months ago
Correct
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet#error-handling:~:text=use%20the%20life%20cycle%20notifications%20to%20detect%20underlying%20failures.
upvoted 2 times

warchoon 8 months ago
but
1. (about the stuck) - Dead-letter and 2. (about error handling) - notifications
upvoted 3 times

Esward 9 months, 1 week ago
Given answers are correct.
upvoted 2 times

Esward 9 months, 1 week ago
https://thegreenerman.medium.com/at-least-once-delivery-using-the-azure-cosmos-db-change-feed-processor-6dd4130b59e3

upvoted 2 times

Question #36 *Topic 3*

You are developing an application to store business-critical data in Azure Blob storage.

The application must meet the following requirements:

- Data must not be modified or deleted for a user-specified interval.
- Data must be protected from overwrites and deletes.
- Data must be written once and allowed to be read many times.

You need to protect the data in the Azure Blob storage account.

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

NOTE: Each correct selection is worth one point.

A. Configure a time-based retention policy for the storage account.

B. Create an account shared-access signature (SAS).

C. Enable the blob change feed for the storage account.

D. Enable version-level immutability support for the storage account.

E. Enable point-in-time restore for containers in the storage account.

F. Create a service shared-access signature (SAS).

Correct Answer: *AF*

Community vote distribution

AD (85%) Other

linhbg Highly Voted 9 months, 3 weeks ago

I think the answer should be AD
A. Configure a time-based retention policy for the storage account
- A time-based retention policy stores blob data in a Write-Once, Read-Many (WORM) format for a specified interval. When a time-based retention policy is set, clients can create and read blobs, but can't modify or delete them. After the retention interval has expired, blobs can be deleted but not overwritten.
https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-time-based-retention-policy-overview
D. Before you can apply a time-based retention policy to a blob version, you must enable support for version-level immutability.
https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-policy-configure-version-scope?tabs=azure-portal

upvoted 27 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with above answer, score 926

upvoted 3 times

Esward 9 months, 1 week ago

I Agree with you A and D are correct answers

upvoted 3 times

Pulk Most Recent 1 week ago

correct answer is AD

upvoted 1 times

AashirwadSinha 2 months, 2 weeks ago

A and D are the correct answers

upvoted 1 times

andrewnm 3 months, 1 week ago

Selected Answer: AD

A is clearly a correct answer and this documentation confirms https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-policy-configure-version-scope?tabs=azure-portal that D is also required.

upvoted 1 times

deathRac3 5 months, 1 week ago

Selected Answer: AF

Feels correct
upvoted 1 times

aragones 5 months, 3 weeks ago
Got this type 2023-05-12.

My case:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada
upvoted 2 times

adilkhan 7 months ago
A,D chatgpt
upvoted 2 times

Saluk_DE 7 months ago
Question was in Exam 2023-03-30
upvoted 1 times

sarmaria 7 months, 2 weeks ago
Got this on 16/03/23. Went with AD Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.
upvoted 2 times

Ayman99 8 months ago
Was on the exam 27.02.2023
Went for A and D.
Score 870
upvoted 2 times

trini99 8 months ago
Which resources did you use to study?
upvoted 2 times

uffuchsi 8 months, 1 week ago
Selected Answer: AD
100% A and D, other options don't fit into this question (i.e., don't solve it)
upvoted 1 times

mabdo 9 months ago
on exam 02/2023
upvoted 1 times

alexein74 9 months, 1 week ago
Selected Answer: AD
A. Configuring a time-based retention policy for the storage account would ensure that the data in the storage account cannot be modified or deleted for a user-specified interval. This would meet the requirement that data must not be modified or deleted for a user-specified interval.

D. Enabling version-level immutability support for the storage account would ensure that the data in the storage account is protected from overwrites and deletes. This would meet the requirement that data must be protected from overwrites and deletes, and written once and allowed to be read many times.
upvoted 1 times

DP_Bhatt 9 months, 1 week ago
Selected Answer: AD
Should be A and D
upvoted 2 times

menlapwa 9 months, 2 weeks ago
Selected Answer: AD
Should be A and D
upvoted 2 times

ldenouter 9 months, 3 weeks ago
Selected Answer: AD
Answer should be AD
upvoted 2 times

imanonion 9 months, 3 weeks ago
Selected Answer: AD
See linhbg's answer
upvoted 2 times

Question #37 *Topic 3*

You are updating an application that stores data on Azure and uses Azure Cosmos DB for storage. The application stores data in multiple documents associated with a single username.

The application requires the ability to update multiple documents for a username in a single ACID operation.

You need to configure Azure Cosmos DB.

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

NOTE: Each correct selection is worth one point.

A. Create a collection sharded on username to store documents.

B. Configure Azure Cosmos DB to use the Gremlin API.

C. Create an unsharded collection to store documents.

D. Configure Azure Cosmos DB to use the MongoDB API.

Correct Answer: *CD*

Community vote distribution

CD (100%)

imanonion Highly Voted 9 months, 3 weeks ago

Selected Answer: CD

Agreed!
To update multiple documents for a username in a single ACID operation in Azure Cosmos DB, you need to ensure that the documents are stored in the same logical partition.

To do this, you should perform the following actions:

Create an unsharded collection to store documents. This will ensure that all documents are stored in the same logical partition.
Configure Azure Cosmos DB to use the MongoDB API. The MongoDB API supports multi-document ACID transactions, which allow you to update multiple documents in a single atomic operation.

upvoted 18 times

Dotar 9 months, 3 weeks ago

How do you know that the documents should be stored in the same logical partition. Is this a requirement for MongoDB API?

upvoted 2 times

surprise0011 6 months, 2 weeks ago

all operations like executing stored procedures or udf are effecting only items within single partition regardless of used API

upvoted 1 times

[Removed] Most Recent 3 months, 2 weeks ago

Correct answer is AD

upvoted 1 times

Bokomi 5 months, 3 weeks ago

Selected Answer: CD

To update multiple documents for a username in a single ACID operation in Azure Cosmos DB, you should perform the following actions:

C. Create an unsharded collection to store documents, since sharded collections do not support multi-document transactions.
D. Configure Azure Cosmos DB to use the MongoDB API, since it is the only API that supports multi-document transactions.
Therefore, the correct actions are C and D.

upvoted 3 times

Xin0613 6 months ago

Agreed!

upvoted 1 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 3 times

Esward 8 months, 2 weeks ago

It was there in 13 Feb 2023 exam
upvoted 2 times

Esward 9 months, 1 week ago

CD are correct
update multiple documents for a username in a single ACID operation.
upvoted 3 times

Esward 9 months, 1 week ago

https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/use-multi-document-transactions#requirements
upvoted 2 times

g2000 9 months, 2 weeks ago

Selected Answer: CD

multi-document transaction is only supported within unsharded collection.
https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/use-multi-document-transactions#requirements
upvoted 2 times

Question #38 *Topic 3*

You develop Azure solutions.

You must connect to a No-SQL globally-distributed database by using the .NET API.

You need to create an object to configure and execute requests in the database.

Which code segment should you use?

A. database_name = 'MyDatabase'
database = client.create_database_if_not_exists(id=database_name)

B. client = CosmosClient(endpoint, key)

C. container_name = 'MyContainer'
container = database.create_container_if_not_exists(
id=container_name, partition_key=PartitionKey(path="/lastName"), offer_throughput=400)

Correct Answer: *C*

Community vote distribution

B (100%)

imanonion Highly Voted 9 months, 3 weeks ago

Selected Answer: B

CosmosClient has to be created before you can do option A and C to create databases and execute requests.

```
client = CosmosClient(endpoint, key)
database_name = 'MyDatabase'
database = client.create_database_if_not_exists(id=database_name)
container_name = 'MyContainer'
container = database.create_container_if_not_exists(
id=container_name, partition_key=PartitionKey(path="/lastName"), offer_throughput=400 )
```

upvoted 14 times

[Removed] 3 months, 2 weeks ago

It doesn't matter. Question is asking about .NET API. A and C are not .NET code

upvoted 5 times

tempacc4nk Most Recent 2 months ago

I think its B, but its missing new keyword.

```
CosmosClient cosmosClient = new CosmosClient(
"connection-string-from-portal",
new CosmosClientOptions()
{
ApplicationRegion = Regions.EastUS2,
});
```

upvoted 1 times

[Removed] 3 months, 2 weeks ago

B is only one .NET code here. So, B is a correct answer

upvoted 2 times

kabbas 4 months, 3 weeks ago

Selected Answer: B

You need to create the client before anything else

upvoted 3 times

deathRac3 5 months, 1 week ago

Selected Answer: B

correct

upvoted 1 times

rcardenas72 8 months, 1 week ago

The correct code segment to configure and execute requests in the No-SQL globally-distributed database by using the .NET API would be option B:
client = CosmosClient(endpoint, key)

This code initializes the CosmosClient class and connects to the database with the specified endpoint and key parameters. This is the first step in creating a client to communicate with the Azure Cosmos DB.

Option A shows how to create a database, and option C shows how to create a container inside a database, but they do not create the client object that is needed to communicate with the database. Therefore, they are not correct options for this scenario.

upvoted 4 times

uffuchsi 8 months, 1 week ago

Has to be B - How could you do either A or C before B?

upvoted 1 times

BrettusMaximus 8 months, 2 weeks ago

A
You need to" create an object" to configure and execute requests in the "database".
Not B
Sure B must be done before A, but A is the actual database Object
Not C
C is the container and not the database object

upvoted 1 times

LeoIT 9 months, 1 week ago

B as “by using the .NET API”

upvoted 2 times

Esward 9 months, 1 week ago

B is the correct answer
https://learn.microsoft.com/en-us/python/api/azure-cosmos/azure.cosmos.cosmosclient?view=azure-python

upvoted 1 times

AmitPaul 9 months, 2 weeks ago

Selected Answer: B

CosmosClient has to be created first

upvoted 2 times

Kavin235 9 months, 3 weeks ago

Selected Answer: B

I believe Cosmos client type object can be used to configure and execute requests in the Azure Cosmos DB database service.
https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient?view=azure-dotnet

upvoted 1 times

Skelletcz 9 months, 3 weeks ago

Selected Answer: B

I vote for CosmosClient, which can connect to No-SQL database

upvoted 3 times

Nhiendo 9 months, 3 weeks ago

i think answer B

upvoted 2 times

chettir01 9 months, 3 weeks ago

i don't think so, in my point of view, cosmos client can be used to create your datababe, but not to manipulate it.
to do that, you have to create a container from this database, and use this object to do some requests !

upvoted 1 times

Question #39 *Topic 3*

You develop a web application that provides access to legal documents that are stored on Azure Blob Storage with version-level immutability policies. Documents are protected with both time-based policies and legal hold policies. All time-based retention policies have the AllowProtectedAppendWrites property enabled.

You have a requirement to prevent the user from attempting to perform operations that would fail only when a legal hold is in effect and when all other policies are expired.

You need to meet the requirement.

Which two operations should you prevent? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. adding data to documents

B. deleting documents

C. creating documents

D. overwriting existing documents

Correct Answer: *BD*

Community vote distribution

BD (100%)

minYoongi Highly Voted 6 months, 2 weeks ago

The legal hold policies prevent the deletion of the legal documents that are stored on Azure Blob Storage, but they do not prevent other types of operations such as read, write, or update. Therefore, you need to prevent the following two operations that would fail only when a legal hold is in effect and when all other policies are expired:

Delete operation: When a legal hold is in effect and all other policies are expired, attempting to delete the document will fail. Therefore, you should prevent the user from attempting to delete a document.

Overwrite operation: When a legal hold is in effect and all other policies are expired, attempting to overwrite the document with a new version will fail. Therefore, you should prevent the user from attempting to overwrite a document with a new version.

By preventing these two operations, you can ensure that the legal documents are not accidentally deleted or overwritten when they are protected by legal hold policies.
CHatgpt answer. Is this correct?

upvoted 12 times

mihailos 2 months ago

Got it in exam 28/08/23. Went with this answer. Scored 912

upvoted 1 times

IvanIco 4 months, 1 week ago

ty for the comment made it much clearer

upvoted 1 times

juanckar Highly Voted 3 months, 3 weeks ago

This was on the exam (July 2023). Went with BD. Scored 917

upvoted 5 times

b1992 3 months, 3 weeks ago

What was your casestudy?

upvoted 1 times

macobuzi 2 months, 1 week ago

What does it mean casestudy?

upvoted 1 times

training_bit2bit Most Recent 1 month, 1 week ago

legal hold policies: prevents delete and overwrite actions.
the question said that "legal hold" is in effect an the others expired, we should prevent other operations from user. so the answer should not be: A and C ?
B and D are already prevented with the legal hold policies, right?

upvoted 1 times

[Removed] 3 months, 2 weeks ago

I don't really understand this question

upvoted 3 times

adilkhan 6 months, 2 weeks ago

b,d IS CORRECT

upvoted 3 times

halfway 6 months, 2 weeks ago

Selected Answer: BD

BD. Ref: https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-legal-hold-overview

upvoted 3 times

Question #40 *Topic 3*

HOTSPOT
-

You provisioned an Azure Cosmos DB for NoSQL account named account1 with the default consistency level.

You plan to configure the consistency level on a per request basis. The level needs to be set for consistent prefix for read and write operations to account1.

You need to identify the resulting consistency level for read and write operations.

Which levels should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Operation type	Resulting consistency level
Read operations	strong / session / consistent prefix
Write operations	strong / session / consistent prefix

Correct Answer:

Answer Area

Operation type	Resulting consistency level
Read operations	strong / session / **consistent prefix**
Write operations	strong / **session** / consistent prefix

adilkhan Highly Voted 6 months, 1 week ago

Just passed the exam 26/04/2023. All the questions were from Exam topics. Got 970 marks!

upvoted 20 times

[Removed] 3 months, 2 weeks ago

useless comment

upvoted 6 times

LucreziaB 3 months, 2 weeks ago

Still more useful than yours :)

upvoted 20 times

UzrShkl 5 months, 1 week ago

Bro you had the contributor access?

upvoted 4 times

Rodikito 1 month, 4 weeks ago

of course

upvoted 1 times

adilkhan Highly Voted 6 months, 2 weeks ago

both should be consistent prefix

upvoted 14 times

Makonee 3 months, 1 week ago

Wrong. The Given answer is correct. Only Read Consistency can be set by client. Write stays at server setting which is session by default.
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level

upvoted 8 times

Vladimir_Gajinov 2 months, 1 week ago

You are right, the given answer is correct.
Microsoft documentation states:
Overriding the default consistency level only applies to reads within the SDK client. An account configured for strong consistency by default will still write and replicate data synchronously to every region in the account. When the SDK client instance or request overrides this with Session or weaker consistency, reads will be performed using a single replica.

upvoted 1 times

b1992 3 months, 3 weeks ago

could you explain why?

upvoted 1 times

richardc7 Most Recent 1 month ago

got in 28/09/2023

upvoted 1 times

seanzhenxiao 2 months, 1 week ago

Sorry, my first cmmt is error
I think should be :
read: consistent prefix,
write: session

1. default level is session
2. override only work for read
3. To move from weaker to stronger consistency, update the default consistency for the Azure Cosmos DB account

strong > boundary > session > consistent prefix > eventually

https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level

upvoted 2 times

seanzhenxiao 2 months, 1 week ago

I think should be session, session
1. default level is session
2. override only work for read
3. To move from weaker to stronger consistency, update the default consistency for the Azure Cosmos DB account

https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level

upvoted 1 times

Dianahu 3 months, 2 weeks ago

answer is here: https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level
Overriding the default consistency level only applies to reads within the SDK client.
defaut consistency is session

upvoted 6 times

[Removed] 3 months, 2 weeks ago

Now I get it! Appreciated.

upvoted 2 times

Mark16 4 months, 1 week ago

stupid question xD

upvoted 2 times

dy0917 6 months ago

Any explanation?
For read operations with the requirement of consistent prefix, the resulting consistency level option is C. Consistent prefix.
Following is what ChatGPT said, but I am not convinced.
For write operations, there is no consistency level option called "Consistent prefix". The available options are Strong, Bounded staleness, Session, and Eventual. If you want to ensure consistency prefix for write operations, you should choose Strong consistency level.

upvoted 2 times

Question #41 *Topic 3*

DRAG DROP
-

You are developing an application to store millions of images in Azure blob storage. The images are uploaded to an Azure blob storage container named companyimages contained in an Azure blob storage account named companymedia. The stored images are uploaded with multiple blob index tags across multiple blobs in the container.

You must find all blobs whose tags match a search expression in the container. The search expression must evaluate an index tag named status with a value of final.

You need to construct the GET method request URI.

How should you complete the URI? To answer, drag the appropriate parameters to the correct request URI targets. Each parameter may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Correct Answer:

adilkhan Highly Voted 6 months, 2 weeks ago
Answer is correct!
upvoted 16 times

Dixavado Most Recent 1 month ago
It was on my exam today (2023-09-26) I went with the examtopics answer - score 850
upvoted 1 times

lyggwtwtczxnhznebw 3 months, 2 weeks ago
Correct
upvoted 2 times

DonH 4 months, 2 weeks ago
Just for information: I just had this question on my AZ204 exam - 16-jun-2023.
I barely made it (with only 767 points) so I can't inform anyony if this answer is correct or not, just stating that this is an actual exam question.
upvoted 4 times

NavinD 3 months, 2 weeks ago
Do you have contributor access? what other websites did you study from?
upvoted 1 times

WorkFromHome 5 months ago
Got this question on the exam at 2023/05/31
upvoted 3 times

Question #42 *Topic 3*

HOTSPOT
-

You develop two Python scripts to process data.

The Python scripts must be deployed to two, separate Linux containers running in an Azure Container Instance container group. The containers must access external data by using the Server Message Block (SMB) protocol. Containers in the container group must run only once.

You need to configure the Azure Container Instance.

Which configuration value should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Configuration Setting	Configuration Value
External data volume	Secret / Empty directory / Cloned git repo / Azure file share
Container restart policy	Never / Always / OnFailure

Correct Answer:

Answer Area

Configuration Setting	Configuration Value
External data volume	Secret / Empty directory / Cloned git repo / **Azure file share**
Container restart policy	**Never** / Always / OnFailure

junix_ Highly Voted 3 months, 3 weeks ago

I believe the answer is correct.

Azure file share: https://learn.microsoft.com/en-us/azure/storage/files/files-smb-protocol?tabs=azure-portal

upvoted 9 times

ay_m Highly Voted 2 months, 2 weeks ago

Azure File Share is the only option that supports SMB

upvoted 6 times

Vladimir_Gajinov 2 months, 1 week ago

You are right

upvoted 1 times

arpannaprap Most Recent 3 months, 1 week ago

Answer is correct.

upvoted 3 times

[Removed] 3 months, 1 week ago

Incorrect
Correct Answer: Empty Directory, Never

upvoted 1 times

Papee 2 months, 4 weeks ago

why Empty Directory ?

upvoted 1 times

Question #43 *Topic 3*

HOTSPOT
-

You are developing a static website hosted on Azure Blob Storage. You create a storage account and enable static website hosting.

The website must support the following requirements:

- Custom domain name
- Custom header values for all responses
- Custom SSL certificate

You need to implement the static website.

What should you configure? To answer, select the appropriate values in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Configuration Value
Custom domain name	Blob index tags Azure Content Delivery Network (CDN) Cross-Origin Resource Sharing (CORS) Azure Storage Service Encryption (SSE)
Custom header values	Blob index tags Azure Content Delivery Network (CDN) Cross-Origin Resource Sharing (CORS) Azure Storage Service Encryption (SSE)
Custom SSL certificate	Blob index tags Azure Content Delivery Network (CDN) Cross-Origin Resource Sharing (CORS) Azure Storage Service Encryption (SSE)

Correct Answer:

Answer Area

Requirement	Configuration Value
Custom domain name	Blob index tags **Azure Content Delivery Network (CDN)** Cross-Origin Resource Sharing (CORS) Azure Storage Service Encryption (SSE)
Custom header values	Blob index tags **Azure Content Delivery Network (CDN)** Cross-Origin Resource Sharing (CORS) Azure Storage Service Encryption (SSE)
Custom SSL certificate	Blob index tags **Azure Content Delivery Network (CDN)** Cross-Origin Resource Sharing (CORS) Azure Storage Service Encryption (SSE)

AbidooKing Highly Voted 3 months, 3 weeks ago

Answer seems correct.
"Static websites have some limitations. For example, If you want to configure headers, you'll have to use Azure Content Delivery Network (Azure CDN)" and " To enable HTTPS, you'll have to use Azure CDN because Azure Storage doesn't yet natively support HTTPS with custom domains"

https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-static-website

upvoted 13 times

ericguo2020 Most Recent 1 month, 3 weeks ago

Using Azure CDN to Specify Custom HTTP Headers for an Azure Static Website Hosted SPA: https://medium.com/datadigest/using-azure-cdn-to-specify-custom-http-headers-for-an-azure-static-website-hosted-spa-41a9b9ec1674

upvoted 1 times

Vladimir_Gajinov 2 months, 1 week ago

The answers are correct.
The most appropriate answer is Azure CDN, as it can fulfill all three requirements: custom domain name, custom header values, and a custom SSL certificate.

upvoted 2 times

Papee 2 months, 3 weeks ago

1. Custom Domain Name:
- Configuration: Azure Content Delivery Network (CDN)
Azure CDN can be configured to map your custom domain name (e.g., www.yourdomain.com) to the Azure Blob Storage static website endpoint. This allows you to access your static website using your custom domain.
2. Custom Header Values for All Responses:
- Configuration: Blob index tag
You can use blob index tags to set custom header values for all responses from your static website. Blob index tags allow you to define metadata at the container or blob level that can be used as headers for the website content.
3. Custom SSL Certificate:
- Configuration: Azure Content Delivery Network (CDN)
To use a custom SSL certificate for your custom domain, you can configure Azure CDN to use the custom SSL certificate. This will enable secure communication between the user's browser and the CDN endpoint serving your static website content.

upvoted 3 times

NightshadeRC 3 months, 1 week ago

Had this question in today's exam: 2023-07-26

upvoted 2 times

Question #44 *Topic 3*

You are developing an inventory tracking solution. The solution includes an Azure Function app containing multiple functions triggered by Azure Cosmos DB. You plan to deploy the solution to multiple Azure regions.

The solution must meet the following requirements:

• Item results from Azure Cosmos DS must return the most recent committed version of an item.
• Items written to Azure Cosmos DB must provide ordering guarantees.

You need to configure the consistency level for the Azure Cosmos DB deployments.

Which consistency level should you use?

A. consistent prefix

B. eventual

C. bounded staleness

D. strong

E. session

Correct Answer: *D*

Community vote distribution

D (88%) 13%

LuchoGonzalez Highly Voted 3 months, 3 weeks ago

Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.
Source: https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels

upvoted 8 times

Xardas Highly Voted 3 months, 2 weeks ago

I believe its C 'bounded staleness'
Because strong consistency is not supported for multiple regions

upvoted 7 times

tarek0811 3 months, 1 week ago

Azure Cosmos DB accounts configured with multiple write regions can't be configured for strong consistency as it isn't possible for a distributed system to provide an RPO of zero and an RTO of zero
https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels

upvoted 3 times

RikinPatel Most Recent 4 weeks, 1 day ago

Selected Answer: D

strong is correct ans.:
Strong consistency: Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Bounded staleness consistency:
With Bounded Staleness consistency, reads issued against a non-primary region may not necessarily return the most recent version of the data globally, but are guaranteed to return the most recent version of the data in that region, which will be within the maximum staleness boundary globally.

https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels

upvoted 1 times

tamerm 1 month, 2 weeks ago

Selected Answer: D

I think it's D.
Strong consistency is not supported for multiple WRITE regions, and I don't see any such requirements in the question.
https://learn.microsoft.com/en-us/azure/cosmos-db/high-availability#multiple-write-regions

upvoted 1 times

RikinPatel 4 weeks, 1 day ago

as per Que.: You plan to deploy the solution to multiple Azure regions.

upvoted 1 times

Vladimir_Gajinov 2 months, 1 week ago

Selected Answer: C

I think it's Bounded Staleness because it's not mentioned that read is executed right after write operation so I think it's acceptable to have small latency between write and read operations. I would go with this!

upvoted 1 times

drkpnsh 2 months, 2 weeks ago

But for Strong consistency, it also says Bounded Staleness in a multi-write account is an anti-pattern. So I really don't understand what the correct answer is.

Bounded Staleness or Strong?

upvoted 2 times

Papee 2 months, 4 weeks ago

C. bounded staleness
Item results return the most recent committed version of an item: Bounded staleness ensures that reads are served from a version of the data that is within a specified lag from the current write operation. This means that you can read the most recent committed version of an item while still providing a level of consistency.

Items written to Azure Cosmos DB must provide ordering guarantees: Bounded staleness guarantees that writes are ordered and committed before being read.

While other consistency levels like "Strong" and "Consistent prefix" also provide strong consistency, "Bounded staleness" offers a balance between consistency and availability, which is suitable for scenarios where the most recent data is needed with some level of delay.

upvoted 3 times

ay_m 2 months, 2 weeks ago

"Reads when using Bounded Staleness returns the latest data available in that region by reading from two available replicas in that region."

With bounded staleness there is a certain amount of lag, so it does not return the most recently committed write

upvoted 1 times

Xardas 3 months ago

Correct answer is C 'bounded staleness'

Question mentions deployment to multiple regions, and strong consistency is not supported for that.

upvoted 2 times

junix_ 3 months, 3 weeks ago

Selected Answer: D

I believe `Strong` consistency is correct.

upvoted 3 times

75daltonic 3 months, 3 weeks ago

Selected Answer: D

I think is strong

upvoted 2 times

Question #45 *Topic 3*

HOTSPOT
-

You are developing an application that runs in several customer Azure Kubernetes Service clusters. Within each cluster, a pod runs that collects performance data to be analyzed later. A large amount of data is collected so saving latency must be minimized.

The performance data must be stored so that pod restarts do not impact the stored data. Write latency should be minimized.

You need to configure blob storage.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
apiVersion: storage.k8s.io/v1
kind: [ ]
        PodStorage
        StorageClass
        PersistentVolume
        PersistentVolumeClaim
metadata:
  name: data-store
provisioner: kubernetes.io/ [ ]
                              azure-disk
                              azure-file
                              portworx-volume
                              scaleio
parameters:

  skuName: Premium_LRS
reclaimPolicy: [ ]
                 local
                 retain
                 delete
```

Correct Answer:

Answer Area

```
apiVersion: storage.k8s.io/v1
kind: [ ]
metadata:
  name: data-store
provisioner: kubernetes.io/[ ]
parameters:
  skuName: Premium_LRS
reclaimPolicy: [ ]
```

kind: PodStorage, **StorageClass**, PersistentVolume, PersistentVolumeClaim

provisioner: kubernetes.io/ **azure-disk**, azure-file, portworx-volume, scaleio

reclaimPolicy: local, **retain**, delete

richardc7 1 month ago

got in 28/09/23 , totally out of scope

upvoted 3 times

Evoila_TrainingMaterial 2 months, 2 weeks ago

Sadly got this question 08/17/23 in the exam

upvoted 2 times

macobuzi 1 month, 3 weeks ago

Thanks for mentioning, so this question is definitely not out of scope! In fact, I don't even believe "scope" is in Microsoft's dictionary

upvoted 7 times

oCaio 2 months, 2 weeks ago

Out of scope. Suggested answer seems correct. Ref:
https://learn.microsoft.com/en-us/azure/aks/azure-csi-disk-storage-provision
https://kubernetes.io/docs/concepts/storage/storage-classes/#azure-disk

Ref azure-file
https://learn.microsoft.com/en-us/azure/aks/azure-csi-files-storage-provision#create-a-storage-class
(Create a storage class)
https://kubernetes.io/docs/concepts/storage/storage-classes/#azure-file

Ref PodStorage / Persistent Volume / PersistentVolumeClaim
https://kubernetes.io/docs/tasks/configure-pod-container/configure-persistent-volume-storage/

Ref portworx-volume
https://docs.portworx.com/portworx-enterprise/operations/operate-kubernetes/storage-operations/kubernetes-storage-101/volumes

Ref Reclaim Policy
https://kubernetes.io/docs/concepts/storage/storage-classes/#reclaim-policy

upvoted 4 times

macobuzi 2 months, 1 week ago

As someone mentioned to received this question recently, it is not out of scope. Actually, I am no longer care if it is out of scope because Microsoft doesn't give a f about it!

upvoted 2 times

Mpho7 2 months, 4 weeks ago

Crazy that people are getting out-of-scope questions in the official exam. Perhaps its not counted in the final score. We wont know

upvoted 4 times

macobuzi 2 months, 1 week ago

Or maybe, "out-of-scope" is not within Microsoft dictionary!

upvoted 3 times

BaoNguyen2411 3 months ago

Got this question 29/06/2023

upvoted 3 times

[Removed] 3 months, 3 weeks ago

Out of scope! Nevertheless, seems to be correct.

upvoted 3 times

Question #46 *Topic 3*

HOTSPOT
-

Case study
-

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

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

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

To start the case study
-
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background
-

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment
-

Corporate website
-

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations
-

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements
-

The application components must meet the following requirements:

Corporate website

-

• Secure the website by using SSL.
• Minimize costs for data storage and hosting.
• Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).
• Distribute the website content globally for local use.
• Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.
• The website must have 99.95 percent uptime.

Retail store locations

-

• Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.
• Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services

-

• Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.
• Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services

-

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to
include read-only access to the data.

Security

-

• All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
• Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues

-

Retail Store Locations

-

• You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
• Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

You need to implement the delivery service telemetry data.

How should you configure the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

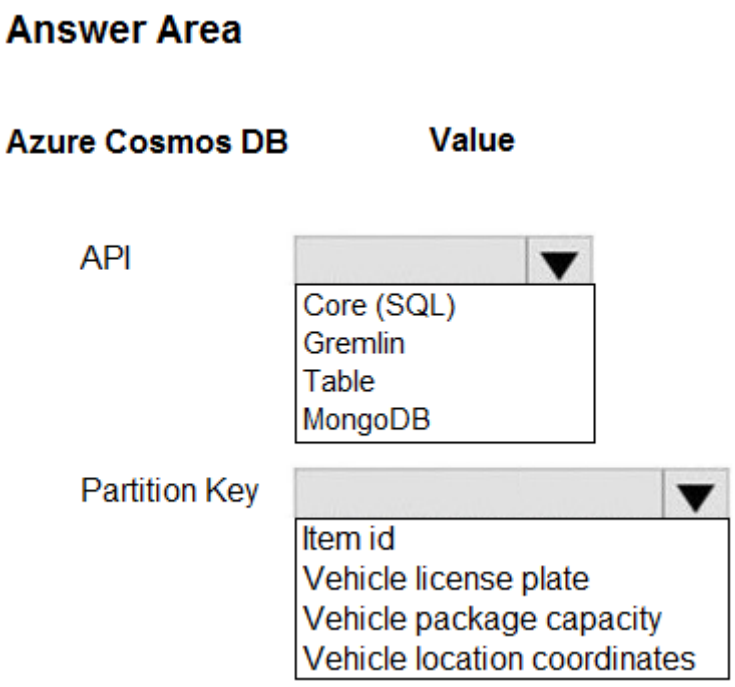

Correct Answer:

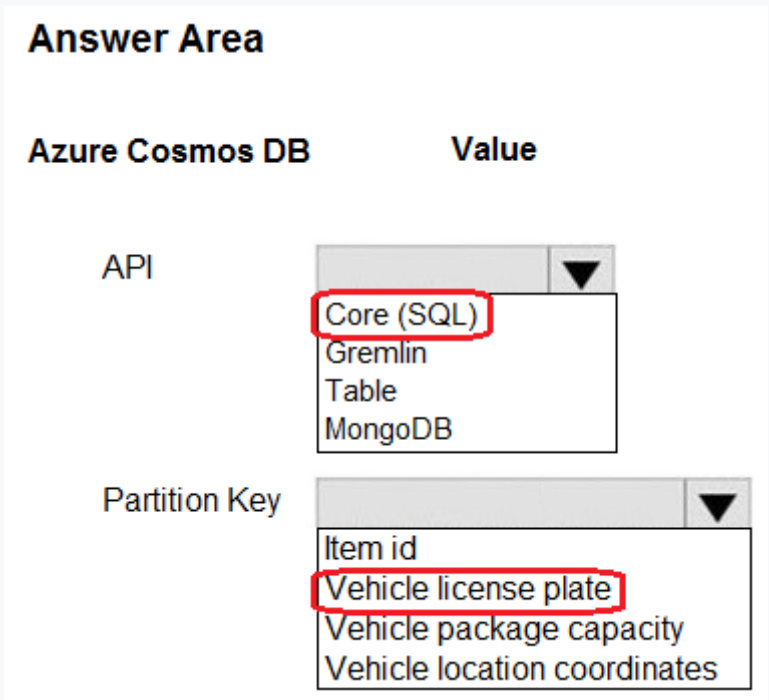

devex Highly Voted 2 months, 2 weeks ago

API: "Core (SQL)"
Partition Key: "Item id"

See: https://learn.microsoft.com/en-us/azure/cosmos-db/partitioning-overview#use-item-id-as-the-partition-key
upvoted 13 times

lyggwtwtczxnhznebw 2 months, 2 weeks ago

That makes a lot more sense, thanks for the link
upvoted 1 times

WH16 Highly Voted 2 months, 3 weeks ago

Why item.id is not used for Partition Key?
upvoted 7 times

Geirebass Most Recent 2 months, 1 week ago

This should be Core(SQL) and Partition Key (item.id) i belive ?
upvoted 2 times

Question #47 *Topic 3*

You create and publish a new Azure App Service web app.

User authentication and authorization must use Azure Active Directory (Azure AD).

You need to configure authentication and authorization.

What should you do first?

A. Add an identity provider.

B. Map an existing custom DNS name.

C. Create and configure a new app setting.

D. Add a private certificate.

E. Create and configure a managed identity.

Correct Answer: *A*

Community vote distribution

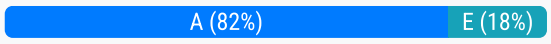

Firo Highly Voted 3 months ago

Selected Answer: A

A- Add an identity provide is correct
https://learn.microsoft.com/en-us/azure/app-service/scenario-secure-app-authentication-app-service#3-configure-authentication-and-authorization

upvoted 6 times

Vladimir_Gajinov 2 months, 1 week ago

Correct - step 3

upvoted 1 times

huuphong91 Most Recent 2 weeks, 5 days ago

I think E is right

upvoted 1 times

nerojdn 1 month, 1 week ago

Selected Answer: E

Why not E. Create and configure a managed identity??

upvoted 2 times

cldmcp 2 weeks, 4 days ago

I think it is because you need to implement user authentication and authorization. Managed identity is used to authenticate Azure resources

upvoted 2 times

lyggwtwtczxnhznebw 3 months ago

Selected Answer: A

Answer seems correct based on https://learn.microsoft.com/en-us/azure/app-service/configure-authentication-provider-aad?tabs=workforce-tenant

upvoted 3 times

Question #48 *Topic 3*

DRAG DROP
-

You have an Azure Cosmos DB for NoSQL account.

You plan to develop two apps named App1 and App2 that will use the change feed functionality to track changes to containers. App1 will use the pull model and App2 will use the push model.

You need to choose the method to track the most recently processed change in App1 and App2.

Which component should you use? To answer, drag the appropriate components to the correct apps. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Components

- Lease container
- Integrated cache
- Continuation token

Answer Area

App	Component
App1	
App2	

Correct Answer:

Answer Area

App	Component
App1	Continuation token
App2	Lease container

oCaio Highly Voted 2 months, 2 weeks ago

Seems correct.

App2 Push Model - Lease container
"When reading from the Azure Cosmos DB change feed, we usually recommend using a push model because you won't need to worry about: ...
Storing state for the last processed change. If you are reading from the change feed processor, state is automatically stored in a lease container.
Ref:
https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/read-change-feed#reading-change-feed-with-a-push-model"

upvoted 5 times

thecrash Most Recent 3 weeks, 4 days ago

Correct!
The change feed pull model allows you to consume the change feed at your own pace. Changes must be requested by the client and there's no automatic polling for changes. If you want to permanently "bookmark" the last processed change (similar to the push model's lease container), you'll need to save a continuation token.

https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/read-change-feed

upvoted 2 times

Ayaz_b 1 month ago

Correct

upvoted 1 times

mario_no_code 3 months ago

The answer is correct. Link: https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-pull-model?tabs=dotnet#compare-to-the-change-feed-processor

upvoted 3 times

Topic 4 - Question Set 4

Question #1 *Topic 4*

You are developing a Java application that uses Cassandra to store key and value data. You plan to use a new Azure Cosmos DB resource and the Cassandra
API in the application. You create an Azure Active Directory (Azure AD) group named Cosmos DB Creators to enable provisioning of Azure Cosmos accounts, databases, and containers.
The Azure AD group must not be able to access the keys that are required to access the data.
You need to restrict access to the Azure AD group.
Which role-based access control should you use?

A. DocumentDB Accounts Contributor

B. Cosmos Backup Operator

C. Cosmos DB Operator

D. Cosmos DB Account Reader

Correct Answer: *C*
Azure Cosmos DB now provides a new RBAC role, Cosmos DB Operator. This new role lets you provision Azure Cosmos accounts, databases, and containers, but can't access the keys that are required to access the data. This role is intended for use in scenarios where the ability to grant access to Azure Active Directory service principals to manage deployment operations for Cosmos DB is needed, including the account, database, and containers.
Reference:
https://azure.microsoft.com/en-us/updates/azure-cosmos-db-operator-role-for-role-based-access-control-rbac-is-now-available/

Community vote distribution
C (100%)

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: C

DocumentDB Account Contributor: Can manage Azure Cosmos DB accounts.

Cosmos Backup Operator: Can submit a restore request for Azure portal for a periodic backup enabled database or a container. Can modify the backup interval and retention on the Azure portal. Cannot access any data or use Data Explorer.

Cosmos DB Operator: Can provision Azure Cosmos accounts, databases, and containers. Cannot access any data or use Data Explorer.

Cosmos DB Account Reader: Can read Azure Cosmos DB account data.
CosmosRestoreOperator: Can perform restore action for Azure Cosmos DB account with continuous backup mode.

Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/role-based-access-control
upvoted 54 times

OPT_001122 1 year, 2 months ago
very good and helpful explanation
upvoted 2 times

RaviKS Highly Voted 2 years, 10 months ago
Answer is correct
upvoted 38 times

richardc7 Most Recent 1 month ago
got 28/09/2023
upvoted 1 times

richardc7 1 month ago
got on 28/09/2023
upvoted 1 times

mihailos 2 months ago
Got it in exam 28/08/23. Went with proposed answer. Scored 912
upvoted 2 times

basquiat 2 months, 2 weeks ago

Got this one in 2023-08-08

upvoted 3 times

JH81 4 months ago

Selected Answer: C

Got this on 6/28/2023 and passed with 850. Went with answer.

upvoted 3 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

My case:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 1 times

Esward 9 months, 1 week ago

Given answer is correct

https://azure.microsoft.com/en-us/updates/azure-cosmos-db-operator-role-for-role-based-access-control-rbac-is-now-available/

upvoted 1 times

kkop 11 months ago

DocumentDB Account Contributor Can manage Azure Cosmos DB accounts.
Cosmos DB Account Reader Can read Azure Cosmos DB account data.

CosmosBackupOperator Can submit a restore request in the Azure portal for a periodic backup enabled database or a container. Can modify the backup interval and retention in the Azure portal. Cannot access any data or use Data Explorer.

CosmosRestoreOperator Can perform a restore action for an Azure Cosmos DB account with continuous backup mode.

Cosmos DB Operator Can provision Azure Cosmos DB accounts, databases, and containers. Cannot access any data or use Data Explorer.

So correct C

upvoted 1 times

capable 11 months ago

got in exam 27/11/2022

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: C

C. Cosmos DB Operator

upvoted 1 times

DP_Bhatt 1 year ago

Selected Answer: C

Correct

upvoted 1 times

elequiel 1 year ago

Got it in exam 20/10/2022

upvoted 2 times

OPT_001122 12 months ago

Thanks for mentioning date

upvoted 1 times

[Removed] 1 year, 1 month ago

I have to also remember the pricing tiers, and also the RBAC roles and permissions? What kind of developers want to remember information like this? I do not even know these things in things I actually use because it is so easy to check when I actually set it up (and I do that once or twice in a lifetime). Microsoft doesn't get how to ask questions of relevance. They ask based on what they know at that point of time. I can ask them a 1000 questions they will not know.

upvoted 7 times

Perplex 1 year, 3 months ago

Selected Answer: C

C is correct, for an explanation refer to the comment by mlantonis.

upvoted 2 times

Question #2 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.
You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.
You need to configure authorization.
Solution: Configure the Azure Web App for the website to allow only authenticated requests and require Azure AD log on.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Instead in the Azure AD application's manifest, set value of the groupMembershipClaims option to All.

Reference:

https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/

Community vote distribution

B (75%) A (25%)

fadikh Highly Voted 2 years, 8 months ago

Answer is correct. The proposed solution lacks the authorization part.

upvoted 32 times

ferut 2 years, 5 months ago

Aggree. Setting up the access on AD is one thing. The application should be set up to differentiate different roles and access.

upvoted 4 times

edengoforit 1 year, 9 months ago

Answer is NO

upvoted 3 times

kondapaturi Highly Voted 2 years, 4 months ago

Answer – No, Here you need to create an application in Azure AD. Then set the groupMembershipClaims claims. Then inspect the token in the application to see if the user is part of that group.
Hence - B is correct

upvoted 13 times

Dixavado Most Recent 1 month ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 2 times

Tarajee 1 month, 2 weeks ago

On my exam 2023sept

upvoted 1 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 4 times

sarmaria 7 months, 2 weeks ago

Got this on 16/03/23. Chosen No. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

Esward 9 months, 1 week ago

Given answer is correct
https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups

upvoted 1 times

capable 11 months ago

Got this case study in my exam 27/11/2022

upvoted 2 times

gmishra88 1 year ago

If you think about it having the groups claim is also not enough. You need to check for it. So, even if the group claims was there it could be no. Because of that reason it could have been yes, because you need to do that step as well as groups claim and also check if the claim is authorized. These kind of questions are just Microsoft way of confusion

upvoted 2 times

bhakk 1 year, 4 months ago

Answer is B

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer

upvoted 1 times

pradipbobhate 1 year, 7 months ago

Selected Answer: A

correct Answer

upvoted 1 times

Freidrich 1 year, 8 months ago

Selected Answer: B

The correct answer is B: No.

upvoted 1 times

edengoforit 1 year, 9 months ago

This is the correct answer:
Solution:
☞ Create a new Azure AD application. In the application's manifest, set value of the groupMembershipClaims option to All.
☞ In the website, use the value of the groups claim from the JWT for the user to determine permissions.

upvoted 5 times

edengoforit 1 year, 9 months ago

Provided answer and explanation is correct.

upvoted 1 times

ehurfheiz 1 year, 9 months ago

Selected Answer: B

B seems to be the correct answer

upvoted 1 times

tramlong888 1 year, 10 months ago

Correct Answer is No

upvoted 1 times

Question #3 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.
You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.
You need to configure authorization.
Solution:
☞ Create a new Azure AD application. In the application's manifest, set value of the groupMembershipClaims option to All.
☞ In the website, use the value of the groups claim from the JWT for the user to determine permissions.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*
To configure Manifest to include Group Claims in Auth Token
1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:
2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.
3. Locate the ג€groupMembershipClaimsג€ setting. Set its value to either ג€SecurityGroupג€ or ג€Allג€. To help you decide which:
☞ ג€SecurityGroupג€ - groups claim will contain the identifiers of all security groups of which the user is a member.
☞ ג€Allג€ - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member
Now your application will include group claims in your manifest and you can use this fact in your code.
Reference:
https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/

Community vote distribution
A (100%)

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: Yes

Reference:
https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups
upvoted 44 times

kondapaturi Highly Voted 2 years, 4 months ago
Answer – Yes, this is the correct approach. The Application manifest has a setting for groupMembershipClaims. By setting this to all, the Azure AD groups that the user belongs to will be returned as part of the claims in the JWT token.
upvoted 9 times

Dixavado Most Recent 1 month ago
Selected Answer: A
It was on my exam today (2023-09-26) I went with the examtopics answer - score 850
upvoted 1 times

Tarajee 1 month, 2 weeks ago
On my exam 2023sept
upvoted 1 times

NightshadeRC 3 months, 1 week ago
Had this question in today's exam: 2023-07-26
upvoted 2 times

MysticalSam 4 months, 3 weeks ago
This question was in today's exam on 10-June-2023
upvoted 2 times

sarmaria 7 months, 2 weeks ago
Got this on 16/03/23. Chosen yes. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

BrettusMaximus 8 months, 2 weeks ago

B
Sure, A would technically work but not the best solution as unauthorized users can hit the website.
Best practice is to use groupMembershipClaims in the App Registration to restrict access at the App level.
upvoted 2 times

Esward 9 months, 1 week ago

Given answer A is correct as per MS docs

https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups
upvoted 2 times

serpevi 1 year, 1 month ago

Got this in 09/22 , went just with this as YES, score 927.
upvoted 4 times

Eltooth 1 year, 4 months ago

Selected Answer: A

A is correct answer.
upvoted 3 times

PieroFranco 1 year, 7 months ago

Selected Answer: A

The answer is correct. If you do not add the groupMembershipClaims to the manifest you wont see any AD group in the auth token.
upvoted 2 times

Freidrich 1 year, 8 months ago

Selected Answer: A

The answer is correct.
upvoted 1 times

sujitwarrier11 2 years, 4 months ago

I think the answer is no. group claims are at Azure Ad tenant level. The question mentions that the authorization should be for the app. So roles would be the better option here.
upvoted 1 times

ZodiaC 2 years, 3 months ago

thats not true, look link plz: https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups
upvoted 1 times

UnknowMan 2 years, 5 months ago

Answer is correct
upvoted 1 times

glam 2 years, 5 months ago

A. Yes
upvoted 1 times

clarionprogrammer 2 years, 6 months ago

A. Yes
Note: Azure AD caps at 200 the number of groups that can be sent via JWT format.
https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups
upvoted 5 times

Question #4 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.
You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.
You need to configure authorization.
Solution:
☞ Create a new Azure AD application. In the application's manifest, define application roles that match the required permission levels for the application.
☞ Assign the appropriate Azure AD group to each role. In the website, use the value of the roles claim from the JWT for the user to determine permissions.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
To configure Manifest to include Group Claims in Auth Token
1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:
2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.
3. Locate the ز€groupMembershipClaimsز€ setting. Set its value to either ز€SecurityGroupز€ or ز€Allز€. To help you decide which:
☞ ز€SecurityGroupز€ - groups claim will contain the identifiers of all security groups of which the user is a member.
☞ ز€Allز€ - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member
Now your application will include group claims in your manifest and you can use this fact in your code.
Reference:
https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/

Community vote distribution
A (58%) B (42%)

[Removed] Highly Voted 2 years, 9 months ago
I agree that this solution should work as well. The roles get assigned by AD groups, so the requirement "A user's Azure AD group membership must be used to determine the permission level" is met.

This solution should be answered with "yes".

This scenario has 2 solutions provided as the approach using the "groupMembershipClaims" is possible as well.
That's OK as it says "Some question sets might have more than one correct solution, while others might not have a correct solution."
upvoted 56 times

Mal22002 1 year, 7 months ago
There are many solutions when you're writing the code but the question is whether provided solution works or not and it will not work. The second part of the solution failed the requirement "A user's Azure AD group membership must be used to determine the permission level." as it's saying "use the value of roles claim".
The answer is No.
upvoted 8 times

Adiu 8 months, 3 weeks ago
But AD groups are used to assign correct roles to them. So the roles claim is based on user’s AD groups membership.
upvoted 1 times

Adrian1405 Highly Voted 2 years, 10 months ago
This should be the correct answer: https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps
upvoted 18 times

ALittleBunny 2 years, 10 months ago
You're totally right. It's just managing permissions in two different ways, groupMembershipClaims and assigning the appropriate Azure AD group to each role, which are both correct.
upvoted 6 times

iiiihhhh 2 years, 10 months ago

I think that without the "groupMembershipClaims" setting you will not be able to obtain group membership information and so the proposal doesn't meet the goal.

upvoted 6 times

bugimachi 2 years, 10 months ago

I am not too sure about this. "groupMembershipClaims" will include AAD groups' IDs in the JWT "group" array, which is one way to go. The other way would be to add a app role in the manifest (which is not necessary for the first approach!) and then add groups to this role. This will include the app roles (but not the group IDs) in the JWT's "roles" array (but not in the "groups" array!).

I liked this blog post to understand the difference: https://joonasw.net/view/using-groups-vs-using-app-roles-in-azure-ad-apps

upvoted 12 times

clarionprogrammer 2 years, 6 months ago

Excellent post! A. Yes.
Final answer.

upvoted 6 times

FatSeal 1 year, 2 months ago

A user's Azure AD group membership must be used to determine the permission level!!!!!!!!!!!

upvoted 2 times

warchoon 8 months ago

You can use gropes for roles
https://learn.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps#:~:text=In%20the%20Create%20app%20role%20pane%2C%20enter%20the%20settings%20for%20the%20role.

upvoted 1 times

ahadjithoma 2 years, 10 months ago

Requirement: A user's Azure AD group membership must be used to determine the permission level.
Does that solution meets the above requirement?

upvoted 3 times

bugimachi 2 years, 10 months ago

I think so; you can add groups to app roles defined in manifest.

upvoted 4 times

rdemontis 2 years, 7 months ago

I think you can't instead, because AppRoles are intended for RBAC access control and not AAD group membership as required in the question. See this for more info: https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps

upvoted 2 times

Dixavado Most Recent 1 month ago

Selected Answer: B

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 1 times

ayush66203 1 week, 3 days ago

These examtopics questions were sufficient for you to score 850? or you prepared from other sources too?

upvoted 1 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 2 times

Zoriark 6 months, 4 weeks ago

ChatGPT
A. Yes, the solution meets the goal.

The solution correctly suggests creating a new Azure AD application and defining application roles that match the required permission levels. By assigning the appropriate Azure AD groups to each role, the application can determine a user's permission level based on their group membership. The roles claim from the JWT token can then be used in the website to determine the user's permissions. This approach allows for granular access control and ensures that users are assigned the correct permission level based on their group membership.

upvoted 2 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23

upvoted 1 times

Jeeaak 8 months ago

does the answer is yes

upvoted 2 times

Esward 9 months, 1 week ago

yes is the correct answer

https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps
upvoted 1 times

andrvelich 1 year, 6 months ago

Selected Answer: A

should work as well
upvoted 1 times

Azprep 1 year, 6 months ago

Answer should be yes
upvoted 1 times

Mal22002 1 year, 7 months ago

Selected Answer: B

The answer is No, the explanation is useless though.
Requirement: A user's Azure AD group membership must be used to determine the permission level.
The solution that failed the requirement: Assign the appropriate Azure AD group to each role. In the website, use the value of the roles claim from the JWT for the user to determine permissions.
Azure AD group membership can be accessed in groups claim not in roles claim.

Reference: https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps
upvoted 4 times

mmdex 8 months, 2 weeks ago

Read carefully: "A user's Azure AD group membership must be used to determine the permission level."

You are not trying to determine user's group membership. You are trying to determine their permission level. You don't need groups claim, roles is enough.
upvoted 1 times

Freidrich 1 year, 8 months ago

Selected Answer: A

I think the correct answer is A: Yes.
upvoted 3 times

idroj 1 year, 8 months ago

Selected Answer: A

The correct answer is yes
upvoted 3 times

Elemata 1 year, 8 months ago

Yes. I agree with "melli": The roles get assigned by AD groups, so the requirement "A user's Azure AD group membership must be used to determine the permission level" is met.
This scenario has 2 solutions provided as the approach using the "groupMembershipClaims" is possible as well.
upvoted 1 times

Saurabh_Kulkarni 2 years, 1 month ago

'No' is correct answer. Application roles are assigned which are provided in 'app claims' and get removed if application get deleted. But in the requirement, it is mentioned that 'group membership' must be used, means 'group claims' must be done, which doesn't changes as per application lifecycle. Even though the solution will work but the requirement is different. (this type of requirement is necessary when there will be multiple apps in future).
upvoted 3 times

RajMasilamani 2 years, 1 month ago

The answer is Yes.
upvoted 3 times

Kalaisuran 2 years, 3 months ago

As per the explaination in the reference link application role base authentication for the AD group also possible
Reference Link : https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-fed-group-claims
upvoted 1 times

sujitwarrier11 2 years, 4 months ago

The answer should be yes. They have mentioned that the claims are for the website. group claims are at Azure AD level while role claims are at app level. since they have mentioned that the authrizations is for the website. this is the right answer and not group claims.

https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps
upvoted 3 times

dreamcoder 1 year, 10 months ago

i agree that this yes is the correct answer. Group claims are not enough. The app roles need to be assigned to appropriate AD groups from "Enterprise Apps".
upvoted 1 times

Question #5 *Topic 4*

DRAG DROP -

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

The application must allow recovery of an accidental deletion of the key vault or key vault objects. Key vault objects must be retained for 90 days after deletion.

You need to protect the key vault and key vault objects.

Which Azure Key Vault feature should you use? To answer, drag the appropriate features to the correct actions. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Features

- Access policy
- Purge protection
- Soft delete
- Shared access signature

Answer Area

Action	Feature
Enable retention period and accidental deletion.	Feature
Enforce retention period and accidental deletion.	Feature

Correct Answer:

Features

- Access policy
- Purge protection
- Soft delete
- Shared access signature

Answer Area

Action	Feature
Enable retention period and accidental deletion.	Soft delete
Enforce retention period and accidental deletion.	Purge protection

Box 1: Soft delete -

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

Box 2: Purge protection -

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

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

Reference:

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

andsol Highly Voted 2 years, 7 months ago

The answer is correct.

upvoted 56 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: Soft delete

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

This can be achieved with the help of the soft-delete feature of the key vault.

Box 2: Purge protection

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

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

This can be achieved with the help of the purge protection feature of the key vault.

Reference:

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

upvoted 38 times

minh1302 1 year, 3 months ago

Thank mlantonis, learn from you a lot (AZ-104, AZ-204)

upvoted 3 times

leviatas Most Recent 1 month ago

Got this question in exam. 2023.09.25

upvoted 1 times

applepie 3 months ago

got this question today, go with the provided answer - 7/30/2023, score 895/1000

upvoted 3 times

NightshadeRC 3 months, 1 week ago

Had this question in today's exam: 2023-07-26

upvoted 2 times

murat12345 3 months, 1 week ago

If Box 1 is about having a retention period (meaning: you cant hard delete before the 90 days pass), what additional advantage does box 2 / purge protection give?

upvoted 1 times

basquiat 2 months, 2 weeks ago

The first is related to key vault secrets and the second to the keyvault itself, from what I understood.

upvoted 1 times

Juhi_Az204 7 months, 1 week ago

Got this on 26March,2023

upvoted 1 times

Esward 9 months, 1 week ago

Given answers are correct
https://learn.microsoft.com/en-us/azure/key-vault/general/soft-delete-overview

upvoted 1 times

Lucifer14 10 months ago

In 28-12-2022 exam

upvoted 1 times

xRiot007 1 year, 4 months ago

You use Soft delete to *enable* and Purge protection to *enforce*

upvoted 3 times

Eltooth 1 year, 4 months ago

Soft delete
Purge protection

upvoted 1 times

Rini100 1 year, 6 months ago

Got this on 20 Apr 2022

upvoted 3 times

meoukg 1 year, 7 months ago

ot it on 03/2022, I chose Soft delete => Purge protection

upvoted 4 times

francis6170 2 years, 4 months ago

got this in the exam :)

upvoted 7 times

UnknowMan 2 years, 5 months ago

The answer is correct.

upvoted 2 times

glam 2 years, 5 months ago

correct

upvoted 3 times

Question #6 *Topic 4*

You provide an Azure API Management managed web service to clients. The back-end web service implements HTTP Strict Transport Security (HSTS).
Every request to the backend service must include a valid HTTP authorization header.
You need to configure the Azure API Management instance with an authentication policy.
Which two policies can you use? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Basic Authentication

B. Digest Authentication

C. Certificate Authentication

D. OAuth Client Credential Grant

Correct Answer: *CD*

Community vote distribution

AC (56%) AD (44%)

Nielson Highly Voted 2 years, 8 months ago

A, C
https://www.examtopics.com/exams/microsoft/az-203/view/15/
https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies
upvoted 93 times

BrettusMaximus 2 years, 6 months ago

Must be A,C As the API documentation only allows 3 options. It states: >>>>
Authentication policies
Authenticate with Basic - Authenticate with a backend service using Basic authentication.
Authenticate with client certificate - Authenticate with a backend service using client certificates.
Authenticate with managed identity - Authenticate with the managed identity for the API Management service.
upvoted 21 times

BrettusMaximus 2 years, 6 months ago

It can't be D as we need to authenticate to the backend and not from client.
upvoted 7 times

fadikh 2 years, 8 months ago

but client certificate does not use the authorization header
upvoted 6 times

vb3d 2 years, 7 months ago

Client certificate does use the x-arr-clientcert header
As Kitkit pointed out below, check this link
https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies.
upvoted 10 times

vb3d 2 years, 7 months ago

I think this means correct answer is A and D
upvoted 7 times

hobob Highly Voted 2 years, 7 months ago

A, D
The question states that the solution MUST use the HTTP authorization header. The only two options that do are Basic Authentication (where the header would be "Authorization: Basic <auth-base64-encoded-string> or "Bearer <bearer-token-string>).
upvoted 38 times

pmsiva 2 years, 7 months ago

AD is correct. Certificate authentication will not send http auth header
https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#Basic
upvoted 6 times

syfool 2 years, 1 month ago

D is not the answer, as that is not even an AUTHENTICATION POLICY!!!

The link below lists 3 authentication policies:
https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies

upvoted 3 times

MiraA 2 years, 1 month ago

It seems "OAuth Client Credential Grant" exists and could be considered as the specific type of "Authenticate with managed identity" which sends the authorization header?
https://oauth.net/2/grant-types/client-credentials/
https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies

upvoted 4 times

ENGs Most Recent 1 month ago

Selected Answer: AD

OAuth Client Credentials Grant (D), when used, will use an Authorization header being sent containing a bearer token.
Basic Authentication (A) will send an Authorization header.
--
Client Certificate (C) could be passed through an HTTP header but not necessarily the Authorization header.

upvoted 1 times

ck_malviya 1 month, 2 weeks ago

Selected Answer: AD

AD. The solution must use Authorization header
Basic auth and OAUTH user Bearer token auth header.

upvoted 1 times

Vladimir_Gajinov 2 months ago

Selected Answer: AD

I would go with A and D.
A: allows you to require clients to include a valid username and password in the HTTP authorization header. While it's a simple method, it might not be the most secure option for modern applications, as the credentials are sent with each request in base64-encoded format.
D: OAuth Client Credential Grant is a flow where a client (in this case, your API Management instance) uses its own credentials (client ID and client secret) to obtain an access token from an OAuth authorization server. This access token can be included in the authorization header of requests to the backend service. This method is more secure than Basic Authentication and provides better control over authentication and authorization.

B: is a challenge-response mechanism that involves a server-provided nonce. It doesn't directly match the requirement for sending an HTTP authorization header with each request.
C: Certificate Authentication involves using client certificates to authenticate the client to the server. While this can be secure, it's not the same as sending an HTTP authorization header.

upvoted 1 times

applepie 3 months ago

got this question today, answer C,D without confidence - 7/30/2023, score 895/1000

upvoted 1 times

uncledana 3 months, 2 weeks ago

I believe it is:

A - Sets the HTTP Authorization header to a corresponding value in the policy request
D - Authenticates to the backend with a valid HTTP authorization header

https://learn.microsoft.com/en-us/azure/api-management/authentication-managed-identity-policy

upvoted 1 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 1 times

minYoongi 6 months, 2 weeks ago

AD is correct. https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad#configure-a-jwt-validation-policy-to-pre-authorize-requests

The following example policy, when added to the <inbound> policy section, checks the value of the audience claim in an access token obtained from Azure AD that is presented in the Authorization header. Straight from the documentation

upvoted 2 times

Zoriark 6 months, 4 weeks ago

ChatGPT
A. Basic Authentication and D. OAuth Client Credential Grant are the two policies that can be used.

Basic Authentication policy can be used to require an HTTP Basic authentication header with each request. This policy requires a username and password to be sent in the Authorization header with each request.

OAuth Client Credential Grant policy can be used to require an OAuth access token with each request. This policy is useful when calling a web service that requires authentication and you want to validate the identity of the client making the request.

Digest Authentication policy cannot be used as it is not supported by the Azure API Management instance.

Certificate Authentication policy can be used to require a client certificate with each request. However, this policy is not applicable in this scenario as the back-end web service does not implement certificate-based authentication.

upvoted 2 times

rcardenas72 8 months, 1 week ago

to configure the Azure API Management instance with an authentication policy to enforce the requirement that every request to the backend service includes a valid HTTP authorization header, you can use the following two policies:

A. Basic Authentication: This policy allows you to require that clients include a valid HTTP authorization header that contains a username and password. This policy can be used to authenticate users against a user store, such as Azure AD, or an external system.

C. Certificate Authentication: This policy allows you to require that clients present a valid client certificate as part of their request. This policy can be used to authenticate clients that possess a valid client certificate signed by a trusted authority.

Both of these policies can be used to enforce the requirement that every request to the backend service includes a valid HTTP authorization header. The choice between them would depend on the specific requirements and constraints of the scenario

upvoted 1 times

baroo1 8 months, 2 weeks ago

Selected Answer: AD

A + D;
A: as clear as https://learn.microsoft.com/en-us/azure/api-management/authentication-basic-policy
D: EAI can use client credentials grant type that use OAuth 2.0. Which grants an authorization header.
Source: https://learn.microsoft.com/en-us/azure/api-management/authorizations-overview.

upvoted 1 times

Esward 9 months, 1 week ago

A,C are the correct answers

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

upvoted 1 times

rolling_potato_ 10 months ago

Was on my exam today (03-01-2023)

upvoted 3 times

hubekpeter 11 months, 1 week ago

Selected Answer: AC

A,C no brainer.

upvoted 1 times

coffecold 1 year ago

C and D

- "The back-end web service implements HTTP Strict Transport Security (HSTS)".
This ensures you that HTTPS must be used all time.Thus a certificate is needed (C)
see : https://en.wikipedia.org/wiki/HTTP_Strict_Transport_Security
- "Every request to the backend service must include a valid HTTP authorization header."
This ensures you that anyhow the level of authorization must be determined by the backend using a header.
Here OAuth 2.0 must be applied to get the token.

A and B, Basic Authentication and Digest Authentication use HTTP.

upvoted 3 times

tarek0811 3 months ago

i believe this is the correct answer

upvoted 1 times

deva2022 1 year, 2 months ago

Basic and Certificate

upvoted 2 times

Question #7 *Topic 4*

DRAG DROP -

You are developing an ASP.NET Core website that can be used to manage photographs which are stored in Azure Blob Storage containers.

Users of the website authenticate by using their Azure Active Directory (Azure AD) credentials.

You implement role-based access control (RBAC) role permissions on the containers that store photographs. You assign users to RBAC roles.

You need to configure the website's Azure AD Application so that user's permissions can be used with the Azure Blob containers.

How should you configure the application? To answer, drag the appropriate setting to the correct location. Each setting can be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

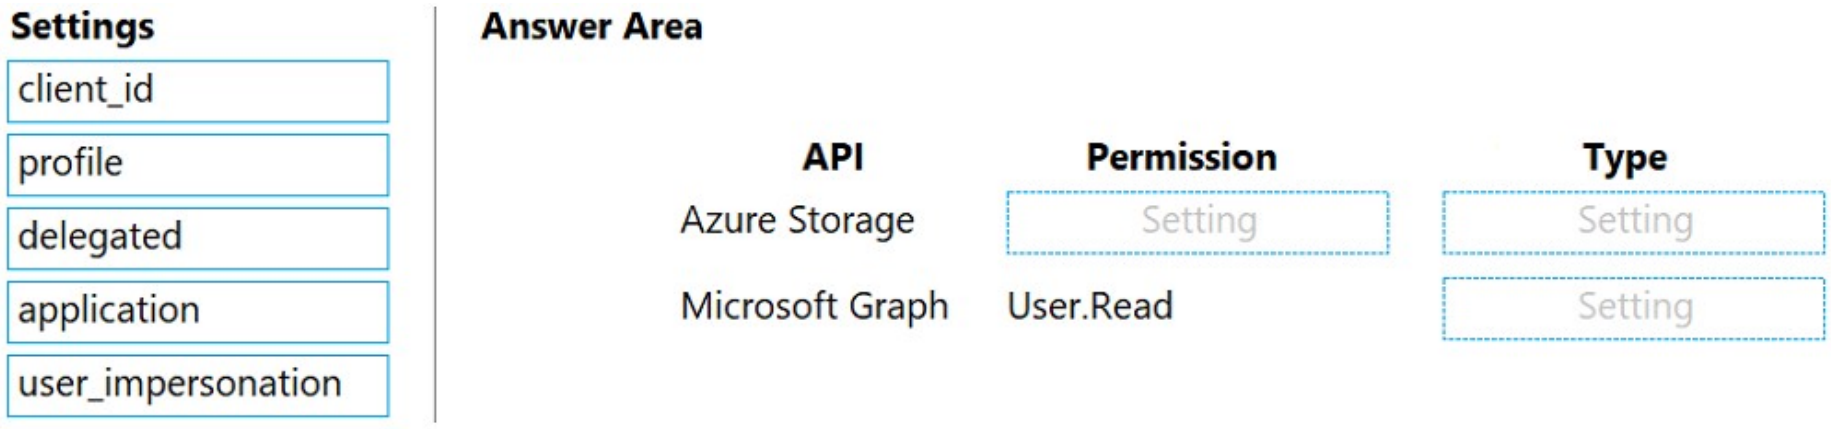

Correct Answer:

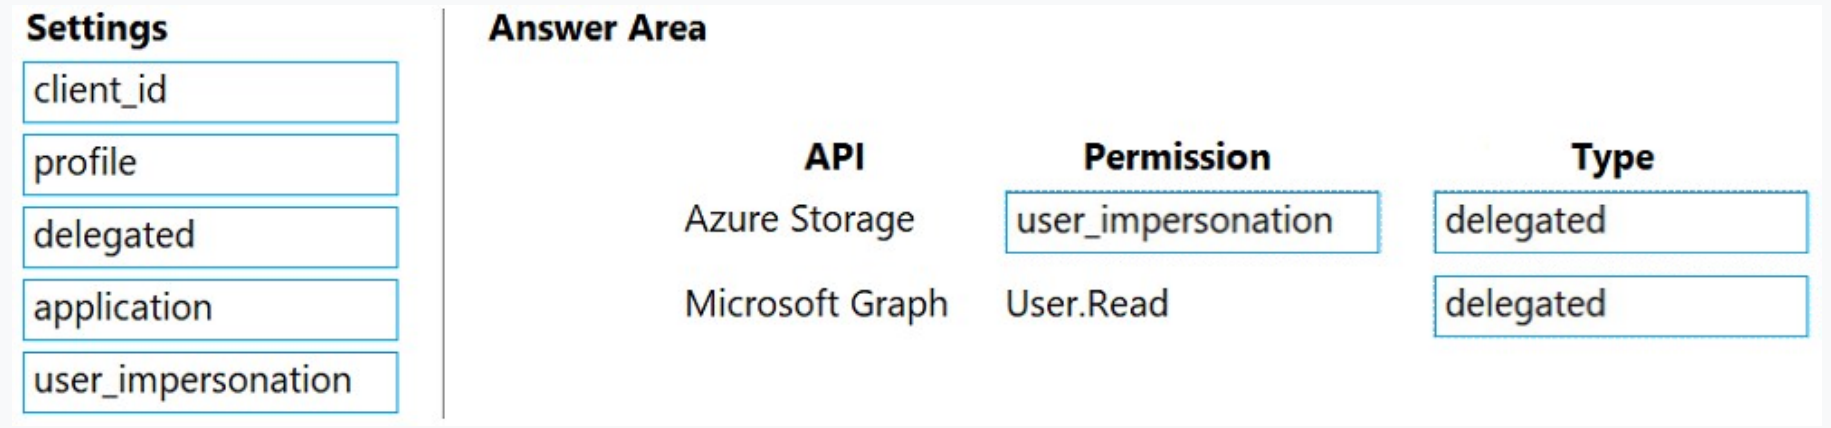

Box 1: user_impersonation -

Box 2: delegated -

Example:

1. Select the API permissions section
2. Click the Add a permission button and then:

Ensure that the My APIs tab is selected

3. In the list of APIs, select the API TodoListService-aspnetcore.
4. In the Delegated permissions section, ensure that the right permissions are checked: user_impersonation.
5. Select the Add permissions button.

Box 3: delegated -

Example -

1. Select the API permissions section
2. Click the Add a permission button and then,

Ensure that the Microsoft APIs tab is selected

3. In the Commonly used Microsoft APIs section, click on Microsoft Graph
4. In the Delegated permissions section, ensure that the right permissions are checked: User.Read. Use the search box if necessary.
5. Select the Add permissions button

Reference:

https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect-aspnetcore/calling-a-web-api-in-an-aspnet-core- web-application-using-azure-ad/

aswqe338 Highly Voted 2 years, 7 months ago

The given answer is correct.

upvoted 45 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000

upvoted 10 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: user_impersonation
The built-in user_impersonation scope indicates that the token is being requested on behalf of the user. Azure Storage exposes a single delegation scope named user_impersonation that permits applications to take any action allowed by the user.

Box 2: delegated

Box 3: delegated

Reference:
https://stackoverflow.com/questions/31404128/azure-ad-app-application-permissions-vs-delegated-permissions
https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-permissions-and-consent
https://docs.microsoft.com/en-us/azure/storage/common/storage-auth-aad-app?tabs=dotnet
https://docs.microsoft.com/en-us/rest/api/storageservices/authorize-with-azure-active-directory

upvoted 33 times

p2006 Most Recent 1 month, 2 weeks ago

https://learn.microsoft.com/en-us/information-protection/develop/concept-api-permissions#delegated-permissions

upvoted 1 times

Dianahu 3 months ago

https://www.andrew-best.com/posts/please-sir-can-i-have-some-auth/ first image

upvoted 1 times

Dats1987 6 months, 2 weeks ago

That's Correct.

upvoted 1 times

Priya0703 8 months, 1 week ago

On exam today 20-02-2023

upvoted 1 times

Esward 9 months, 1 week ago

Given answers are correct

https://stackoverflow.com/questions/31404128/azure-ad-app-application-permissions-vs-delegated-permissions

upvoted 1 times

serpevi 1 year, 1 month ago

Got this in 09/22 , went with the most voted answers, score 927.

upvoted 6 times

Eltooth 1 year, 4 months ago

User
Delegated
Delegated

upvoted 3 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 3 times

monaindia 1 year, 4 months ago

what answer you have chosen

upvoted 1 times

edengoforit 1 year, 11 months ago

Application Permissions: Your application needs to access the web API directly as itself (no user context). This type of permission requires administrator consent and is also not available for native client applications.
Delegation Permissions: Your application needs to access the web API as the signed-in user, but with access limited by the selected permission. This type of permission can be granted by a user unless the permission is configured as requiring administrator consent.

upvoted 4 times

lugospod 1 year, 10 months ago

Nope, it explicitly states that user is given RBAC permissions, and that each users permissions need to be used to access storage. So Application is not an option.

upvoted 1 times

jungaster 2 years, 5 months ago

the answer is correct.

upvoted 1 times

roybunt 2 years, 5 months ago

The answer is correct. For granting registered app permissions to Azure Storage, find in reference: https://docs.microsoft.com/en-us/azure/storage/common/storage-auth-aad-app?tabs=dotnet#grant-your-registered-app-permissions-to-azure-storage

upvoted 4 times

glam 2 years, 5 months ago

correct

upvoted 2 times

kwaazaar 2 years, 6 months ago

why Graph API?

upvoted 4 times

Sachini 1 year, 8 months ago

"Permissions are granted to Microsoft Graph automatically when you first register your app with Azure AD."

upvoted 2 times

Question #8 *Topic 4*

HOTSPOT -

You are developing an ASP.NET Core app that includes feature flags which are managed by Azure App Configuration. You create an Azure App Configuration store named AppFeatureFlagStore that contains a feature flag named Export.

You need to update the app to meet the following requirements:

- Use the Export feature in the app without requiring a restart of the app.
- Validate users before users are allowed access to secure resources.
- Permit users to access secure resources.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
public void Configure(IApplicationBuilder app, IWebHostEnvironment env)
{
if (env.IsDevelopment())
{
   app.UseDeveloperExceptionPage();
}
else
{
   app.UseExceptionHandler("/Error");
}
app. [ ] ();
app. [ ] ();
app. [ ] ();
app.UseEndpoint(endpoints =>
{
   endpoints.MapRazorPages();
});
}
```

app. ( ) ;
- UseAuthentication
- UseStaticFiles
- UseSession
- UseCookiePolicy

app. ( ) ;
- UseAuthorization
- UseHttpsRedirection
- UseSession
- UseCookiePolicy

app. ( ) ;
- UseAzureAppConfiguration
- UseRequestLocalization
- UseCors
- UseStaticFiles

Correct Answer:

Answer Area

```
public void Configure(IApplicationBuilder app, IWebHostEnvironment env)
{
if (env.IsDevelopment())
{
   app.UseDeveloperExceptionPage();
{
else
{
   app.UseExceptionHandler("/Error");
}
```

app. [ ] ();

- UseAuthentication
- UseStaticFiles
- UseSession
- UseCookiePolicy

app. [ ] ();

- UseAuthorization
- UseHttpsRedirection
- UseSession
- UseCookiePolicy

app. [ ] ();

- UseAzureAppConfiguration
- UseRequestLocalization
- UseCors
- UseStaticFiles

```
app.UseEndpoint(endpoints =>
{
   endpoints.MapRazorPages();
});
}
```

Box 1: UseAuthentication -

Need to validate users before users are allowed access to secure resources.

UseAuthentication adds the AuthenticationMiddleware to the specified IApplicationBuilder, which enables authentication capabilities.

Box 2: UseAuthorization -

Need to permit users to access secure resources.

UseAuthorization adds the AuthorizationMiddleware to the specified IApplicationBuilder, which enables authorization capabilities.

Box 3: UseStaticFiles -

Need to use the Export feature in the app without requiring a restart of the app.

UseStaticFiles enables static file serving for the current request path

Reference:

https://docs.microsoft.com/en-us/dotnet/api/microsoft.aspnetcore.builder.iapplicationbuilder?view=aspnetcore-5.0

hobob Highly Voted 2 years, 7 months ago

UseAuthentication, Use Authorization, UseAzureAppConfiguration.
The last option allows dynamic configuration to be reloaded without an app restart.
https://docs.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core5x

upvoted 151 times

Pooochie 2 years, 7 months ago

According to the site provided by you it should be
1. UseAzureAppConfiguration
2.UseAuthentication
3. Use Authorization

upvoted 12 times

xRiot007 1 year, 4 months ago

It is irrelevant in this case, the boxes only have ONE correct option each and those are, box1: authentication, box2: authorization, box3: app config.

upvoted 2 times

hobob 2 years, 7 months ago

And that's what I said. Just in a different order.

upvoted 8 times

OPT_001122 1 year, 2 months ago

i have doubt regarding UseAuthentication

UseAuthentication - this is not in the code snippet in the above link

public void Configure(IApplicationBuilder app, IWebHostEnvironment env)
{
.....
app.UseAzureAppConfiguration();

app.UseHttpsRedirection();

app.UseStaticFiles();

app.UseRouting();

app.UseAuthorization();
}

upvoted 1 times

avanthasiriwardana 2 years, 7 months ago

Order matters

upvoted 11 times

Kerrits 2 years, 6 months ago

The UseAzureAppConfiguration is not available in the first list of options, only the last list of options.
Also, for configuring the app the order only matters sometimes. It doesn't in this case.

upvoted 6 times

rdemontis 2 years, 7 months ago

correct, you're right!

upvoted 4 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: UseAuthentication
Need to validate users before users are allowed access to secure resources.
UseAuthentication adds the AuthenticationMiddleware to the specified IApplicationBuilder, which enables authentication capabilities. To verify users, we need to use the Authentication middleware.

Box 2: UseAuthorization
Need to permit users to access secure resources.
UseAuthorization adds the AuthorizationMiddleware to the specified IApplicationBuilder, which enables authorization capabilities. To authorize users to access resources , we need to use the Authorization middleware.

Box 3: UseAzureAppConfiguration
Adding the UseAzureAppConfiguration middleware to allow the configuration settings registered for refresh to be updated while the ASP.NET Core web app continues to receive requests. For using feature flags, you need to make use of the Azure App Configuration service. To ensure the configuration settings are refreshed without the need to restart the web app, you can use the middleware of UseAzureAppConfiguration();

upvoted 54 times

mlantonis 2 years, 5 months ago

Reference:

https://docs.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core5x

https://docs.microsoft.com/en-us/dotnet/api/microsoft.aspnetcore.builder.iapplicationbuilder?view=aspnetcore-5.0

upvoted 11 times

HimGarg Most Recent 3 weeks, 4 days ago

Why no one from ExamTopics bother to correct the answers?

upvoted 2 times

gmishra88 1 year ago

And Microsoft claims this exam does not require specific knowledge of .Net SDK.

upvoted 6 times

macobuzi 2 months, 1 week ago

I doubt it ever stated that

upvoted 1 times

tushar1812 1 year, 3 months ago

Drop down 1 UseAuthentication

Drop down 2 Use Authorization

Drop down 3 UseAzureAppConfiguration.

The last option allows dynamic configuration to be reloaded without an app restart.

Refer

https://docs.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core5x

upvoted 2 times

Eltooth 1 year, 4 months ago

Authentication
Authorisation
Appconfig

upvoted 2 times

tramlong888 1 year, 10 months ago

the correct answer should be UseAuthentication, Use Authorization, UseAzureAppConfiguration.

upvoted 2 times

doudou123 2 years, 1 month ago

according Udemy,
1. UseAuthentication
2.Use Authorization
3. UseAzureAppConfiguration

upvoted 5 times

burakcive 1 year, 4 months ago

Which course on udemy?

upvoted 1 times

kondapaturi 2 years, 4 months ago

app.UseAuthentication(); - To verify users ,we need to use the Authentication middleware.
app.UseAuthorization(); -To authorize users to access resources , we need to use the Authorization middleware.
app.UseAzureAppConfiguration(); -For using feature flags, you need to make use of the Azure App Configuration service. To ensure the configuration settings are refreshed without the need to restart the web app, you can use the middleware of UseAzureAppConfiguration();

upvoted 2 times

UnknowMan 2 years, 5 months ago

1.UseAuthentication, -> Login
2.Use Authorization, -> Authorization
3.UseAzureAppConfiguration. -> Dynamic App settings

upvoted 3 times

glam 2 years, 5 months ago

UseAuthentication, Use Authorization, UseAzureAppConfiguration.

upvoted 3 times

pavan555manjunath 2 years, 5 months ago

Any one please confirm the answer is correct
1.UseAuthentication,
2.Use Authorization,
3.UseAzureAppConfiguration.

upvoted 5 times

Question #9 *Topic 4*

You have an application that includes an Azure Web app and several Azure Function apps. Application secrets including connection strings and certificates are stored in Azure Key Vault.
Secrets must not be stored in the application or application runtime environment. Changes to Azure Active Directory (Azure AD) must be minimized.
You need to design the approach to loading application secrets.
What should you do?

A. Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity.

B. Create a single Azure AD Service Principal with permission to access Key Vault and use a client secret from within the App Services to access Key Vault.

C. Create a system assigned Managed Identity in each App Service with permission to access Key Vault.

D. Create an Azure AD Service Principal with Permissions to access Key Vault for each App Service and use a certificate from within the App Services to access Key Vault.

Correct Answer: *C*

Use Key Vault references for App Service and Azure Functions.

Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.

Reference:

https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

Community vote distribution

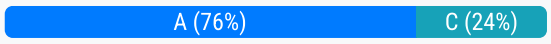

AssilAbdulrahim Highly Voted 2 years, 7 months ago

I think the answer is A. Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity.
Because we have more than one App (Web App and other Function Apps) , So we agree it is going to be a managed identity but should I create one for each app or one for all apps?
If I create system MI then there should be one for each App.
If I create user MI then I can re-use it for any App I want with minimum change to AD

upvoted 105 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000

upvoted 7 times

Basu525 2 years, 7 months ago

Yes exactly. The question emphasizes on having multiple resources having to access Key vault. So using User Assigned MI is more pragmatic.

upvoted 3 times

SlavMar 2 years, 4 months ago

Well from security standpoint you could assign different secrets to different parts of application (if they are using system managed identity) soi if one component is compromised - not all secrets are available to hackers

upvoted 4 times

midhy 2 years, 7 months ago

User assigned cannot be used Look at the link provided.

upvoted 4 times

vb3d 2 years, 7 months ago

I just tried this, I was able to give access to a user assigned managed identity to a key vault.

upvoted 8 times

vb3d 2 years, 7 months ago

I think that note in the link is only for the usage of key vault references, which is a new thing in key vaults. I think a user assigned identity should still have access to the key vault.

upvoted 3 times

Zsolt72 2 years, 7 months ago

Yes, key vault reference is an App Configuration feature referencing a secret from key vault.

upvoted 1 times

Tom87 2 years, 6 months ago

That's true.
https://kasunkodagoda.com/2019/06/09/using-user-assigned-managed-identity-to-access-azure-key-vault-from-azure-app-service/

upvoted 1 times

vtomy 2 years, 4 months ago

User assigned MI is not supported for keyvault.
System assigned MI is correct.

upvoted 6 times

zzt 2 years, 3 months ago

This is not true: see: https://kasunkodagoda.com/2019/06/09/using-user-assigned-managed-identity-to-access-azure-key-vault-from-azure-app-service/

upvoted 6 times

PhilLI 1 year, 10 months ago

I don't read this on that page t.b.h.

upvoted 1 times

mlantonis Highly Voted 2 years, 5 months ago

Correct Answer: C

Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.

Reference:
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview#managed-identity-types
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

upvoted 37 times

TakumaK 2 years, 4 months ago

You are wrong. your second link OBVISOULY says
"Key Vault references will use the app's system assigned identity by default, but you can specify a user-assigned identity."

upvoted 18 times

eX900 1 year ago

Granting your app access to Key Vault
2.Create a managed identity for your application.
"Key Vault references will use the app's system assigned identity by default, but you can specify a user-assigned identity."

Your second reference:
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references.

So, the answer is : A

upvoted 4 times

zzt 2 years, 3 months ago

Not correct: see link https://kasunkodagoda.com/2019/06/09/using-user-assigned-managed-identity-to-access-azure-key-vault-from-azure-app-service/

upvoted 5 times

altafpatel1984 1 year, 11 months ago

Don't misguide people. In the same link you give, this statement exists "Some apps need to reference secrets at creation time, when a system-assigned identity would not yet be available. In these cases, a user-assigned identity can be created and given access to the vault in advance."

Your given reference: https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

upvoted 8 times

junkz Most Recent 1 week, 2 days ago

there is a saying in my country: one fool throws a rock in the pond, and 10 good people do their best to pull it out. this is exactly how i feel looking at all of us trying to "guess" what "minimizing AD changes" realy means in this context...

upvoted 1 times

drkpnsh 2 months, 2 weeks ago

I think in the past, you cannot use the user-assigned identity to access to the Azure Key-Vault. But it is possible now: "Key vault references use the app's system-assigned identity by default, but you can specify a user-assigned identity."

https://learn.microsoft.com/en-us/azure/app-service/app-service-key-vault-references?tabs=azure-cli

So I think the answer should be A.

upvoted 1 times

Dianahu 3 months ago

if I want to add an user-assigned managed identity I need to create it so I wouldnt say that it is minimum change to AD and it is not a good pattern to have the same identity for all of them

upvoted 1 times

AzSimuz 6 months, 4 weeks ago

Managed identities are of the following two types:

System-assigned -- this is activated directly on the Azure cloud service. For example, you can create a managed identity for your web app that will be automatically deleted once the web app instance is deleted
User-assigned -- this is a stand-alone Azure resource that is created independent of the actual service instance. You can assign this identifier to one or more Azure resources simultaneously. As an example, you can set this identity to both an Azure Web App and an Azure Function at the same time

- https://www.loginradius.com/blog/engineering/guest-post/using-azure-key-vault-with-an-azure-web-app-in-c-sharp/

upvoted 2 times

maqsoodshah 7 months ago

Correct Answer: A

System-assigned -- this is activated directly on the Azure cloud service. For example, you can create a managed identity for your web app that will be automatically deleted once the web app instance is deleted

User-assigned -- this is a stand-alone Azure resource that is created independent of the actual service instance. You can assign this identifier to one or more Azure resources simultaneously. As an example, you can set this identity to both an Azure Web App and an Azure Function at the same time

upvoted 1 times

JamieS 7 months, 3 weeks ago

I would have thought it was system assigned identity, but it says we have multiple Azure Function apps, and one of the requirements is to minimise changes to AD.

Only the user identity has a one to many mapping (https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview#managed-identity-types).

upvoted 3 times

manopeydakon 8 months, 2 weeks ago

Key Vault references will use the app's system assigned identity by default, but you can specify a user-assigned identity.
User-assigned managed identity Can be shared!
The same user-assigned managed identity can be associated with more than one Azure resource.
So A and C both works but less work with A in this case.

upvoted 1 times

Esward 9 months, 1 week ago

User assigned Identity
https://thecodeblogger.com/2020/06/13/user-assigned-managed-identity-with-azure-key-vault/

upvoted 1 times

Esward 9 months, 1 week ago

Also refer
https://yourazurecoach.com/2022/02/04/use-app-service-key-vault-references-with-user-assigned-identities/

upvoted 1 times

Esward 9 months, 1 week ago

https://stackoverflow.com/questions/55101525/how-can-i-give-access-to-key-vault-to-a-user-assigned-identity

upvoted 1 times

daothanhyen 9 months, 1 week ago

Selected Answer: C

C.
Because user assigned will take more codes and configurations for the key.
System assigned for each app service more security and noneed to have any configuration and the same code

upvoted 3 times

imanonion 9 months, 3 weeks ago

Selected Answer: C

By creating a system-assigned Managed Identity for each App Service, you can avoid storing secrets in the application or runtime environment, as the Managed Identity provides a way for the App Service to authenticate to Key Vault without the need for hard-coded credentials. Additionally, using a separate Managed Identity for each App Service allows for fine-grained access control to the secrets stored in Key Vault, as you can assign permissions to specific identities rather than granting broad permissions to a single identity.

upvoted 1 times

imanonion 9 months, 3 weeks ago

Option A, "Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity" also can be a good solution, But it can't provide fine-grained access control as all App Service using the same identity.

Option B "Create a single Azure AD Service Principal with permission to access Key Vault and use a client secret from within the App Services to access Key Vault" would not be the recommended approach, as it would involve storing the client secret in the App Services, which is a violation of the requirement to avoid storing secrets in the application or runtime environment.

Option D "Create an Azure AD Service Principal with Permissions to access Key Vault for each App Service and use a certificate from within the App Services to access Key Vault" is also a valid approach but it is more difficult to manage the certificate in each App Service and need more code to handle the certificate in the code.

upvoted 1 times

warchoon 8 months ago

User-assigned Managed Identity is an additional resource. System-assigned is not

upvoted 1 times

oceane0316 10 months, 2 weeks ago

The correct answer is A.
Source : https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/managed-identity-best-practice-recommendations#using-user-assigned-identities-to-reduce-administration

Read the article but in summary "When a user-assigned identity is associated with the four virtual machines, only two role assignments are required, compared to eight with system-assigned identities. If the virtual machines' identity requires more role assignments, they'll be granted to all the resources associated with this identity."

The key here is "Changes to Azure Active Directory (Azure AD) must be minimized" and the paragraph is called "Using user-assigned identities to reduce administration".

So A is the correct answer because of the "change must be minimized". C would work but it means you have much more changes.

upvoted 3 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: A

A. Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity.

upvoted 2 times

micro9000 11 months, 3 weeks ago

The answer should be A,
we have this requirements: Changes to Azure Active Directory (Azure AD) must be minimized.
If A, we only need to create single User Assigned MI, and allow that MI to access KV, then assign that MI to App Service and Function Apps

upvoted 3 times

hubekpeter 11 months, 1 week ago

Ask yourself please what kind of change in AD you are doing in C vs A. C is the most correct secure and straightforward answer. You don't need to touch AD at all.

upvoted 1 times

gmishra88 1 year ago

" Changes to Azure Active Directory (Azure AD) must be minimized" For what practical reasons? These questions doesn't make sense

upvoted 1 times

serpevi 1 year, 1 month ago

Got this in 09/22 , went with the most voted answer, score 927.

upvoted 2 times

VACAS 1 year, 1 month ago

what do you answer?

upvoted 1 times

Question #10 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms.
If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised.
You need to store the intake forms according to the requirements.
Solution:
1. Create an Azure Key Vault key named skey.
2. Encrypt the intake forms using the public key portion of skey.
3. Store the encrypted data in Azure Blob storage.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*

Community vote distribution

A (94%) 6%

pac1311 Highly Voted 2 years, 8 months ago

Correct!

upvoted 37 times

bingohah Highly Voted 2 years, 7 months ago

The answer is right, but not good. Because we should use symmetric key to encrypt file, instead of public key. The best practise is: encrypt file by symmetric key, and the symmetric key is protected by public key.

upvoted 27 times

vb3d 2 years, 7 months ago

Not sure what do you mean, but check the link below.
Symmetric key means the same key is used for encryption and decryption. When you use an Asymmetric key, you encrypt with the public key and decrypt with a secret key, which is what the solution is talking about and is better.
https://www.ssl2buy.com/wiki/symmetric-vs-asymmetric-encryption-what-are-differences#:~:text=Symmetric%20encryption%20uses%20a%20single,and%20decrypt%20messages%20when%20communicating.

upvoted 11 times

aarrtturas 2 years, 5 months ago

We should use symmetric key not asymmetric key. Asymmetric key is used for small amount's of data to encrypt. Usually asymmetric key encrypts symmetric key, and symmetric key encrypts data

upvoted 5 times

fearoffree 1 year, 8 months ago

The asymmetric key is used to sign the intake so it could not be compromised!

upvoted 1 times

richardc7 Most Recent 1 month ago

Selected Answer: A

Got 29/09/2023

upvoted 2 times

notedo 1 month ago

Selected Answer: A

Correct

upvoted 1 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 3 times

Saluk_DE 7 months ago

Selected Answer: A

Question was in Exam 2023-03-30

upvoted 2 times

sarmaria 7 months, 2 weeks ago

Got this on 16/03/23. Chosen yes. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 1 times

Bear_Polar 11 months, 1 week ago

Selected Answer: A

AZ-900, is that you? :)))

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: A

A. Yes

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: A

A. Yes

upvoted 2 times

coffecold 1 year ago

The anwser is B (NO)

This is a business case of digital signing.
In digital signing the sender's private key encrypts the data -- this is the digital signature -- and the receiver uses the public key to decrypt it and verify it matches the attachment.
"2. Encrypt the intake forms using the public key portion of skey."
must be
"2. Encrypt the intake forms using the private key portion of skey."

Digital signing is opposed to encryption : Encryption uses a public key to ensure the ciphertext cannot be deciphered by anyone but the authorized recipient with the private key.

upvoted 3 times

gmishra88 1 year ago

Client side encryption is a valid method to do this. Asymmetric key can be used (RSA in keyvault). If using Asymetric key then client side encryption should be done with the public key because it can only be decrypted with a private key (which other people do not have). One will encrypt with private key only as a signature to prove she has the private key so that it can be verified with the public key.
But still a difficult question in a certification exam with a time limit if you used their SDK to do the encryption and decryption.

upvoted 1 times

gmishra88 1 year ago

And, I will always use a symmetric key for client side encryption. But god knows what Microsoft want as an answer and what the question creator experienced in life.

upvoted 1 times

Dani_ac7 1 year, 1 month ago

Selected Answer: A

Correct

upvoted 1 times

Knightie 1 year, 2 months ago

Selected Answer: B

"If the stored intake forms are downloaded from storage by a third party,"

The third party can not access the decrypted format of the content if you encrypt with keys from your vault.... next question is the correct solution.

upvoted 1 times

FatSeal 1 year, 2 months ago

"If the stored intake forms are downloaded from storage by a third party, the contents of the forms must NOT be compromised"

according to the whole sentence, the third party should not be able to decrypt it

upvoted 2 times

Eltooth 1 year, 4 months ago

Selected Answer: A

A is correct answer.

upvoted 2 times

xRiot007 1 year, 4 months ago

Selected Answer: A

The answer is Yes. Encrypted content will remain encrypted upon storing it and will be delivered as encrypted to anyone who requests it.

upvoted 1 times

Rini100 1 year, 6 months ago

Got this on 20 Apr 2022

upvoted 2 times

Question #11 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms.
If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised.
You need to store the intake forms according to the requirements.
Solution:
1. Create an Azure Cosmos DB database with Storage Service Encryption enabled.
2. Store the intake forms in the Azure Cosmos DB database.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

Community vote distribution

B (82%) A (18%)

cbn Highly Voted 2 years, 8 months ago

I think the rationale is Azure Cosmos is not the best place to store blob data, like scanned patient intake forms.

https://stackoverflow.com/questions/49343675/how-can-i-store-images-in-azure-cosmos-db

upvoted 30 times

iamstudying 1 year, 7 months ago

Although it's not ideal, it is still a valid solution (sort of)... media files are stored in blob storage https://docs.microsoft.com/en-us/azure/cosmos-db/database-encryption-at-rest

The key thing to point here is that the solution requires encrypted files when downloaded, but the solution here is only encryption at rest

upvoted 9 times

gmishra88 1 year ago

One of those typical Microsoft questions that punishes people who know something. If you do not understand the question then you will answer No, correctly (don't store images in cosmos), but if you know encryption to the next level you think it is encrypted so it is Yes. But one more level if you think then server side encryption (at rest) is not enough if it is already downloaded.

upvoted 2 times

pieronegri 2 years, 8 months ago

agreed

upvoted 2 times

mlantonis Highly Voted 2 years, 5 months ago

Correct Answer: B - No

If storage is encrypted by default but you need to encrypt the content so when downloaded, its useless to unauthorised party.

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

upvoted 29 times

TMakki 1 year, 2 months ago

Absolutely right. Data is encrypted at rest but when downloaded it is not.

upvoted 3 times

MysticalSam Most Recent 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 3 times

Saluk_DE 7 months ago

Selected Answer: B

Question was in Exam 2023-03-30

upvoted 1 times

sarmaria 7 months, 2 weeks ago

Got this on 16/03/23. Chosen No. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

hubekpeter 11 months, 1 week ago

Selected Answer: B

For a god sake. This is an encryption at rest, where you are encrypting the underlying storage doesn't matter if it's with a azure or customer managed keys instead of encrypting each document itself on an upper layer.

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: B

No - correct ans

upvoted 1 times

eX900 1 year ago

If I have any doubts about the question, I would check these users (who are the top commenters) to be sure:
- Eltooth
- mlantonis
- glam

Thank you guys. 👍

upvoted 3 times

TMakki 1 year, 2 months ago

Selected Answer: B

Data is encrypted at rest but when downloaded it is not. So when downloaded by third party it is not encrypted so instead use an Azure Key vault and public key encryption

upvoted 3 times

Knightie 1 year, 2 months ago

Selected Answer: A

It is encrypted at rest, but available (not encrypted) for third party to download the copy, so the third party can read, come on. regarding encryption on transfer, it's the job of https.. so that's the way for third party to get the scanned copy, right?

upvoted 2 times

alcideus 1 year, 2 months ago

Then why did you put A as your answer lol

upvoted 3 times

xRiot007 1 year, 4 months ago

The answer is No. Encrypting the storage itself only. If someone requests it, they will get plain data.

upvoted 3 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.

upvoted 3 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 2 times

altafpatel1984 1 year, 11 months ago

I guess Cosmos DB is for semi-structured data, not for unstructured data. Here scanned copies of forms are pdf/image files which are unstructured files, hence they cannot be stored in CosmosDB. Instead should be stored in Blob Storage which also supports security for data in-transit and data at rest. And because of this B - No is correct.

upvoted 2 times

glam 2 years, 5 months ago

B. No..

upvoted 5 times

Frakandel 2 years, 5 months ago

Question is unclear... However, if I interpret the question as a third part trying to hack the database by directly downloading (outside of the Cosmos environment) the stored documents-, the answer is Yes. because the data is encrypted at rest/is stored encrypted...

upvoted 3 times

jokergester 2 years, 7 months ago

Encryption at rest should not be confused with encrypting the data. The former is encrypting the SSD/HDD whereas the latter is data is being encrypted. The third party specified can still download the encrypted data from blob but will not be able read it plainly without having the decryption or private key.

Blob encryption - https://docs.microsoft.com/en-us/azure/storage/blobs/storage-encrypt-decrypt-blobs-key-vault?tabs=dotnet#encrypt-blob-and-upload

upvoted 7 times

SummerWarrior 1 year, 9 months ago

How can third party have access to the data if Cosmos DB access is restricted to only certain identities?

upvoted 2 times

raelo 2 years, 6 months ago

Encryption at rest is a phrase that commonly refers to the encryption of data on nonvolatile storage devices, such as solid state drives (SSDs) and hard disk drives (HDDs)

From the quote above, it is the data that is being encrypted. https://docs.microsoft.com/en-us/azure/cosmos-db/database-encryption-at-rest#:~:text=With%20the%20release%20of%20encryption,is%20very%20easy%20to%20use.

upvoted 1 times

Question #12 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms.
If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised.
You need to store the intake forms according to the requirements.
Solution: Store the intake forms as Azure Key Vault secrets.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

Community vote distribution

B (100%)

clarionprogrammer Highly Voted 2 years, 6 months ago

B. No
You'd definitely not want to store "data" in Azure Key Vault.

upvoted 25 times

mlantonis Highly Voted 2 years, 5 months ago

Correct Answer: B - No

You'd definitely not want to store "data" in Azure Key Vault.

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

upvoted 17 times

richardc7 Most Recent 1 month ago

Selected Answer: B

Got 28/09/2023
No

upvoted 1 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 2 times

Saluk_DE 7 months ago

Selected Answer: B

Question was in Exam 2023-03-30

upvoted 1 times

sarmaria 7 months, 2 weeks ago

Got this on 16/03/23. Chosen No. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

NombreFalso 8 months, 2 weeks ago

No, that's crazy.

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: B

B. No -blob can not be stored in azure key vault

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.

upvoted 5 times

Rini100 1 year, 6 months ago

Got this on 20 Apr 2022

upvoted 3 times

Freidrich 1 year, 8 months ago

Selected Answer: B

The correct answer is B: No.

upvoted 1 times

Kvm1 2 years, 4 months ago

B. No
Technically possible.. but KeyValut is intended to store securely Application secrets, certificates etc. and not to store application data like the Intake forms and other documents.

upvoted 7 times

glam 2 years, 5 months ago

B. No...

upvoted 3 times

jokergester 2 years, 7 months ago

Technically possible but the documentation insists on using the Key Vault as storage for Application data or config not user data.

upvoted 3 times

Question #13 *Topic 4*

HOTSPOT -

You plan to deploy a new application to a Linux virtual machine (VM) that is hosted in Azure.

The entire VM must be secured at rest by using industry-standard encryption technology to address organizational security and compliance requirements.

You need to configure Azure Disk Encryption for the VM.

How should you complete the Azure CLI commands? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
az provider register -n Microsoft.KeyVault
resourcegroup="myResourceGroup"
az group create --name $resourcegroup --location westus
keyvault_name=myvaultname$RANDOM
az [ ▼ ] create \
```

- vm
- keyvault
- keyvault key
- vm encryption

```
--name $keyvault_name \
--resource-group $resourcegroup \
--location eastus \
--enabled-for-disk-encryption True
az [ ▼ ] create \
```

- vm
- keyvault
- keyvault key
- vm encryption

```
--vault-name $keyvault_name \
--name Name1 \
--protection software
az [ ▼ ] create \
```

- vm
- keyvault
- keyvault key
- vm encryption

```
--resource-group $resourcegroup \
--name Name2 \
--image Canonical:UbuntuServer:16.04-LTS:latest \
--admin-username azureuser \
--generate-ssh-keys \
--data-disk-sizes-gb 5
az [ ▼ ] enable\
```

- vm
- keyvault
- keyvault key
- vm encryption

```
--resource-group $resourcegroup \
--name Name2 \
--disk-encryption-keyvault $keyvault_name \
--key-encryption-key Name1 \
--volume-type [ ▼ ]
```

- all
- data
- os

Answer Area

```
az provider register -n Microsoft.KeyVault
resourcegroup="myResourceGroup"
az group create --name $resourcegroup --location westus
keyvault_name=myvaultname$RANDOM
az [ ▼ ] create \
    --name $keyvault_name \
    --resource-group $resourcegroup \
    --location eastus \
    --enabled-for-disk-encryption True
az [ ▼ ] create \
    --vault-name $keyvault_name \
    --name Name1 \
    --protection software
```

vm / keyvault / keyvault key / vm encryption

vm / keyvault / keyvault key / vm encryption

Correct Answer:

```
az [ ▼ ] create \
    --resource-group $resourcegroup \
    --name Name2 \
    --image Canonical:UbuntuServer:16.04-LTS:latest \
    --admin-username azureuser \
    --generate-ssh-keys \
    --data-disk-sizes-gb 5
az [ ▼ ] enable\
    --resource-group $resourcegroup \
    --name Name2 \
    --disk-encryption-keyvault $keyvault_name \
    --key-encryption-key Name1 \
    --volume-type [ ▼ ]
```

vm / keyvault / keyvault key / vm encryption

vm / keyvault / keyvault key / vm encryption

all / data / os

Box 1: keyvault -

Create an Azure Key Vault with az keyvault create and enable the Key Vault for use with disk encryption. Specify a unique Key Vault name for keyvault_name as follows: keyvault_name=myvaultname$RANDOM az keyvault create \

--name $keyvault_name \

--resource-group $resourcegroup \

--location eastus \

--enabled-for-disk-encryption True

Box 2: keyvault key -

The Azure platform needs to be granted access to request the cryptographic keys when the VM boots to decrypt the virtual disks. Create a cryptographic key in your Key Vault with az keyvault key create. The following example creates a key named myKey: az keyvault key create \

--vault-name $keyvault_name \

--name myKey \

--protection software

Box 3: vm -

Create a VM with az vm create. Only certain marketplace images support disk encryption. The following example creates a VM named myVM using an Ubuntu
16.04 LTS image:
az vm create \
--resource-group $resourcegroup \
--name myVM \
--image Canonical:UbuntuServer:16.04-LTS:latest \
--admin-username azureuser \
--generate-ssh-keys \

Box 4: vm encryption -
Encrypt your VM with az vm encryption enable:
az vm encryption enable \
--resource-group $resourcegroup \
--name myVM \
--disk-encryption-keyvault $keyvault_name \
--key-encryption-key myKey \
--volume-type all
Note: seems to an error in the question. Should have enable instead of create.

Box 5: all -
Encrypt both data and operating system.
Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/linux/disk-encryption-cli-quickstart

Tom87 Highly Voted 2 years, 6 months ago
The answer is correct.
Other possiblities doesn't make sense. Although it seems strange to me to create the VM when the question says that we already have one.
upvoted 56 times

Tom87 2 years, 6 months ago
Ok, now I see that the question doesn't say that the VM is already created. Sorry for that :)
upvoted 7 times

glam Highly Voted 2 years, 5 months ago
correct
upvoted 7 times

Juhi_Az204 Most Recent 7 months, 1 week ago
Got this on 26March,2023
upvoted 5 times

sarmaria 7 months, 2 weeks ago
Got this on 16/03/23. Went with proposed solution. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.
upvoted 2 times

Esward 9 months, 1 week ago
VM
Keyvault
Keyvault key
VM Enryption
ALL

Refer: https://learn.microsoft.com/en-us/azure/virtual-machines/linux/disk-encryption-cli-quickstart
upvoted 3 times

manopeydakon 8 months, 2 weeks ago
Look at the whole picture son! Answer is correct
upvoted 2 times

rolling_potato_ 10 months ago
Was on my exam today (03-01-2023)
upvoted 6 times

Lucifer14 10 months ago
In 28-12-2022 exam

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Key vault

Key vault key
VM
VM encryption
All

upvoted 3 times

at_sayali 11 months, 3 weeks ago

The answer is correct.
Got this in 10/2022

upvoted 2 times

gmishra88 1 year ago

The documentation here does not mention the step to create the keyvault key.
https://learn.microsoft.com/en-us/azure/virtual-machines/linux/disk-encryption-cli-quickstart.
But it is okay

upvoted 1 times

serpevi 1 year, 1 month ago

Got this in 09/22 , went with the most voted answers, score 927.

upvoted 2 times

Eltooth 1 year, 4 months ago

Answer is correct.
Keyvault
Keyvault key
VM
VM encryption
All

upvoted 3 times

Baskman 1 year, 7 months ago

Got this in the exam 03/22

upvoted 5 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with the given answer

upvoted 5 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 5 times

tramlong888 1 year, 10 months ago

The answer is correct.

upvoted 1 times

debanjan10 2 years ago

Order: keyvault > keyvault key > vm > vm encryption

upvoted 4 times

Question #14 *Topic 4*

Your company is developing an Azure API hosted in Azure.

You need to implement authentication for the Azure API to access other Azure resources. You have the following requirements:

- All API calls must be authenticated.
- Callers to the API must not send credentials to the API.

Which authentication mechanism should you use?

A. Basic

B. Anonymous

C. Managed identity

D. Client certificate

Correct Answer: *C*

Azure Active Directory Managed Service Identity (MSI) gives your code an automatically managed identity for authenticating to Azure services, so that you can keep credentials out of your code.

Note: Use the authentication-managed-identity policy to authenticate with a backend service using the managed identity. This policy essentially uses the managed identity to obtain an access token from Azure Active Directory for accessing the specified resource. After successfully obtaining the token, the policy will set the value of the token in the Authorization header using the Bearer scheme.

Incorrect Answers:

A: Use the authentication-basic policy to authenticate with a backend service using Basic authentication. This policy effectively sets the HTTP Authorization header to the value corresponding to the credentials provided in the policy.

B: Anonymous is no authentication at all.

D: Your code needs credentials to authenticate to cloud services, but you want to limit the visibility of those credentials as much as possible. Ideally, they never appear on a developer's workstation or get checked-in to source control. Azure Key Vault can store credentials securely so they aren't in your code, but to retrieve them you need to authenticate to Azure Key Vault. To authenticate to Key Vault, you need a credential! A classic bootstrap problem.

Reference:

https://azure.microsoft.com/en-us/blog/keep-credentials-out-of-code-introducing-azure-ad-managed-service-identity/

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

Community vote distribution

C (100%)

clarionprogrammer Highly Voted 2 years, 6 months ago

C is correct.

upvoted 26 times

nnvuf Highly Voted 2 years, 4 months ago

A and D will send credential to the API.
B is not meet the requirement.
only C can be selected

upvoted 17 times

OPT_001122 Most Recent 11 months, 2 weeks ago

Selected Answer: C

C. Managed identity

upvoted 2 times

gmishra88 1 year ago

Another pointless requirement made up with some answer in mind and it is impossible to guess what the Microsoft guy had in his mind at that point in time (of course nobody reviews relevance of the questions once created). The requirement that the credentials are not sent is weird. It depends on whether a client certificate is seen as a credential. That gets in conflict with whether the API is called from a client that does not have a managed identity. For example, from postman or curl or a Single Page Application. You cannot use Managed identity in that case. So, depending on how much an unfortunate question taker knows there is 50% chance (excuse my probability knowledge) that he can make a mistake or not. Microsoft guy wins

upvoted 4 times

gmishra88 1 year ago

My mistake, the requirement did say one azure resource to access another azure resource. In that case Managed Identity is clearly a winner. One must really read the question properly

upvoted 2 times

Nokaido 1 year, 1 month ago

Selected Answer: C

Tested it with a App Service running my C# code while using the Azure identity packages.

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: C

C is the correct answer.

upvoted 2 times

Eltooth 1 year, 4 months ago

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

upvoted 1 times

AZ204Cert 1 year, 6 months ago

Got this on 04/05/22 (selected Managed Identity)

upvoted 7 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 2 times

ning 2 years, 2 months ago

Question is about from APIM to other azure resource, C is correct managed identity. Question is NOT about how end users access APIM

upvoted 6 times

francis6170 2 years, 4 months ago

got this in the exam :)

upvoted 5 times

businesskasper 2 years, 5 months ago

Managed identity only works for internal api calls. The requirement is "all api calls must be authenticated". Certificate would meet that requirement?

upvoted 4 times

glam 2 years, 5 months ago

correct

upvoted 3 times

MrZoom 2 years, 7 months ago

D would work though, and it meets all the criteria (the private key of the certificate is never sent to the server during https handshake). If multiple answers are possible, I'd go for C and D, but if only one answer is possible, C is better for the reason stated in the answer explanation section.

upvoted 2 times

MrZoom 2 years, 7 months ago

Also, using certificates for authentication isn't typical use with API's hosted in Azure (which _is_ part of the question).

upvoted 2 times

Spooky7 2 years, 5 months ago

How ClientCerficate would help Azure API to authenticate in other Azure resources? The question is not about client authentication but Azure API authentication to use other Azure resources.

upvoted 4 times

Question #15 *Topic 4*

DRAG DROP -

You are developing an application. You have an Azure user account that has access to two subscriptions.

You need to retrieve a storage account key secret from Azure Key Vault.

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

Select and Place:

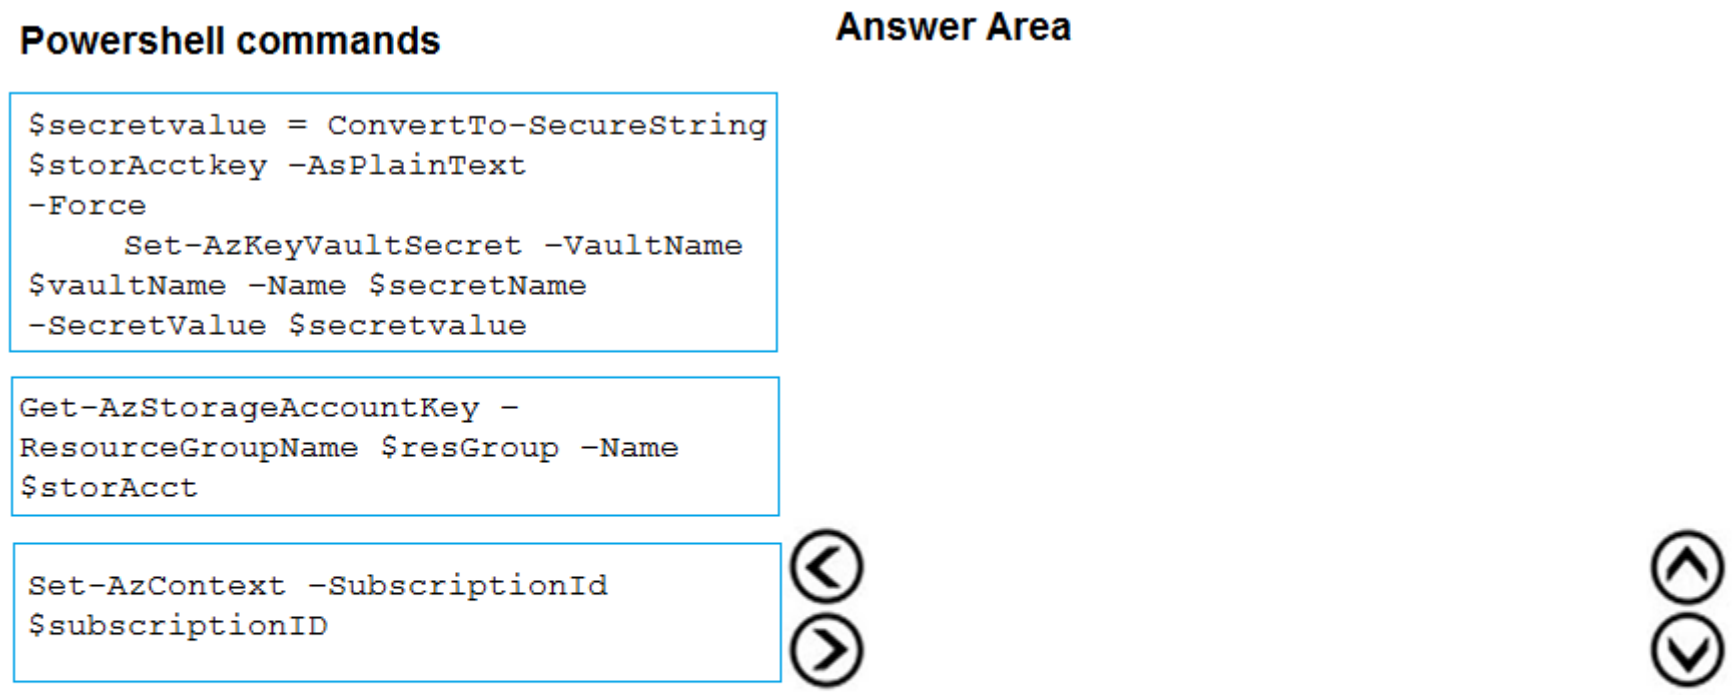

Correct Answer:

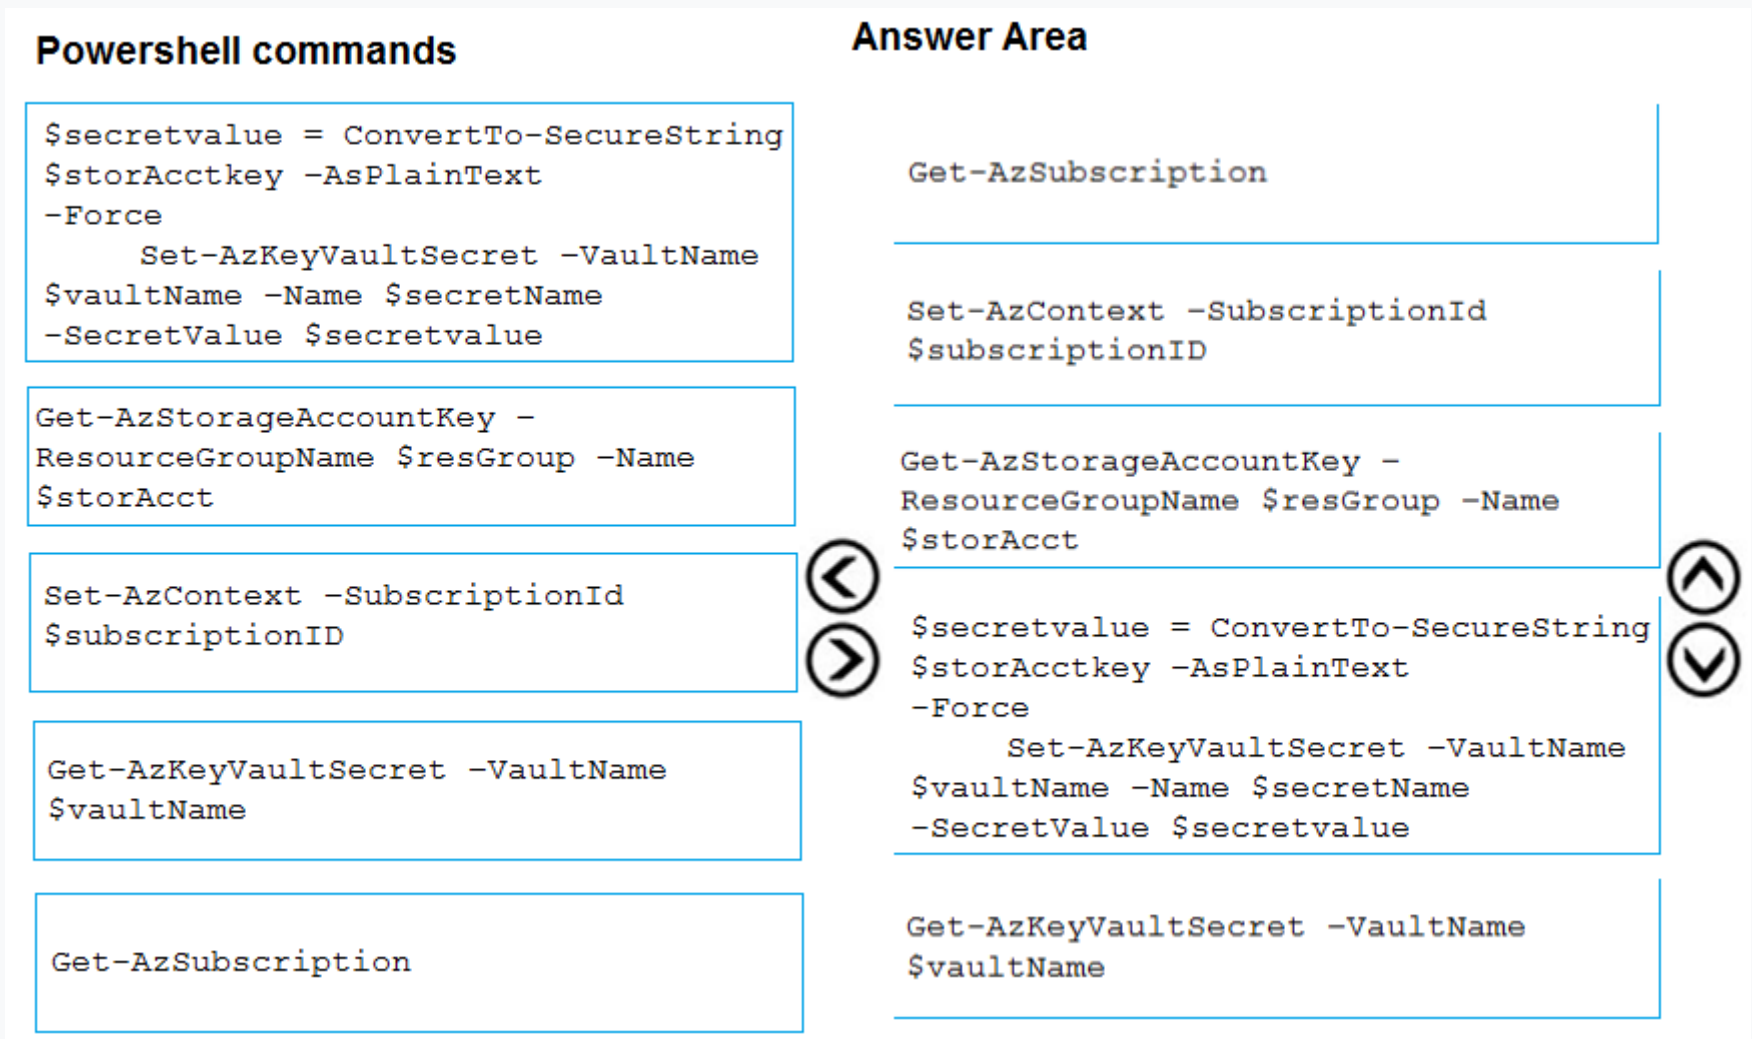

Step 1: Get-AzSubscription -

If you have multiple subscriptions, you might have to specify the one that was used to create your key vault. Enter the following to see the subscriptions for your account:

Get-AzSubscription -

Step 2: Set-AzContext -SubscriptionId

To specify the subscription that's associated with the key vault you'll be logging, enter:

Set-AzContext -SubscriptionId <subscriptionID>

Step 3: Get-AzStorageAccountKey -
You must get that storage account key.
Step 4: $secretvalue = ConvertTo-SecureString <storageAccountKey> -AsPlainText -Force
Set-AzKeyVaultSecret -VaultName <vaultName> -Name <secretName> -SecretValue $secretvalue
After retrieving your secret (in this case, your storage account key), you must convert that key to a secure string, and then create a secret with that value in your key vault.

Step 5: Get-AzKeyVaultSecret -
Next, get the URI for the secret you created. You'll need this URI in a later step to call the key vault and retrieve your secret. Run the following PowerShell command and make note of the ID value, which is the secret's URI:
Get-AzKeyVaultSecret ɜ€"VaultName <vaultName>
Reference:
https://docs.microsoft.com/bs-latn-ba/Azure/key-vault/key-vault-key-rotation-log-monitoring

AnonymousJhb Highly Voted 2 years, 10 months ago
Answer looks correct
upvoted 52 times

Bear_Polar 11 months, 1 week ago
Correct! Reference: https://learn.microsoft.com/bs-latn-ba/azure/key-vault/secrets/tutorial-rotation-dual?tabs=azurepowershell#add-the-storage-account-access-keys-to-key-vault-secrets
upvoted 2 times

Frakandel Highly Voted 2 years, 5 months ago
Question does not explain where the variables come from... Assuming I set the variables myself (as a "side effect" of the command), the solution below does make sense:
1. Get-AzSubscription
2. Set-AzContext –SubscriptionId $subscriptionID
3. Get-AzKeyVaultSecret –VaultName $vaultName
4. Get-AzStorageAccountKey –ResourceGroupName $resGroup –Name $storAcct
5. $secretvalue = ConvertTo-SecureString $storAcctkey –AsPlainText –Force
Set-AzKeyVaultSecret –VaultName $vaultName –Name $secretName –SecretValue $secretvalue
upvoted 35 times

p2006 Most Recent 1 month, 1 week ago
https://learn.microsoft.com/en-us/azure/key-vault/secrets/tutorial-rotation-dual?tabs=azurepowershell#add-storage-account-access-key-to-key-vault-secrets
upvoted 4 times

suku555 1 year, 5 months ago
Given Answer is correct only. "Get-AzKeyVaultSecret -VaultName $vaultName" is just to make sure that we stored the key in vault (verification/double check). Placing this command anywhere else in the order does not make any sense.
upvoted 3 times

suku555 1 year, 5 months ago
Also the main purpose of "Get-AzKeyVaultSecret -VaultName $vaultName" at last is for the future usage..
upvoted 2 times

Azprep 1 year, 6 months ago
Given answer us correct
upvoted 2 times

MiraA 2 years, 1 month ago
I believe the assignment wants to select a proper subscription, then to retrieve the storage account key, then to store this key into the KeyVault and finally to check the secret was inserted properly. So:

1. Get-AzSubscription ($subscriptionID = (...).SubscriptionName)
2. Set-AzContext
3. Get-AzStorageAccountKey ($storAcctkey = "(...)[0].Value")
4. ConvertTo-SecureString, Set-AzKeyVaultSecret
5. Get-AzKeyVaultSecret (list the secrets stored in the KeyVault - check only)
upvoted 27 times

ning 2 years, 2 months ago
Does the question miss some part? Why it has to save the key before retrieval? No where in the question mentioned that ... First two
Get-AzSubscription
Set-AzContext –SubscriptionId $subscriptionID
are absolutely correct, since there are two subscriptions, you have to point to the correct one ... but I am confused with 3, 4, 5 steps, what are those?

upvoted 12 times

glam 2 years, 5 months ago

correct.

upvoted 3 times

sushikid112 2 years, 9 months ago

Agree with some people saying the question and picture does not make sense. The question only says how to retrieve the secret from Key Vault. Based on the options, we also have to retrieve the value and insert it to Key Vault first. The question/picture lacks information based on what we have.

upvoted 8 times

Rodashar 2 years, 10 months ago

Is this question actually on the exam? Even the url provided to explain the answer says this code is meant to cycle the storage keys stored in the vault not retrieve a key from the vault. Either the question is wrong or the answer is wrong.

upvoted 6 times

Tealon 2 years, 10 months ago

I agree, the question does not match the options to give the answer. If the question was to store the account key to the key vault and then retrieve it from the key vault, the answer would be correct.

upvoted 5 times

iiiihhhh 2 years, 10 months ago

There is another discussion for this question: https://www.examtopics.com/discussions/microsoft/view/22272-exam-az-204-topic-3-question-5-discussion/

upvoted 2 times

iiiihhhh 2 years, 10 months ago

The question is: How to get secret from key vault, when there are two subscriptions available?

Available components(in order of boxes):
1. converting $storageAcctkey to secure string and storing az secret with name $secretName
2. getting storage accout key from storage account (but without storing to variable $storageAcctkey)
3. setting subscription context for use in current session (https://docs.microsoft.com/en-us/powershell/module/az.accounts/set-azcontext?view=azps-5.1.0)
4. list secrets in the vault
5. list subscriptions available

So for retrieving secret we (I think) need 5,3,4 and we need also SecureStringToBSTR (https://docs.microsoft.com/cs-cz/azure/key-vault/secrets/quick-create-powershell).
However the question seems to require using all actions available.

Is the question text really matching the picture?

upvoted 7 times

Question #16 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop Azure solutions.
You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.
You need to obtain an Azure Resource Manager access token.
Solution: Use an X.509 certificate to authenticate the VM with Azure Resource Manager.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
Instead run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm

Community vote distribution
B (100%)

KeerthiKP Highly Voted 3 years ago
No is the right answer:
Explanation :
Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.
link : https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm
upvoted 31 times

clarionprogrammer 2 years, 6 months ago
Using Invoke-WebRequest vs Invoke-RestMethod doesn't matter. The point is that using a certificate is wrong for Identity Management.
upvoted 13 times

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: B - No

A certificate cannnot be used to authenticate.

Instead run the Invoke-RestMethod or Invoke-WebRequest cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

https://docs.microsoft.com/en-us/azure/role-based-access-control/overview

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm
upvoted 16 times

Esward Most Recent 9 months, 1 week ago
Correct Answer: B - No

Azure RBAC is an authorization system built on Azure Resource Manager that provides fine-grained access management of Azure resources. Here we need to authenticate.

Here we need to make use of managed identities for the virtual machine. Role-based access control is used for authorization and not authentication.

Instead run the Invoke-RestMethod or Invoke-WebRequest cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:
https://docs.microsoft.com/en-us/azure/role-based-access-control/overview
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm
upvoted 3 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: B

Correct Answer: B
upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

glam 2 years, 5 months ago

B. No...
upvoted 2 times

Frakandel 2 years, 5 months ago

Yes, because a certificate can be used to authenticate... See also: https://devblogs.microsoft.com/premier-developer/centralized-vm-certificate-deployment-across-multiple-regions-with-arm-templates/
upvoted 1 times

TakumaK 2 years, 5 months ago

Yes, your right. But can you explain this part in the question?
"You need to obtain an Azure Resource Manager access token"
upvoted 2 times

RahulKate 2 years, 10 months ago

Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.
upvoted 1 times

Tealon 2 years, 10 months ago

The given answer is correct.
upvoted 5 times

cyberbull 3 years, 1 month ago

No , is the correct Answer
upvoted 8 times

Question #17 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop Azure solutions.
You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.
You need to obtain an Azure Resource Manager access token.
Solution: Use the Reader role-based access control (RBAC) role to authenticate the VM with Azure Resource Manager.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
Instead run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm

Community vote distribution
B (80%) A (20%)

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: B - No

Azure RBAC is an authorization system built on Azure Resource Manager that provides fine-grained access management of Azure resources. Here we need to authenticate.

Here we need to make use of managed identities for the virtual machine. Role-based access control is used for authorization and not authentication.

Instead run the Invoke-RestMethod or Invoke-WebRequest cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:
https://docs.microsoft.com/en-us/azure/role-based-access-control/overview
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm
upvoted 35 times

KeerthiKP Highly Voted 3 years ago
ITs not a invoke-RestMethod , it is: Invoke-WebRequest cmdlet

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm
upvoted 20 times

oversio 3 years ago
Yes! you all right
upvoted 2 times

Tom87 2 years, 6 months ago
Both Invoke-RestMethod and Invoke-WebRequest can be used. Their behaviour is very similar.
https://blog.truesec.com/2020/07/29/invoke-webrequest-vs-invoke-restmethod/
upvoted 5 times

clarionprogrammer 2 years, 6 months ago
Agreed. The real reason it is wrong is RBAC can't be used for Identity Management.
upvoted 2 times

clarionprogrammer 2 years, 6 months ago
It's true that the Reader role is required so the VM can read/access the Azure Resource Manager. But, it's not used to authenticate the VM with the Azure Resource Manager.
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm
upvoted 3 times

hubekpeter Most Recent 11 months, 1 week ago
Grant access to specific resource groups in Azure Resource Manager - resource based access control is perfectly usable, but a key here is to limit the access to a specific group, which is not possible through reader role, you need to implement ABAC policy in this case. So the answer is NO.
upvoted 1 times

coffecold 1 year ago

Selected Answer: A

Answer A - because this particular setup is appropriate to get the token (among others).

User-assigned managed identity with RBAC is just one way to obtain an Azure Resource Manager access token. Having done that still you have to get the access token by an Invoke-WebRequest.

step 1 : get the access token

$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&client_id=<value rof the user-assigned managed identity>&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}

step 2 : access Azure Resource Manager with the access token

See https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-ua-arm

upvoted 1 times

coffecold 1 year ago

although ... the Reader role-based access control (RBAC) role is not used here to get the token, it's the client ID in the invoke web-request. So to be save : answer B

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.

upvoted 3 times

sawipef270 1 year, 10 months ago

Selected Answer: B

correct

upvoted 1 times

glam 2 years, 5 months ago

B. No...

upvoted 2 times

ccinetto 2 years, 10 months ago

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm
Here it says that you first need to go to the Resource Group, and in the IAM tab grant access to the VM as a Reader Role
Then, when you are inside the VM, you can use the command Invoke-WebRequest cmdlet in order to get the Token
4 is Correct because it is the first step you need to do in order to get the Token
and question 5 is correct too

upvoted 4 times

azahran 2 years, 9 months ago

The question is you need an access token? Q 4 is preparation step and did not retrieve an access token.

upvoted 3 times

bugimachi 2 years, 10 months ago

I'm not sure. What you do in IAM / RBAC is authorizing (!) the VM to access the resource group. You do NOT use IAM / RBAC for authENTICATION.

I mean, you are right in saying that it is needed to grant access for the VM. But strictly speaking, the question asks for authentication only.

upvoted 4 times

Gathix444 2 years, 9 months ago

Yes you do. https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-linux-vm-access-arm

upvoted 1 times

27close 2 years, 11 months ago

Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager. answer is NO

upvoted 2 times

blueturtle 2 years, 11 months ago

Yup the updated version is to use Invoke-WebRequest cmdlet.

upvoted 1 times

Question #18 *Topic 4*

HOTSPOT -

You are building a website that is used to review restaurants. The website will use an Azure CDN to improve performance and add functionality to requests.

You build and deploy a mobile app for Apple iPhones. Whenever a user accesses the website from an iPhone, the user must be redirected to the app store.

You need to implement an Azure CDN rule that ensures that iPhone users are redirected to the app store.

How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
"conditions": [ {
    "name": "IsDevice",
    "parameters": {
      "@odata.type": "#Microsoft.Azure.Cdn.Models.[ ▼ ]",
      "operator": "Equal",
      "matchValues": [ "[ ▼ ]" ]
} },
  {
    "name": "RequestHeader",
    "parameters": {
      "@odata.type": "#Microsoft.Azure.Cdn.Models.[ ▼ ]",
      "operator": "Contains",
      "selector": "[ ▼ ]",
      "matchValues": [ "[ ▼ ]" ]
  } }
]
```

DeliveryRuleIsDeviceConditionParameters
DeliveryRuleCookiesConditionParameters
DeliveryRulePostArgsConditionParameters
DeliveryRuleRequestHeaderConditionParameters

iOS
Mobile
iPhone
Desktop

DeliveryRuleIsDeviceConditionParameters
DeliveryRuleCookiesConditionParameters
DeliveryRulePostArgsConditionParameters
DeliveryRuleRequestHeaderConditionParameters

FROM
PRAGMA
X-POWERED-BY
HTTP_USER_AGENT

iOS
Mobile
iPhone
Desktop

Correct Answer:

Answer Area

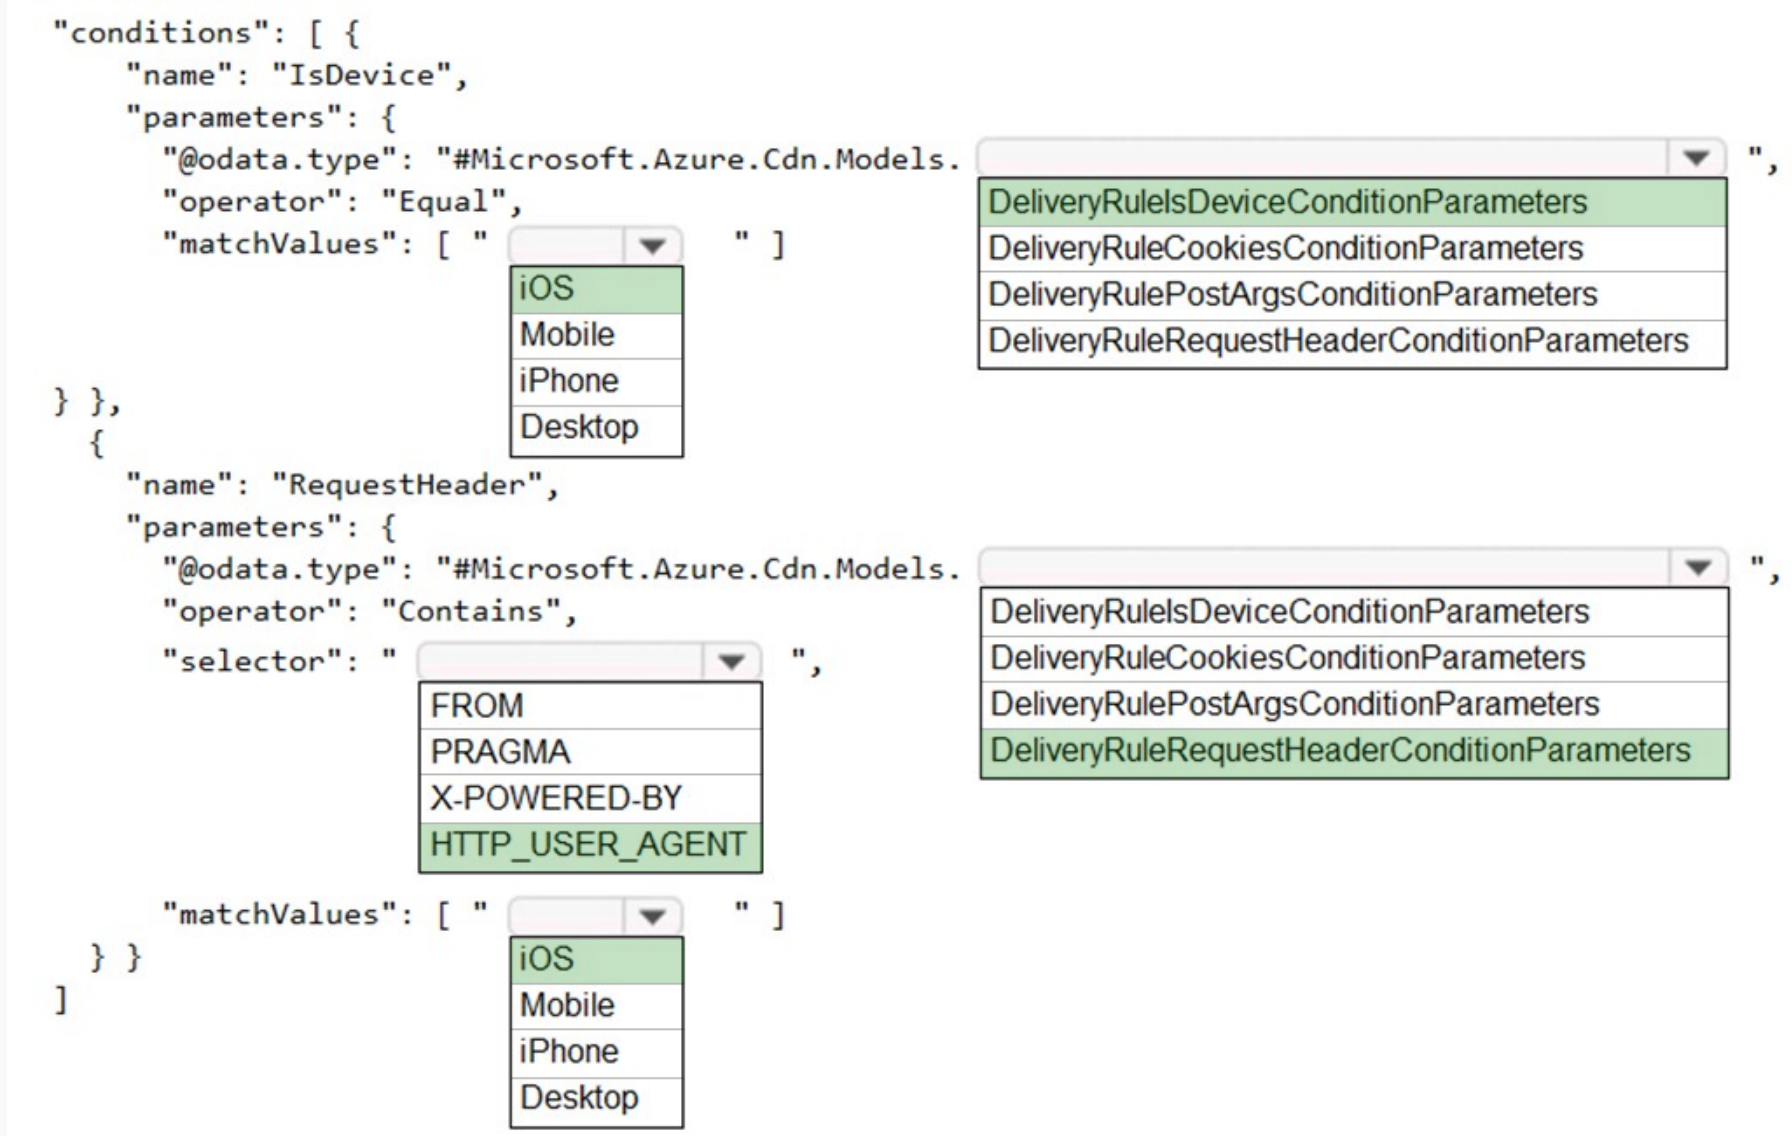

Box 1: iOS -

Azure AD Conditional Access supports the following device platforms:

☞ Android

☞ iOS

☞ Windows Phone

☞ Windows

macOS

▪

Box 2: DeliveryRuleIsDeviceConditionParameters

The DeliveryRuleIsDeviceCondition defines the IsDevice condition for the delivery rule. parameters defines the parameters for the condition.

Box 3: HTTP_USER_AGENT -

Incorrect Answers:

☞ The Pragma HTTP/1.0 general header is an implementation-specific header that may have various effects along the request-response chain. It is used for backwards compatibility with HTTP/1.0 caches.

☞ "X-Powered-By" is a common non-standard HTTP response header (most headers prefixed with an 'X-' are non-standard).

Box 4: DeliveryRuleRequestHeaderConditionParameters

DeliveryRuleRequestHeaderCondition defines the RequestHeader condition for the delivery rule. parameters defines the parameters for the condition.

Box 5: iOS -

The Require approved client app requirement only supports the iOS and Android for device platform condition.

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/concept-conditional-access-conditions

https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/concept-conditional-access-grant

Zidimirite Highly Voted 2 years, 6 months ago

First box: Mobile
https://docs.microsoft.com/en-us/python/api/azure-mgmt-cdn/azure.mgmt.cdn.models.isdevicematchconditionparametersmatchvaluesitem

Third box: iPhone. We are not looking for other Apple devices running iOS. It's the iPhone we're looking for.

upvoted 90 times

Netspud 1 year, 8 months ago

https://deviceatlas.com/blog/list-of-user-agent-strings#iphone
For 5 it has to be iPhone (this is not something Microsoft we are matching, but the details pulled from the device).
So I get:
1. DeliveryRuleDeviceConditionParameters
2. Mobile
(as below: https://docs.microsoft.com/en-us/python/api/azure-mgmt-cdn/azure.mgmt.cdn.models.isdevicematchconditionparametersmatchvaluesitem?view=azure-python)
3.DeliveryRequestHeaderConditionParameters
4. HTTP_USER_AGENT
5. iPhone

upvoted 17 times

Esward 9 months, 1 week ago

I agree with you!

upvoted 2 times

baroo1 8 months, 2 weeks ago

First box: Mobile as per documentation provided by this user.

Yup! Confirmed.
Third box: iPhone.
HTTP_USER_AGENT for iPhones is like the one below. iOS does not exist.
Mozilla/5.0 (iPhone; U; CPU like Mac OS X; en) AppleWebKit/420+ (KHTML, like Gecko) Version/3.0 Mobile/1A543a Safari/419.3

upvoted 3 times

titombo 2 years, 6 months ago

You are right, also the HTTP_USER_AGENT can be iPhone, iPad, etc... There isn't iOS. Even because iPhone and iPad I believe both could have iOS running in it.

upvoted 3 times

atomicicebreaker 2 years, 6 months ago

I agree, question is about CDN and answer explanation is all about Azure AD...

upvoted 3 times

kondapaturi Highly Voted 2 years, 4 months ago

1 DeliveryRuleDeviceConditionParameters - we are first checking for a device condition, hence we need to use the condition of DeliveryRuleDeviceConditionParameters
2.Mobile - The devices can be either Desktop or Mobile. These are the two accepted values. Here since we need to route requests based on mobile devices, we need to choose the value of Mobile.
3.DeliveryRequestHeaderConditionParameters. we need to understand the type of operating system running on the device. We can get this information from the request headers. Hence, we need to use the parameter of DeliveryRequestHeaderConditionParameters.
4.HTTP_USER_AGENT - we can check the HTTP_USER_AGENT property in the request header. In the user agent property of the request header, you will normally get information about the environment where the request is originating from. An example is given below where I am showing the request header from my own machine when I browse to a site.
5.iOS - we need to check the operating system which will be iOS.

upvoted 40 times

ivan0590 1 year, 6 months ago

I agree on the first four points, but the fifth point is incorrect.

The correct answer is iPhone.
Here you can see lots of iOS User Agent examples:
https://developers.whatismybrowser.com/useragents/explore/operating_system_name/ios/
"iPhone" appears all the time and "Mac OS X" is used instead of "iOS"

Example:
Mozilla/5.0 (iPhone; CPU iPhone OS 9_3_1 like Mac OS X) AppleWebKit/601.1.46 (KHTML, like Gecko) Version/9.0 Mobile/13E238 Safari/601.1

upvoted 6 times

Knightie 1 year, 2 months ago

Only this one is correct, the value only takes Desktop/Mobile, last one is iOS, Linux...etc

upvoted 1 times

Dixavado Most Recent 1 month ago

It was on my exam today (2023-09-26) I used the examtopics answer, but the last one I choose iPhone - score 850

upvoted 1 times

Dixavado 1 month ago

It was on my exam today (2023-09-26) I used the examtopics answer - score 850

upvoted 1 times

nikipediaa 1 month ago

Question was on exam 2023-09-26

upvoted 1 times

Saluk_DE 7 months ago

1. DeliveryRuleDeviceConditionParameters
2. Mobile

3.DeliveryRequestHeaderConditionParameters
4. HTTP_USER_AGENT
5. iPhone

Question was on exam 2023-03-30

upvoted 9 times

sarmaria 7 months, 2 weeks ago

Got this on 16/03/23. Chosen DeviceRule, Mobile, Header, HTTP Request, iOS. Got 846. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

mabdo 8 months, 3 weeks ago

was on exam 02/2023

upvoted 1 times

mabdo 9 months ago

on exam 02/23

upvoted 1 times

kociuba 10 months ago

got it on my exam 30-12-2022
score: 818

upvoted 1 times

vcfvct 1 year ago

Got this question on 10/21/2022

upvoted 3 times

CloudixExamTopics 11 months, 1 week ago

Answer ?

upvoted 1 times

gmishra88 1 year ago

Really, I should know this to be a developer? This is so specific, Microsoft

upvoted 9 times

DonH 4 months, 2 weeks ago

I agree this is insane.. how could they ask us to remember all this specific stuff.. I mean look at this table alone: https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.management.cdn.models?view=azure-dotnet
and that's only a single question.. :(

upvoted 1 times

Eltooth 1 year, 4 months ago

DeviceConditions
Mobile
HeaderConditions
HTTP
IOS

upvoted 4 times

tosm 1 year, 4 months ago

reference: https://docs.microsoft.com/en-us/azure/cdn/cdn-standard-rules-engine-match-conditions the filter rule for device may only accept the value of mobile and desktop. So mobile/iphone is my vote.

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer:
DeliveryRuleIsDeviceConditionParameters
Mobile
DeliveryRuleRequestHeaderConditionParameters
HTTP_USER_AGENT
iPhone

upvoted 7 times

michaelfastov 10 months, 3 weeks ago

Your link doesn't relate to this question as it describes Azure AD, not CDN.
https://docs.microsoft.com/en-us/python/api/azure-mgmt-cdn/azure.mgmt.cdn.models.isdevicematchconditionparametersmatchvaluesitem?view=azure-python

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose as below:
1. Mobile
2. DeliveryRuleIsDeviceConditionParameters
3. HTTP_USER_AGENT

4. DeliveryRuleRequestHeaderConditionParameters
5. Iphone

upvoted 5 times

szymson 1 year, 7 months ago

DeviceCondition box must be mobile or desktop:
https://docs.microsoft.com/en-us/azure/cdn/cdn-standard-rules-engine-match-conditions

upvoted 3 times

Question #19 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.
You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.
You need to configure authorization.
Solution:
☞ Configure and use Integrated Windows Authentication in the website.
☞ In the website, query Microsoft Graph API to load the groups to which the user is a member.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
Microsoft Graph is a RESTful web API that enables you to access Microsoft Cloud service resources.
Instead in the Azure AD application's manifest, set value of the groupMembershipClaims option to All. In the website, use the value of the groups claim from the
JWT for the user to determine permissions.
Reference:
https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/

Community vote distribution
B (78%) A (22%)

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: B - No

Reference:

https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups

https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication

https://docs.microsoft.com/en-us/graph/api/resources/azure-ad-overview?view=graph-rest-1.0
upvoted 36 times

lugospod 1 year, 9 months ago
I would say yes. Because if you have AD federation, then just by enabling WINDOWS authentication in your web app you gain access to Azure AD... so it is feasible
upvoted 1 times

mattvasc 1 year, 8 months ago
Authentication isn't the same as authorization, what about the roles? The answer is No.
upvoted 3 times

Tom87 Highly Voted 2 years, 6 months ago
I would say the solution meets the goal, according to this:
https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication
upvoted 10 times

coffecold 1 year ago
No, that is only for desktop or mobile application that runs on Windows. This is a web app in Azure.
upvoted 3 times

warchoon 8 months ago
The browser is such an application
upvoted 1 times

MysticalSam Most Recent 4 months, 3 weeks ago
This question was in today's exam on 10-June-2023

upvoted 5 times

hubekpeter 11 months, 1 week ago

Selected Answer: A

It's perfectly doable. Voting for A.

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: B

Correct Answer: B. No

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is the correct answer.
No.

upvoted 2 times

vavra 1 year, 4 months ago

I think it's B, because of this:
https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication
Federated users only, i.e. those created in an Active Directory and backed by Azure Active Directory. Users created directly in AAD, without AD backing - managed users - cannot use this auth flow. This limitation does not affect the Username/Password flow.

upvoted 2 times

minaritochuck 1 year, 5 months ago

Integrated Windows authentication is for "Intranet Application" (Although AD manages it) so don't use it on website

upvoted 2 times

minaritochuck 1 year, 7 months ago

It makes sense for app to use Integrated Windows Authentication (Implicit grant) for authentication and authorized by Microsoft Graph to access AD resources.
https://docs.microsoft.com/en-us/graph/api/resources/directoryrole?view=graph-rest-1.0

upvoted 1 times

EwanoE 1 year, 7 months ago

Selected Answer: B

Question says "website that will run as an Azure Web App" and there's no Integrated Windows Authentication in Azure Web App (aka App Service)

upvoted 2 times

Freidrich 1 year, 8 months ago

Selected Answer: B

The correct answer is B: No.

upvoted 1 times

mattdober 1 year, 9 months ago

Selected Answer: B

I think its B

upvoted 1 times

ScubaDiver123456 1 year, 9 months ago

Selected Answer: A

The answer is YES, assuming the application is hosted on a domain-joined VM. This document shows you how to set it up
https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication

upvoted 1 times

ScubaDiver123456 1 year, 9 months ago

Sorry. It does say it is an Azure Web App which can only run on Azure App Service. You can't join the underlying VMs to Active Directory. Therefore, I'm changing my vote to NO

upvoted 9 times

ning 2 years, 2 months ago

No, windows integrated authentication is NOT azure AD, you need app registration to get Azure AD

upvoted 7 times

kondapaturi 2 years, 4 months ago

Answer – No, Here you need to create an application in Azure AD. Then set the groupMembershipClaims claims. Then inspect the token in the application to see if the user is part of that group.

upvoted 5 times

Idkhow 2 years, 4 months ago

seriously what's the real answer??
in my opinion i think it's a yesn't

Configure and use Integrated Windows Authentication in the website.
(what does it have to do with the Integrated Windows Authentication)

In the website, query Microsoft Graph API to load the group to which the user is a member.
(like what the explanation said that it works)

upvoted 1 times

glam 2 years, 5 months ago

A. Yes

upvoted 1 times

Question #20 *Topic 4*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop Azure solutions.
You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.
You need to obtain an Azure Resource Manager access token.
Solution: Run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*
Get an access token using the VM's system-assigned managed identity and use it to call Azure Resource Manager
You will need to use PowerShell in this portion.
1. In the portal, navigate to Virtual Machines and go to your Windows virtual machine and in the Overview, click Connect.
2. Enter in your Username and Password for which you added when you created the Windows VM.
3. Now that you have created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.
4. Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource
Manager.
Example:
$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm

Community vote distribution
A (90%) 10%

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: Yes
Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.

I believe we can you Invoke-RestMethod too.

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm
upvoted 32 times

princesskay Highly Voted 3 years, 2 months ago
This does make sense since a web-request can be a rest call!
In the link below, it explains how web-request supports HTML content while Rest-Method supports JSON and XML content. The link provided for the solution is correct as well.
https://superuser.com/questions/1235349/what-is-the-difference-between-invoke-webrequest-and-invoke-restmethod
upvoted 25 times

carlosghosn Most Recent 11 months, 1 week ago
Got this in the exam today ! Nov 25, 2022
upvoted 2 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: A
Answer: Yes
upvoted 1 times

Mousavi 1 year, 1 month ago
Selected Answer: A
Absolutely the corrct answer is A
upvoted 2 times

Perplex 1 year, 3 months ago

Selected Answer: A

Answer A - The 2 most upvoted comments say this is correct and they are right.

upvoted 3 times

Azprep 1 year, 6 months ago

Answer is Yes

upvoted 2 times

SivajiTheBoss 1 year, 7 months ago

Selected Answer: B

Need to use Invoke-WebRequest cmdlet not Invoke-RestMethod
Correct Answer: B
Get an access token using the VM's system-assigned managed identity and use it to call Azure Resource Manager
You will need to use PowerShell in this portion.
1. In the portal, navigate to Virtual Machines and go to your Windows virtual machine and in the Overview, click Connect.
2. Enter in your Username and Password for which you added when you created the Windows VM.
3. Now that you have created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.
4. Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource
Manager.
Example:
$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm

upvoted 1 times

mabusalma 1 year, 9 months ago

Selected Answer: A

The answer is Yes correct
Invoke-WebRequest or Invoke-RestMethod will do the same job but the response in RestMethod will be much simpler.

upvoted 3 times

LauraGF 2 years, 1 month ago

The answer is no
Invoke-RestMethod: Grant your VM access to a secret stored in a Key Vault
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-nonaad

Invoke-WebRequest: Grant your VM access to a Resource Group in Azure Resource Manager
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm

upvoted 3 times

kondapaturi 2 years, 4 months ago

Answer is YES, you can make use of the PowerShell commands to make a request to managed identity service of the virtual machine. The Invoke-RestMethod is a generic call that is available in PowerShell for making a call to a web service.

upvoted 4 times

Spooky7 2 years, 5 months ago

In order to grant VM access to ARM you need to:
- assign VM user/system managed identity
- configure permission to ARM for that identity
- and then you can access token by Invoke-WebRequest cmdlet and use to to authenticate in ARM.
There is nothing in the question which indicates that those 2 previous steps are already done, so I would say it is not a correct answer.

upvoted 4 times

glam 2 years, 5 months ago

A. Yes

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

A. Yes
Invoke-RestMethod is merely a rapper commandlet around Invoke-WebRequest.

upvoted 2 times

ferut 2 years, 5 months ago

roflol... do you mean wrapper?

upvoted 2 times

businesskasper 2 years, 5 months ago

I'm pretty sure its a rapper cmdlet

upvoted 1 times

Santileo 2 years, 6 months ago

I code in powershell regularly and Invoke-RestMethod is the correct one to fetch these tokens.

upvoted 4 times

kwaazaar 2 years, 6 months ago

Where does it say the code is running on the VM itself?

upvoted 1 times

Gathix444 2 years, 9 months ago

Question 4 is "yes", Question 5 is "No". This is because the question ask you how to "grant a VM access to specific resource groups in Azure Resource Manager". Question 5 is only "yes" if you assigned the VM with role-base access control first. So Question 4 is more correct then Question 5

upvoted 2 times

brainwave15 2 years, 9 months ago

last line literally says "You need to obtain an Azure Resource Manager access token."
the only way to obtain is to use the invoke method...

upvoted 5 times

Question #21 *Topic 4*

HOTSPOT -

You are building a website to access project data related to teams within your organization. The website does not allow anonymous access.

Authentication is performed using an Azure Active Directory (Azure AD) app named internal.

The website has the following authentication requirements:

☞ Azure AD users must be able to login to the website.

☞ Personalization of the website must be based on membership in Active Directory groups.

You need to configure the application's manifest to meet the authentication requirements.

How should you configure the manifest? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

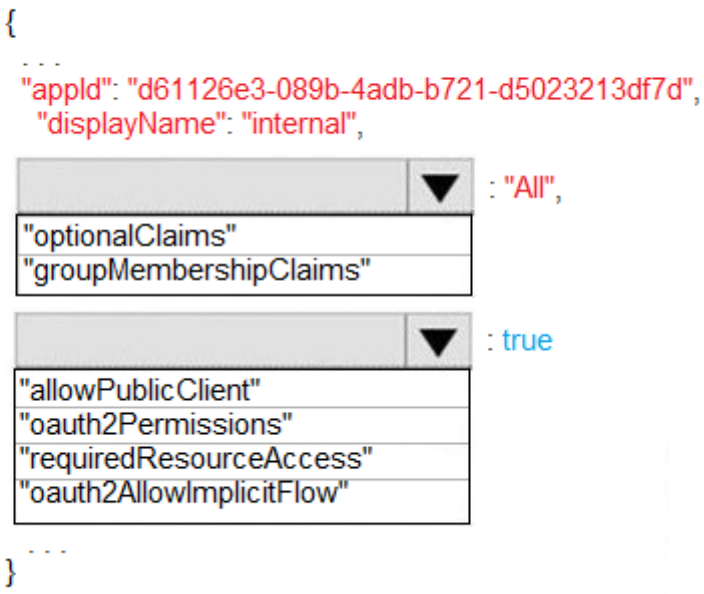

Correct Answer:

Answer Area

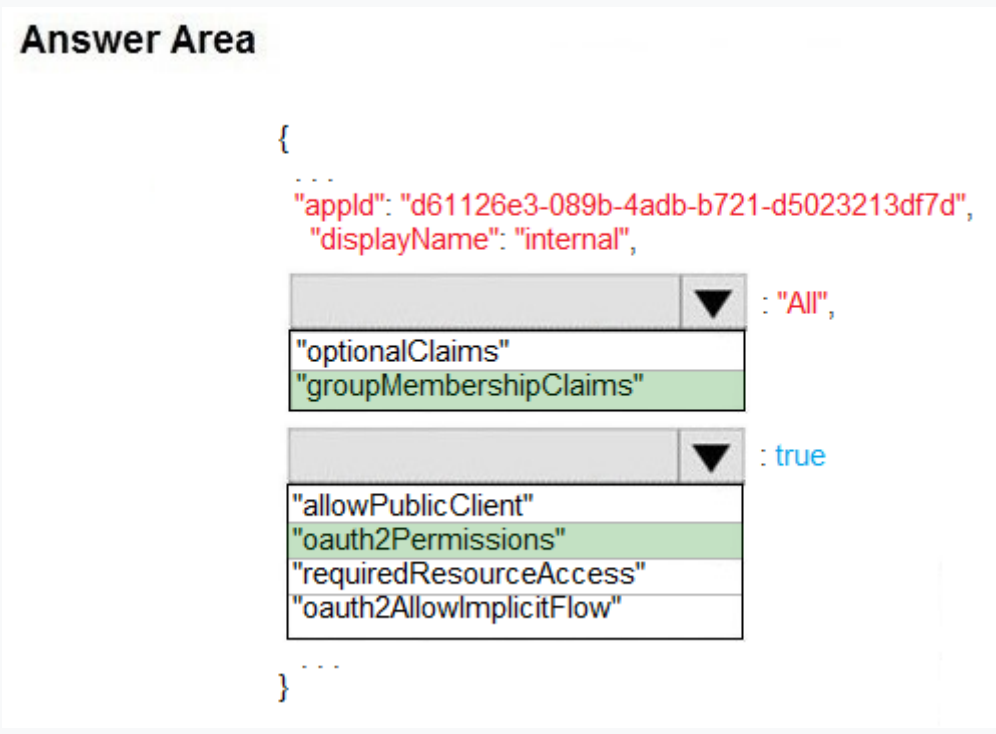

Box 1: groupMembershipClaims -

Scenario: Personalization of the website must be based on membership in Active Directory groups.

Group claims can also be configured in the Optional Claims section of the Application Manifest.

Enable group membership claims by changing the groupMembershipClaim

The valid values are:

"All"

"SecurityGroup"

"DistributionList"

"DirectoryRole"

Box 2: oauth2Permissions -

Scenario: Azure AD users must be able to login to the website. oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

Incorrect Answers:

oauth2AllowImplicitFlow. oauth2AllowImplicitFlow specifies whether this web app can request OAuth2.0 implicit flow access tokens. The

default is false. This flag is used for browser-based apps, like Javascript single-page apps.

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-fed-group-claims

AakashNeedsEmAll Highly Voted 2 years, 11 months ago

Second answer is incorrect. Well, oauth2Permissions can only accept collections value like an array not a boolean. It should be oauth2AllowImplicitFlow.

upvoted 166 times

cbn 2 years, 8 months ago

Agree.
oauth2permissions:
https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#oauth2permissions-attribute

oauth2allowimplicitflow:
https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#oauth2allowimplicitflow-attribute

upvoted 13 times

Asteriz 1 year, 9 months ago

Agreed

upvoted 1 times

cloud_exam1 2 years, 11 months ago

I think so. oauth2Permissions is a list. Should be oauth2AllowImplicitFlow.
https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest

upvoted 11 times

iiiihhhh 2 years, 10 months ago

The problem is that oauth2AllowImplicitFlow is for browser-based apps(SPA), what is not mentioned in the question.

upvoted 1 times

Tealon 2 years, 10 months ago

OAuth2 can also be used in this scenario. It does not matter if you use the normal grant or implicit flow.

upvoted 4 times

AfroYeti 2 years, 9 months ago

It is "You are building a WEBSITE to access project data related to teams within your organization."

I would imagine that, that website would run in a browser

upvoted 9 times

Esward 9 months, 1 week ago

Agreed

upvoted 1 times

Juanlu 2 years, 8 months ago

I Agree with you:
1) groupMembershipClaims
2) oauth2AllowImplicitFlow

upvoted 19 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: groupMembershipClaims
Personalization of the website must be based on membership in Active Directory groups.
Group claims can also be configured in the Optional Claims section of the Application Manifest. Enable group membership claims by changing the groupMembershipClaim
The valid values are:
- "All"
- "SecurityGroup"
- "DistributionList"
- "DirectoryRole"
Here we need to mention that we want to get the groups for the users. Hence we need to mention to set the groupMembershipClaims property to All.

Box 2: oauth2AllowImplicitFlow
Azure AD users must be able to login to the website.
auth2Permissions can only accept collections value like an array, not a boolean. oauth2AllowImplicitFlow accepts boolean value.
Here from the list of options given, if we want the application to fetch the required tokens , we would need to allow Implicit Flow.

upvoted 40 times

mlantonis 2 years, 5 months ago

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#groupmembershipclaims-attribute

https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#oauth2permissions-attribute

https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#oauth2allowimplicitflow-attribute

https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-fed-group-claims

https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-implicit-grant-flow

upvoted 8 times

richardc7 Most Recent 1 month ago

Got this question on 28/09/2023

upvoted 1 times

BaoNguyen2411 3 months ago

got this question on 06-29-2023

upvoted 3 times

MysticalSam 4 months, 3 weeks ago

This question was in today's exam on 10-June-2023

upvoted 4 times

adilkhan 7 months ago

ChatGPT
1) groupMembershipClaims
2) oauth2AllowImplicitFlow
{
"signInAudience": "AzureADMyOrg",
"groupMembershipClaims": "SecurityGroup",
"appId": "<app ID>",
"appRoles": [],
"availableToOtherTenants": false,
"displayName": "<app name>",
"oauth2AllowImplicitFlow": false,
"optionalClaims": {
"idToken": [],
"accessToken": []
},
"publisherDomain": "<your domain>",
"requiredResourceAccess": []
}

upvoted 2 times

carlosghosn 11 months, 1 week ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

DusanS 1 year ago

It must be oauth2AllowImplicitFlow because it is boll type in contrast to collection type. Look carefully picture.

upvoted 1 times

Eltooth 1 year, 4 months ago

1) groupMembershipClaims
2) oauth2AllowImplicitFlow

upvoted 3 times

Evo_Morales 1 year, 6 months ago

Check Point - how many folks knew the correct answer before reading/researching?

upvoted 1 times

ReniRechner 1 year, 8 months ago

According to
https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest

optionalClaims => "All" might not be allowed
groupMembershipClaims => "None", ... "All"

allowPublicClient => boolean
oauth2Permissions => collection
requiredResourceAccess => collection
ouath2AllowImplicitFlow => boolean

So first has to be "groupMembershipClaims" (can also be inferred from the requirements).

"allowPublicClient" seems to be meant for special cases not required here, there is only "oauth2AllowImplicitFlow" which makes sense for a website.

upvoted 3 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with:

1) groupMembershipClaims
2) oauth2AllowImplicitFlow

upvoted 3 times

MasterQuestMaster 1 year, 9 months ago

Got this on the exam!

upvoted 3 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 5 times

maharaju 1 year, 6 months ago

How come after writing exams you guys are visiting each questions and commenting ..SO much time

upvoted 8 times

paulomjeet 2 years, 2 months ago

2nd Answer should be oauth2AllowIdTokenImplicitFlow as it is a Boolean field. On the otherhand oauth2Permissions is an array of permission objects.

upvoted 3 times

ning 2 years, 2 months ago

oauth2Permissions has to be a collection, cannot be Boolean, and only valid is implicitflow

upvoted 2 times

weasel97 2 years, 3 months ago

The answer stated here seems correct implicit flow doesn't work for allcases. Reference: https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-implicit-grant-flow#:~:text=Prefer%20the%20auth,flow%20as%20well.

upvoted 1 times

Question #22 *Topic 4*

You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage containers named Container1 and Container2.
Uploading of videos occurs on an irregular basis.
You need to copy specific blobs from Container1 to Container2 when a new video is uploaded.
What should you do?

A. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API

B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet

C. Use AzCopy with the Snapshot switch to copy blobs to Container2

D. Download the blob to a virtual machine and then upload the blob to Container2

Correct Answer: *B*
The Start-AzureStorageBlobCopy cmdlet starts to copy a blob.

Example 1: Copy a named blob -
C:\PS>Start-AzureStorageBlobCopy -SrcBlob "ContosoPlanning2015" -DestContainer "ContosoArchives" -SrcContainer "ContosoUploads"
This command starts the copy operation of the blob named ContosoPlanning2015 from the container named ContosoUploads to the container named
ContosoArchives.
Reference:
https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps-6.13.0

Community vote distribution
B (100%)

AnonymousJhb Highly Voted 2 years, 10 months ago
B looks correct.
Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow. Asynchronous file uploads are a great fit for events. When changes are infrequent, but your scenario requires immediate responsiveness, event-based architecture can be especially efficient.
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview
upvoted 34 times

Esward 9 months, 1 week ago
I agree with you! B is correct answer
upvoted 1 times

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: B

The Start-AzureStorageBlobCopy cmdlet starts to copy a blob.

Reference:
https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps-6.13.0
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview
upvoted 15 times

Hogwarts Most Recent 6 months, 4 weeks ago
Got this in exam today (5 April 2023)
upvoted 4 times

Saluk_DE 7 months ago
Selected Answer: B
Question was on exam 2023-03-30
upvoted 1 times

NombreFalso 8 months, 1 week ago
Selected Answer: B
Only B is automated, since this happens on an irregular basis it needs to be automated
upvoted 2 times

OPT_001122 10 months, 3 weeks ago

Selected Answer: B

B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet

upvoted 1 times

coffecold 1 year ago

Selected Answer: B

I think the handler then should be a powershell function ?

upvoted 2 times

gmishra88 1 year ago

I do get it is the correct answer but the option looked very weird when I first read it: "B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet". Probably Microsoft guy wanted that effect. I would have said event-grid topic with an Start-AzureStorageBlobCopy event handler. Microsoft wins again.

upvoted 4 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is the correct answer.

upvoted 3 times

arindam1989 1 year, 6 months ago

Selected Answer: B

The automation of the process comes into play. Answer should be (B)

upvoted 1 times

AZ204Cert 1 year, 6 months ago

Got this on 04/05/22 (selected B)

upvoted 2 times

Mev4953 1 year, 9 months ago

Selected Answer: B

Correct Answer B

upvoted 2 times

ehurfheiz 1 year, 9 months ago

Selected Answer: B

B looks correct

upvoted 1 times

hems4all 2 years, 4 months ago

The Start-AzureStorageBlobCopy cmdlet starts to copy a blob

https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps-6.13.0

The Put Blob operation creates a new block, page, or append blob, or updates the content of an existing block blob. So in the given options, only Start-AzureStorageBlobCopy cmdlet copies the blob

upvoted 3 times

glam 2 years, 5 months ago

B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet

upvoted 3 times

faizalzain 2 years, 6 months ago

the answer is C in udemy

upvoted 3 times

AOE 2 years, 2 months ago

You dont have snapshots in the question so C is definitely wrong.

upvoted 3 times

sien 2 years, 6 months ago

I personally think it should be anwer A -> The Put Block operation creates a new block to be committed as part of a blob. -> https://docs.microsoft.com/en-us/rest/api/storageservices/put-block

Main reasoning is that answer B wirh the powershell doesn;t make sense

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

Put Blob is incorrect. It only "[c]reates a new blob or replaces an existing blob within a container." https://docs.microsoft.com/en-us/rest/api/storageservices/blob-service-rest-api

upvoted 1 times

clarionprogrammer 2 years, 6 months ago

It doesn't seem to imply that it can used for copying between containers.

upvoted 1 times

sien 2 years, 6 months ago

a copy between containers is nothing more that taking a blob from container A and then "Create" a blob in container B, isn't it?

upvoted 2 times

Question #23 *Topic 4*

You are developing an ASP.NET Core website that uses Azure FrontDoor. The website is used to build custom weather data sets for researchers. Data sets are downloaded by users as Comma Separated Value (CSV) files. The data is refreshed every 10 hours.
Specific files must be purged from the FrontDoor cache based upon Response Header values.
You need to purge individual assets from the Front Door cache.
Which type of cache purge should you use?

A. single path

B. wildcard

C. root domain

Correct Answer: *A*
These formats are supported in the lists of paths to purge:
☞ Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example, /
[1]
☞ Wildcard purge: Asterisk (*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /*, for example, /pictures/*.
☞ Root domain purge: Purge the root of the endpoint with "/" in the path.
Reference:
https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching

Community vote distribution
A (100%)

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: A

These formats are supported in the lists of paths to purge:

- Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example, /pictures/strasbourg.png;
- Wildcard purge: Asterisk (*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /*, for example, /pictures/*.
- Root domain purge: Purge the root of the endpoint with "/" in the path.

Here since we just need to purge individual assets, we can just make use of the Single path purge as the format for purging assets.

Reference:
https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching
upvoted 29 times

Tealon Highly Voted 2 years, 10 months ago
The given answer is correct.
upvoted 26 times

Esward Most Recent 9 months, 1 week ago
Given answer is correct!

https://learn.microsoft.com/en-us/azure/frontdoor/standard-premium/how-to-cache-purge
upvoted 1 times

Lucifer14 10 months ago
In 28-12-2022 exam
upvoted 3 times

capable 11 months ago
GOT THIS QUE IN EXAM 27/11/2022
upvoted 2 times

TheExamMaster2020 11 months, 2 weeks ago
Did my exam on 15th November 2022. This question was on it.
upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: A

A is correct answer.
Specific file = Single path

upvoted 4 times

[Removed] 1 year, 6 months ago

Selected Answer: A

Got this on 04/29/2022. Correct Answer: A

upvoted 3 times

AZ204Cert 1 year, 6 months ago

Got this on 04/05/22 (selected Single path)

upvoted 2 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 2 times

Manivannan19 1 year, 8 months ago

Is Azure FrontDoor part of this exam?

upvoted 3 times

coffecold 1 year ago

CDN is , Azure FrontDoor includes CDN. CDN has the caching thing.

upvoted 1 times

Freidrich 1 year, 8 months ago

Selected Answer: A

The proposed answer is correct A: Single path.

upvoted 1 times

ehurfheiz 1 year, 9 months ago

Selected Answer: A

A is the correct answer

upvoted 2 times

kondapaturi 2 years, 4 months ago

Single Path is correct , we just need to purge individual assets , we can just make use of the Single path purge as the format for purging assets.

upvoted 1 times

UnknowMan 2 years, 5 months ago

The given answer is correct.

upvoted 1 times

glam 2 years, 5 months ago

correct

upvoted 3 times

azahran 2 years, 9 months ago

individual assets so answer is correct

upvoted 3 times

Question #24 *Topic 4*

Your company is developing an Azure API.
You need to implement authentication for the Azure API. You have the following requirements:
All API calls must be secure.
▪
☞ Callers to the API must not send credentials to the API.
Which authentication mechanism should you use?

A. Basic

B. Anonymous

C. Managed identity

D. Client certificate

Correct Answer: *C*
Use the authentication-managed-identity policy to authenticate with a backend service using the managed identity of the API Management service. This policy essentially uses the managed identity to obtain an access token from Azure Active Directory for accessing the specified resource. After successfully obtaining the token, the policy will set the value of the token in the Authorization header using the Bearer scheme.
Reference:
https://docs.microsoft.com/bs-cyrl-ba/azure/api-management/api-management-authentication-policies

Community vote distribution
C (73%) D (27%)

MasDen Highly Voted 2 years, 12 months ago
If a caller outside Azure doesn't look like we can use Managed Identity in such a situation. I think the client certificate is better.
upvoted 33 times

coffecold 1 year ago
it says : You need to implement authentication for the Azure API.
So the Azure API is the one that needs to have an authenticated access to the backend services. That means Managed Identity.
upvoted 5 times

Knightie 1 year, 2 months ago
there is another similar question, using the cert, you will need to pass the cert or else store the cert in the key vault that you will still need to pass a value to retrieve the secret from key vault, so you will anyway pass a value to get it done. Managed ID doesn't need to pass a value like this.
upvoted 1 times

dluk 2 years, 11 months ago
Seems it's legit because question says "All API calls must be secure"
upvoted 2 times

IsildursHeir 2 years, 10 months ago
They do say "Azure API" so I am guessing they mean its an Azure provided API, not one wrapped around by Azure API management so MI would work..
upvoted 3 times

jay158 2 years, 4 months ago
Outsides user's can authenticate using Facebook /Google Identity providers.
Managed Identity is a good answer.
upvoted 4 times

ning 2 years, 2 months ago
Those are NOT managed identity, those are claim based authentication.
upvoted 6 times

luppittegui Highly Voted 2 years, 10 months ago
Ans is correct: managed identities eliminate the need for developers having to manage credentials by providing an identity for the Azure resource in Azure AD and using it to obtain Azure Active Directory (Azure AD) tokens.
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview
upvoted 29 times

Vady98 2 years ago
If a caller outside Azure doesn't look like we can use Managed Identity.
I think the client certificate is better.

upvoted 2 times

sarmaria Most Recent 7 months, 2 weeks ago

Got this on 16/03/23. Went with proposed solution. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 3 times

YLM007 5 months ago

have u gone through all the questions from exam topics?

upvoted 2 times

NombreFalso 8 months, 2 weeks ago

Selected Answer: C

C because it doesn't send credentials + API that's gonna call it is on Azure so Managed Identity is viable.

upvoted 2 times

Esward 9 months, 1 week ago

Given answer is correct as per MS docs

https://learn.microsoft.com/en-us/azure/api-management/authentication-managed-identity-policy

upvoted 1 times

gmishra88 1 year ago

It can never be managed identity unless the client is an azure resource that supports managed identity. The question does not say that. So, it has to be client-certificate. But then this is Microsoft exam and cannot expect a lot of logic from the question creators.

upvoted 2 times

bamosk 1 year, 3 months ago

Selected Answer: C

C is correct. A lot of people choose D (Client Certificate), but that is also sent to the server using the header X-ARR-ClientCert so that is not correct.

See here: https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth#access-client-certificate

upvoted 2 times

Eltooth 1 year, 4 months ago

C appears best answer based on question text, however if question is missing key words (API Management) then D would be best answer.

Remember the exam question will not be ambiguous nor prove to someone's photo memory of rewriting it here.

upvoted 3 times

tosm 1 year, 4 months ago

my read was: my company develop an Azure API, and need to authenticate the API to other (Azure) services. If this is the case. using managed ID is correct.

upvoted 3 times

suku555 1 year, 5 months ago

key words - developing an Azure API (obviously customization and there is no Azure service named 'Azure API') & Secure API calls --> These leads to the answer 'Client Certificate'.. Also in the Q there is no requirement of managing/rotating credential so surprising why many people vote for 'Managed Identity'..

upvoted 1 times

Azprep 1 year, 6 months ago

Managed identity

upvoted 1 times

iamstudying 1 year, 7 months ago

Selected Answer: D

Agree with MasDen

upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 1 times

DonOnur 1 year, 8 months ago

Selected Answer: C

Since there is no requirement for users outside of Azure, I would choose C.

upvoted 4 times

DonOnur 1 year, 8 months ago

Couldn't find the requirement that it has to be available to users outside azure.

upvoted 4 times

ytingyeu 1 year, 8 months ago

There is no service called "Azure API". Thus, I assume it refers to "App Service > API Apps".
https://azure.microsoft.com/en-us/services/app-service/api

In the page of "Secure app > Use managed identities", it's all about App itself accesses other resources, not about client calling API.
https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity

And the page "AuthN and AuthZ in Azure App Service ", it says "App Service uses federated identity", not managed identities.
https://docs.microsoft.com/en-us/azure/app-service/overview-authentication-authorization

So I suppose the only thing we can do is client certificate
https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth

upvoted 2 times

ytingyeu 1 year, 8 months ago

Plus, in page "Configure your App Service app to use AAD login" which is a part of federated identity, I don't see managed identities either.
https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-provider-aad
https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-provider-aad

upvoted 1 times

ytingyeu 1 year, 8 months ago

imo managed identities is for resource-to-resource not for client-to-resource

upvoted 3 times

ScubaDiver123456 1 year, 9 months ago

Selected Answer: D

Client Certificate as explained here:
https://docs.microsoft.com/en-us/learn/modules/explore-api-management/7-secure-access-api-certificates

upvoted 2 times

Question #25 *Topic 4*

You are a developer for a SaaS company that offers many web services.

All web services for the company must meet the following requirements:

- Use API Management to access the services
- Use OpenID Connect for authentication
- Prevent anonymous usage

A recent security audit found that several web services can be called without any authentication.

Which API Management policy should you implement?

A. jsonp

B. authentication-certificate

C. check-header

D. validate-jwt

Correct Answer: *D*

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

Incorrect Answers:

A: The jsonp policy adds JSON with padding (JSONP) support to an operation or an API to allow cross-domain calls from JavaScript browser-based clients.

JSONP is a method used in JavaScript programs to request data from a server in a different domain. JSONP bypasses the limitation enforced by most web browsers where access to web pages must be in the same domain.

JSONP - Adds JSON with padding (JSONP) support to an operation or an API to allow cross-domain calls from JavaScript browser-based clients.

Reference:

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

Community vote distribution

D (100%)

Cornholioz Highly Voted 2 years, 11 months ago

Correct. It's validate-jwt for Open ID, API, secure authentication.

upvoted 45 times

mlantonis Highly Voted 2 years, 5 months ago

Correct Answer: D

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

Reference:

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

upvoted 17 times

applepie Most Recent 3 months ago

got this question today, answer D - validate-jwt - 7/30/2023, score 895/1000

upvoted 3 times

Videira 8 months ago

On my exam 2023-02-25

upvoted 2 times

sam5678 1 year ago

correct

upvoted 1 times

serpevi 1 year, 1 month ago

Got this in 09/22 , went with the most voted answer, score 927.

upvoted 3 times

Eltooth 1 year, 4 months ago

Selected Answer: D

D is correct answer.

upvoted 3 times

JotaEleEfe 1 year, 7 months ago

Selected Answer: D

I think validate-jwt is the correct option.

upvoted 2 times

iamstudying 1 year, 7 months ago

Selected Answer: D

validate-jwt, buddies

upvoted 1 times

Freidrich 1 year, 8 months ago

Selected Answer: D

The proposed answer is correct D: validate-jwt.

upvoted 1 times

mabusalma 1 year, 9 months ago

Can someone let me understand why a check-header isn't a valid answer?

https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies#CheckHTTPHeader

Use the check-header policy to enforce that a request has a specified HTTP header. You can optionally check to see if the header has a specific value or check for a range of allowed values. If the check fails, the policy terminates request processing and returns the HTTP status code and error message specified by the policy.

upvoted 1 times

coffecold 1 year ago

JWT is a token issued by an identity provider, sent in the request, and that can be verified by the receiving site. A HTTP header is not issued by an identity provider.

upvoted 3 times

mcbc 2 years, 2 months ago

correct

upvoted 2 times

Peter304403 2 years, 4 months ago

The answer might be correct, but in the explanation of the incorrect answers it sounds very confusing to read jsonp described 3-times in a similar way? Since I don't know and I coldn't find anything: Could someone please explain the other two answers, what it is or where to find sothing about that?

upvoted 4 times

UnknowMan 2 years, 5 months ago

correct

upvoted 1 times

glam 2 years, 5 months ago

correct

upvoted 2 times

demius_8 2 years, 6 months ago

Answer is correct. Source: https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad

upvoted 3 times

pac1311 2 years, 9 months ago

Correctomundo!

upvoted 4 times

Question #26 *Topic 4*

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:

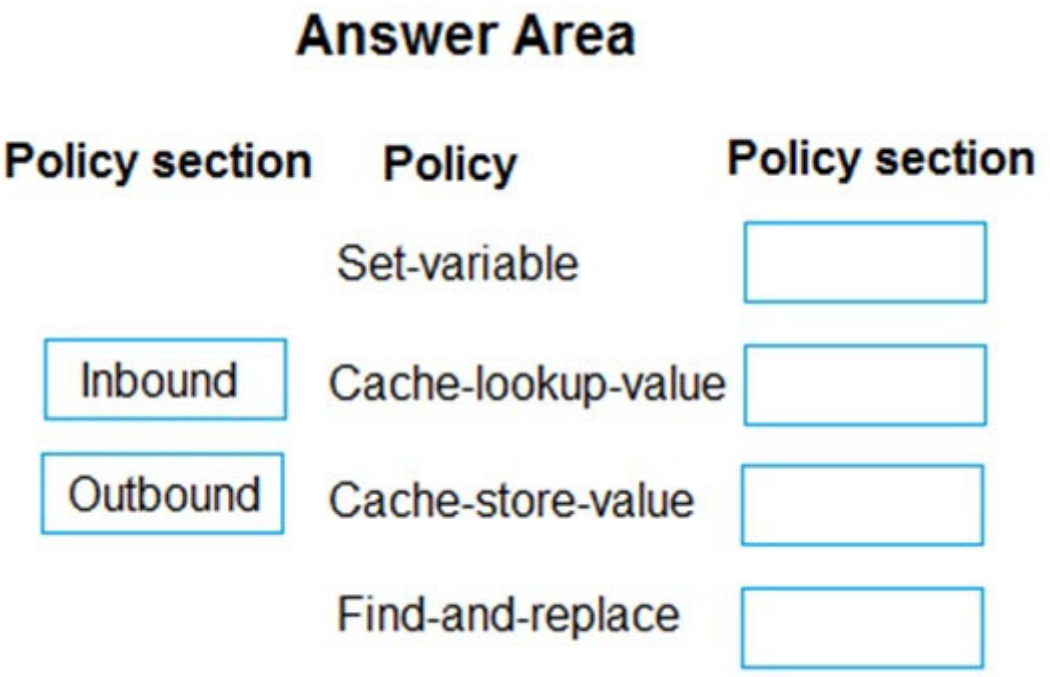

Correct Answer:

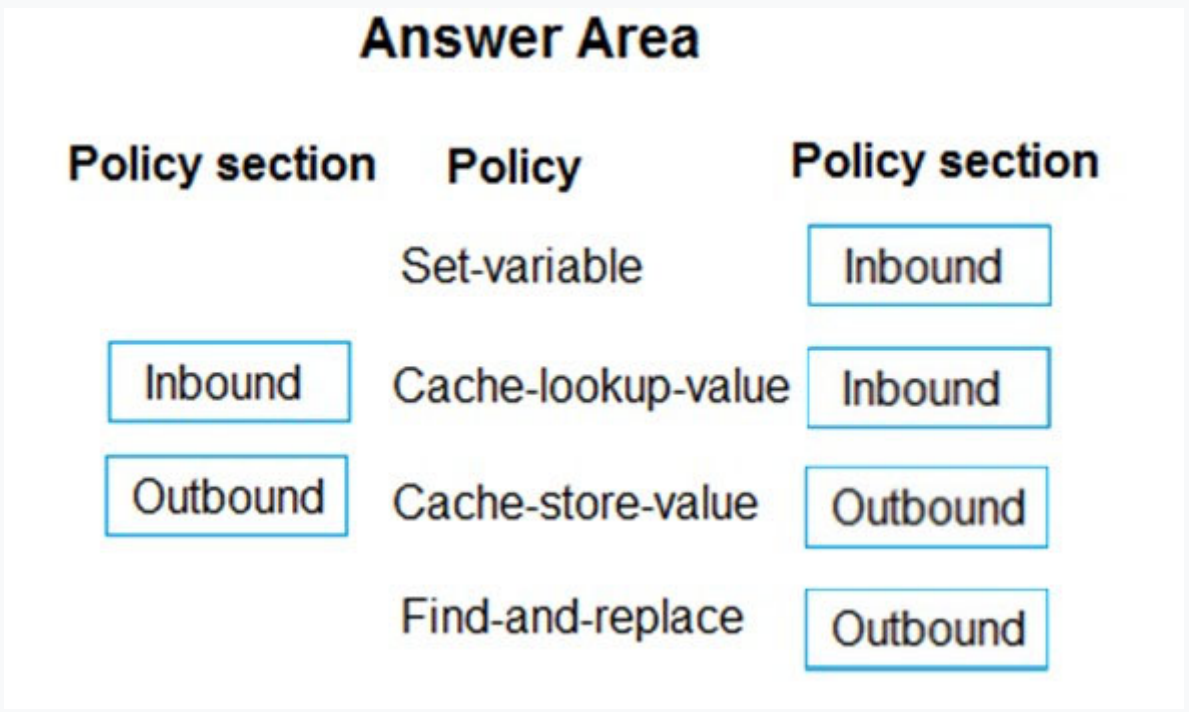

Box 1: Inbound.

A set-variable policy to store the detected user identity.

Example:

```
<policies>
<inbound>
<!-- How you determine user identity is application dependent -->
<set-variable
name="enduserid"
value="@(context.Request.Headers.GetValueOrDefault("Authorization","").Split(' ')[1].AsJwt()?.Subject)" />
```

Box 2: Inbound -

A cache-lookup-value policy -

Example:

<inbound>
<base />
<cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false" downstream-caching-type="none | private | public" must-revalidate="true | false">
<vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several times -->
</cache-lookup>
</inbound>

Box 3: Outbound -
A cache-store-value policy.
Example:
<outbound>
<base />
<cache-store duration="3600" />
</outbound>

Box 4: Outbound -
A find-and-replace policy to update the response body with the user profile information.
Example:
<outbound>
<!-- Update response body with user profile-->
<find-and-replace
from="'$userprofile$'"
to="@((string)context.Variables["userprofile"])" />
<base />
</outbound>
Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key

Ave Highly Voted 2 years, 11 months ago

Inbound
Inbound
Inbound
Outbound

upvoted 164 times

NiceGuyAlberto 1 year, 11 months ago

correct!
-> https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#fragment-caching

upvoted 8 times

[Removed] 3 months, 2 weeks ago

Incorrect. It should be Inbound, Inbound, Outbound, Outbound

upvoted 1 times

ks321 3 months, 1 week ago

Here is the complete policy in MS Documentation:
https://learn.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#:~:text=Once%20you%20combine%20these%20steps%2C%20the%20end%20result%20is%20a%20policy%20that%20looks%20like%20the%20following%20one.

upvoted 2 times

dbobspurfpoo 1 year, 10 months ago

Is this how i should answer on the exam? Is it just this site that's wrong?

upvoted 2 times

rgullini 6 months, 3 weeks ago

Newbie, you have the questions, don´t expect to have everything so easy, at least you need to understand what you are answering. The answers usually are wrong in this site so you need to understand the topic.

upvoted 4 times

taupokk Highly Voted 2 years, 11 months ago

The answer provided is correct.

upvoted 36 times

robertob 2 years, 11 months ago

yes is correct: <policies>
<inbound>
<base />
<cache-lookup vary-by-developer="false" vary-by-developer-groups="false" downstream-caching-type="none" must-revalidate="true" caching-type="internal" >
<vary-by-query-parameter>version</vary-by-query-parameter>
</cache-lookup>
</inbound>
<outbound>
<cache-store duration="seconds" />
<base />
</outbound>
</policies>

upvoted 8 times

cbn 2 years, 8 months ago

This is not cache-store-value. However the answer seems correct as per @profesorklaus answer below.

upvoted 4 times

JH81 Most Recent 4 months ago

Got this on 6/28/2023 and passed with 850. Went with highly voted answer.

upvoted 3 times

NombreFalso 8 months, 2 weeks ago

>> The caching mechanism must detect the user ID of the client that accesses data for a given location and *cache the response for that user ID.*

Thus cache store should be outbound guys

upvoted 3 times

Esward 9 months, 1 week ago

Given answers are correct.
Inbound
Inbound
outbound->in general the response come in oubound section
outbound
https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-cache

upvoted 4 times

kkop 11 months ago

Here is Answer:https://learn.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#architecture

inbound,
inbound,
inbound
outbound

upvoted 4 times

elequiel 1 year ago

Got it in exam 20/10/2022

upvoted 3 times

Knightie 1 year, 2 months ago

see the xml tag of inbound and outbound
https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key

upvoted 1 times

Lucky_me 1 year, 4 months ago

<policies>
<inbound>
<base />
<cache-lookup vary-by-developer="false" vary-by-developer-groups="false" downstream-caching-type="none" must-revalidate="true" caching-type="internal" >
<vary-by-query-parameter>version</vary-by-query-parameter>
</cache-lookup>
</inbound>
<outbound>
<cache-store duration="seconds" />
<base />
</outbound>
</policies>

upvoted 1 times

vavra 1 year, 4 months ago

I think it's B, because of this:
https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication
Federated users only, i.e. those created in an Active Directory and backed by Azure Active Directory. Users created directly in AAD, without AD backing - managed users - cannot use this auth flow. This limitation does not affect the Username/Password flow.

upvoted 1 times

kjfdzkkbsm 1 year, 4 months ago

Your'e at the wrong question, mate

upvoted 8 times

Azprep 1 year, 6 months ago

Inbound
Inbound
Inbound
Outbound

upvoted 3 times

Prasu69 1 year, 7 months ago

Inbound
Inbound
Inbound
Outbound

upvoted 2 times

[Removed] 3 months, 2 weeks ago

Incorrect

upvoted 1 times

petitbilly 1 year, 7 months ago

According to this docs from Microsoft, the third box (cache-store-value) could be: inbound or outbound.
https://docs.microsoft.com/en-us/azure/app-service/faq-app-service-linux#i-m-using-my-own-custom-container--i-want-the-platform-to-mount-an-smb-share-to-the---home---directory-

upvoted 1 times

prastogi 1 year, 9 months ago

https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key

Answer should be
Inbound
Inbound
Inbound
Outbound

upvoted 3 times

[Removed] 3 months, 2 weeks ago

Incorrect

upvoted 1 times

AzureLearning 1 year, 9 months ago

Why everyone is saying Set-Variable is Inbound ? Microsoft site says it can be used in both inbound and outbound.

https://docs.microsoft.com/en-us/azure/api-management/api-management-advanced-policies#usage-12

Set-variable policy can be used in
Policy sections: inbound, outbound, backend, on-error

upvoted 4 times

DonOnur 1 year, 8 months ago

although it can be used in inbound and outbound it only makes sense as inbound

upvoted 2 times

tramlong888 1 year, 10 months ago

cache-store-value should be Outbound
Refer: https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key

upvoted 2 times

Yazhu 1 year, 10 months ago

Refer Link : https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#fragment-caching
1. <Inbound> <set-variable/><cache-lookup-value/><<cache-store-value/></Inbound>
2. <Outbound><find-and-replace/></Outbound>

upvoted 5 times

Question #27 *Topic 4*

DRAG DROP -

You are developing an Azure solution.

You need to develop code to access a secret stored in Azure Key Vault.

How should you complete the code segment? To answer, drag the appropriate code segments to the correct location. Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Correct Answer:

Box 1: SecretClient -

Box 2: DefaultAzureCredential -

In below example, the name of your key vault is expanded to the key vault URI, in the format "https://<your-key-vault-name>.vault.azure.net". This example is using 'DefaultAzureCredential()' class from Azure Identity Library, which allows to use the same code across different environments with different options to provide identity. string keyVaultName = Environment.GetEnvironmentVariable("KEY_VAULT_NAME"); var kvUri = "https://" + keyVaultName + ".vault.azure.net"; var client = new SecretClient(new Uri(kvUri), new DefaultAzureCredential());

Reference:

https://docs.microsoft.com/en-us/azure/key-vault/secrets/quick-create-net

mlantonis Highly Voted 2 years, 5 months ago

Answer is correct.

string keyVaultName = Environment.GetEnvironmentVariable("KEY_VAULT_NAME");
var kvUri = "https://" + keyVaultName + ".vault.azure.net";
var client = new SecretClient(new Uri(kvUri), new DefaultAzureCredential());

Box 1: SecretClient
Box 2: DefaultAzureCredential

Reference:
https://docs.microsoft.com/en-us/azure/key-vault/secrets/quick-create-net

upvoted 72 times

Esward 9 months, 1 week ago

Agreed

upvoted 1 times

oescm Highly Voted 1 year, 8 months ago

Got this one 02/2022. Went with most voted

upvoted 9 times

JH81 Most Recent 4 months ago

Got this on 6/28/2023 and passed with 850. Went with answer.

upvoted 2 times

Saluk_DE 7 months ago

Question was on exam 2023-03-30

upvoted 1 times

mabdo 9 months ago

on exam 02/23

upvoted 1 times

kociuba 10 months ago

got it on my exam 30-12-2022
score: 818

upvoted 3 times

adilkhan 9 months, 2 weeks ago

are all these questions enough to pass the exam?

upvoted 2 times

OPT_001122 10 months, 3 weeks ago

1: SecretClient
2: DefaultAzureCredential

upvoted 2 times

vcfvct 1 year ago

Had this question on 10/21/22

upvoted 2 times

Eltooth 1 year, 4 months ago

Secret
Default

upvoted 2 times

edengoforit 1 year, 10 months ago

DefaultAzureCredential
DefaultAzureCredential is appropriate for most scenarios where the application is intended to run in the Azure Cloud. This is because the DefaultAzureCredential determines the appropriate credential type based of the environment it is executing in. It supports authenticating both as a service principal or managed identity, and can be configured so that it will work both in a local development environment or when deployed to the cloud.

The DefaultAzureCredential will first attempt to authenticate using credentials provided in the environment. In a development environment you can authenticate as a service principal with the DefaultAzureCredential by providing configuration in environment variables as described in the next section.

If the environment configuration is not present or incomplete, the DefaultAzureCredential will then determine if a managed identity is available in the current environment. Authenticating as a managed identity requires no configuration, but does require platform support. See the managed identity documentation for more details on this.

upvoted 8 times

ADJ85 2 years, 2 months ago

Correct

upvoted 4 times

Question #28 *Topic 4*

You are developing an Azure App Service REST API.
The API must be called by an Azure App Service web app. The API must retrieve and update user profile information stored in Azure Active Directory (Azure AD).
You need to configure the API to make the updates.
Which two tools should you use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Microsoft Graph API

B. Microsoft Authentication Library (MSAL)

C. Azure API Management

D. Microsoft Azure Security Center

E. Microsoft Azure Key Vault SDK

Correct Answer: *AC*
A: You can use the Azure AD REST APIs in Microsoft Graph to create unique workflows between Azure AD resources and third-party services. Enterprise developers use Microsoft Graph to integrate Azure AD identity management and other services to automate administrative workflows, such as employee onboarding (and termination), profile maintenance, license deployment, and more.
C: API Management (APIM) is a way to create consistent and modern API gateways for existing back-end services.
API Management helps organizations publish APIs to external, partner, and internal developers to unlock the potential of their data and services.
Reference:
https://docs.microsoft.com/en-us/graph/azuread-identity-access-management-concept-overview

Community vote distribution
AB (100%)

Cholo981 Highly Voted 2 years, 4 months ago
The Answer is Graph API and MSAL.
API managment as nothing to do with the developing of the REST API...
upvoted 98 times

Esward 9 months, 1 week ago
Agreed.
https://learn.microsoft.com/en-us/graph/auth-cloudsolutionprovider?tabs=azuread
upvoted 1 times

AndrewJB 1 year, 11 months ago
Correct. To access Graph API user must be logged in, so MSAL is required. APIM is optional and not mandatory for this scenario
upvoted 7 times

ReniRechner 1 year, 8 months ago
Furthermore the API is not intended for public use but internal use by another web app. So there is no need for APIM.
upvoted 5 times

pazetox 1 year, 3 months ago
MSAL can be used to provide secure access to Microsoft Graph.
https://docs.microsoft.com/en-us/azure/active-directory/develop/msal-overview
upvoted 1 times

mlantonis Highly Voted 2 years, 5 months ago
Answer is correct.

A: Microsoft Graph to retrieve and update user profile info from Azure AD.

C: API Management (APIM)

Reference:

https://docs.microsoft.com/en-us/graph/azuread-identity-access-management-concept-overview

https://docs.microsoft.com/en-us/graph/overview

https://docs.microsoft.com/en-us/azure/api-management/api-management-key-concepts
upvoted 21 times

cooyoote 2 years, 1 month ago
So wrong, API Management is supossed to create API gateways and rules to manage API calls
upvoted 8 times

ensa 2 years ago
Fiest time see you are wrong.
AAAAAAAAABBBBBBBBBB
upvoted 8 times

Dixavado Most Recent 1 month ago
Selected Answer: AB
It was on my exam today (2023-09-26) - score 850
upvoted 1 times

nikipediaa 1 month ago
Question was on exam 2023-09-26. The order of answers was changed.
upvoted 1 times

deathRac3 5 months, 1 week ago
Selected Answer: AB
trythis
upvoted 1 times

Hogwarts 6 months, 4 weeks ago
Got this in exam today (5 April 2023)
upvoted 4 times

NombreFalso 8 months, 2 weeks ago
Selected Answer: AB
Choosing A implies using B too. APIM is totally unrelated.
upvoted 3 times

Jeff8888 9 months ago
Graph Api, Azure APIM
upvoted 1 times

Michael2023 9 months ago
"The API must retrieve and update user profile information stored in Azure Active Directory (Azure AD).
You need to configure the API to make the updates."

this is a hint that you want programmatically update the user data stored in Azure AD...
so Graph API and MSAL is the right answer :)
upvoted 1 times

serpevi 1 year, 1 month ago
Got this in 09/22 , went with MSAL and Graph, score 927.
upvoted 9 times

Azprep 1 year, 6 months ago
Graph Api, Azure APIM
upvoted 2 times

svsv22 1 year, 6 months ago
I think answer should be A and B. MSAL is needed to get the tokens and pass it in Graph API request.
upvoted 1 times

iamstudying 1 year, 7 months ago
Selected Answer: AB
AB, buddies. 981Cholo, AndrweJB and ReniRechner are the mvps
upvoted 3 times

karthik0328 1 year, 8 months ago
Selected Answer: AB
Why is API management required for this?
upvoted 1 times

Chiboy 1 year, 8 months ago
You don't need MSAL to get access to Graph API when you are a web app using azure app service in Azure. The given answer is correct
upvoted 3 times

ehurfheiz 1 year, 9 months ago

Selected Answer: AB

Correct answer seems to be AB

upvoted 1 times

vtomy 2 years, 3 months ago

As per the question it is mentioned 'Configure'. MSAL is for process, nothing related to configuere. I guess API management could be correct.

upvoted 3 times

Ganesh1983 1 year, 5 months ago

yes, that makes sense, since the question specifically talks about configure api.

upvoted 1 times

Question #29 *Topic 4*

You develop a REST API. You implement a user delegation SAS token to communicate with Azure Blob storage.
The token is compromised.
You need to revoke the token.
What are two possible ways to achieve this goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Revoke the delegation key.

B. Delete the stored access policy.

C. Regenerate the account key.

D. Remove the role assignment for the security principle.

Correct Answer: *AB*
A: Revoke a user delegation SAS -
To revoke a user delegation SAS from the Azure CLI, call the az storage account revoke-delegation-keys command. This command revokes all of the user delegation keys associated with the specified storage account. Any shared access signatures associated with those keys are invalidated.
B: To revoke a stored access policy, you can either delete it, or rename it by changing the signed identifier. Changing the signed identifier breaks the associations between any existing signatures and the stored access policy. Deleting or renaming the stored access policy immediately effects all of the shared access signatures associated with it.
Reference:
https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/storage/blobs/storage-blob-user-delegation-sas-create-cli.md
https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy#modifying-or-revoking-a-stored-access-policy

Community vote distribution
AD (100%)

jay158 Highly Voted 2 years, 4 months ago
Given answer is incorrect.
Correct answers are A , D
https://docs.microsoft.com/en-us/rest/api/storageservices/create-user-delegation-sas#revoke-a-user-delegation-sas
upvoted 102 times

Esward 9 months, 1 week ago
I greed with A, D
upvoted 2 times

OPT_001122 1 year, 2 months ago
Thanks https://www.examtopics.com/user/jay158/
no doubt after referring the ink that it is A and D
upvoted 5 times

jkes80 2 years, 3 months ago
Your link literally says you're correct it should be A and D.
"If you believe that a SAS has been compromised, then you should revoke the SAS. You can revoke a user delegation SAS either by revoking the user delegation key, or by changing or removing RBAC role assignments for the security principal used to create the SAS."
upvoted 19 times

maukaba 1 year, 10 months ago
There're two ways to create a SAS:
(1). The "standard" way to generate a SAS token is to use the storage account key.
(2). by using "managed identities" with a technique is called a "user delegation" SAS, and it allows you to sign the signature with Azure AD credentials instead of with the storage account key.
This question is (2) hence A, D is correct
REF: https://markheath.net/post/user-delegation-sas
upvoted 3 times

mlantonis Highly Voted 2 years, 5 months ago
Answer seems correct.

A: Use az storage account revoke-delegation-keys command.
B: To revoke a stored access policy, you can either delete it, or rename it by changing the signed identifier.

Reference:

https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/storage/blobs/storage-blob-user-delegation-sas-create-cli.md

https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy#modifying-or-revoking-a-stored-access-policy

https://docs.microsoft.com/en-us/cli/azure/storage/account?view=azure-cli-latest#az_storage_account_revoke_delegation_keys
upvoted 16 times

Jurgen1234 2 years, 4 months ago
Incorrect it should be D not B, see https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy

Stored access policies are not supported for the user delegation SAS or the account SAS..
upvoted 3 times

Azprep 1 year, 6 months ago
Answers should be A& D
upvoted 1 times

nikipediaa Most Recent 1 month ago
Question was on exam 2023-09-26
upvoted 1 times

Teit_examptopics 1 month ago
If you believe that a SAS has been compromised, you should revoke it. You can revoke a user delegation SAS either by revoking the user delegation key, or by changing or removing RBAC role assignments for the security principal that's used to create the SAS.

https://learn.microsoft.com/en-us/rest/api/storageservices/create-user-delegation-sas
upvoted 1 times

IvanIco 2 months, 1 week ago
C By regenerating the you make the prevues key unusable why isn't that right
upvoted 1 times

RKyawal 4 months ago
https://learn.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy
cannot be B as per this link..not supporeted
upvoted 1 times

DP_Bhatt 9 months, 1 week ago
Selected Answer: AD
Correct answers A and D
upvoted 1 times

rosa_0617 1 year, 2 months ago
Selected Answer: AD
If you believe that a SAS has been compromised, then you should revoke the SAS.
You can revoke a user delegation SAS either by revoking the user delegation key, or by changing or removing RBAC role assignments for the security principal used to create the SAS.
upvoted 3 times

nuiowmokka 1 year, 2 months ago
Can we pass the exam just by studying questions tilll Page 22? I don't have contributor access and can't afford right now.
upvoted 5 times

ASAAAAS2 1 year, 3 months ago
Selected Answer: AD
A & D for sure
upvoted 2 times

Eltooth 1 year, 4 months ago
Selected Answer: AD
A and D are correct answers.
upvoted 2 times

emetos05 1 year, 4 months ago
AD is correct
upvoted 1 times

Azprep 1 year, 6 months ago
Correct answers are A& B
upvoted 1 times

Prasu69 1 year, 7 months ago
Selected Answer: AD
https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy

upvoted 1 times

Prasu69 1 year, 7 months ago

A and D
Stored access policies are not supported for the user delegation SAS or the account SAS..
https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy

upvoted 2 times

avanip 1 year, 7 months ago

Selected Answer: AD

A and D are correct

upvoted 1 times

iamstudying 1 year, 7 months ago

Selected Answer: AD

AD, buddies. @158jay is correct. shout out to maukabababa too

upvoted 1 times

Question #30 *Topic 4*

DRAG DROP -

You are developing an Azure-hosted application that must use an on-premises hardware security module (HSM) key.

The key must be transferred to your existing Azure Key Vault by using the Bring Your Own Key (BYOK) process.

You need to securely transfer the key to Azure Key Vault.

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

Select and Place:

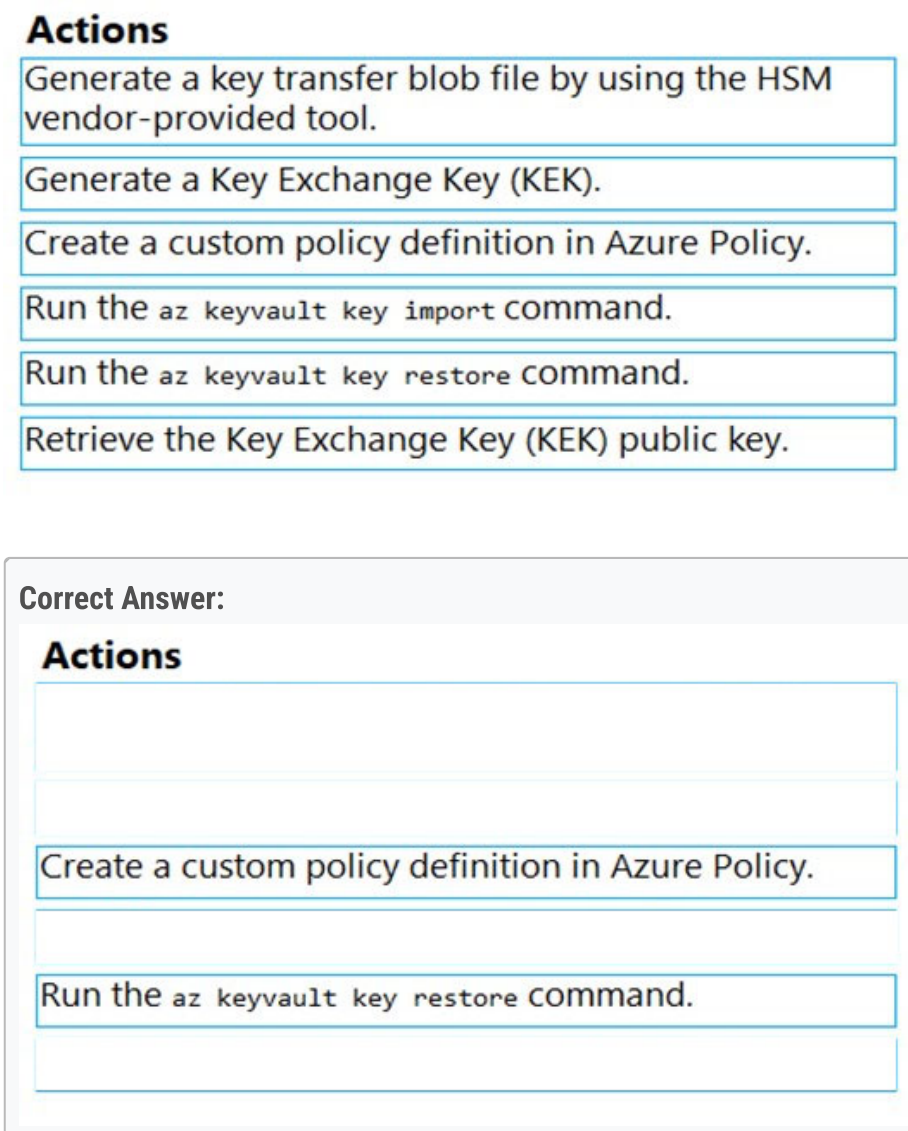

Correct Answer:

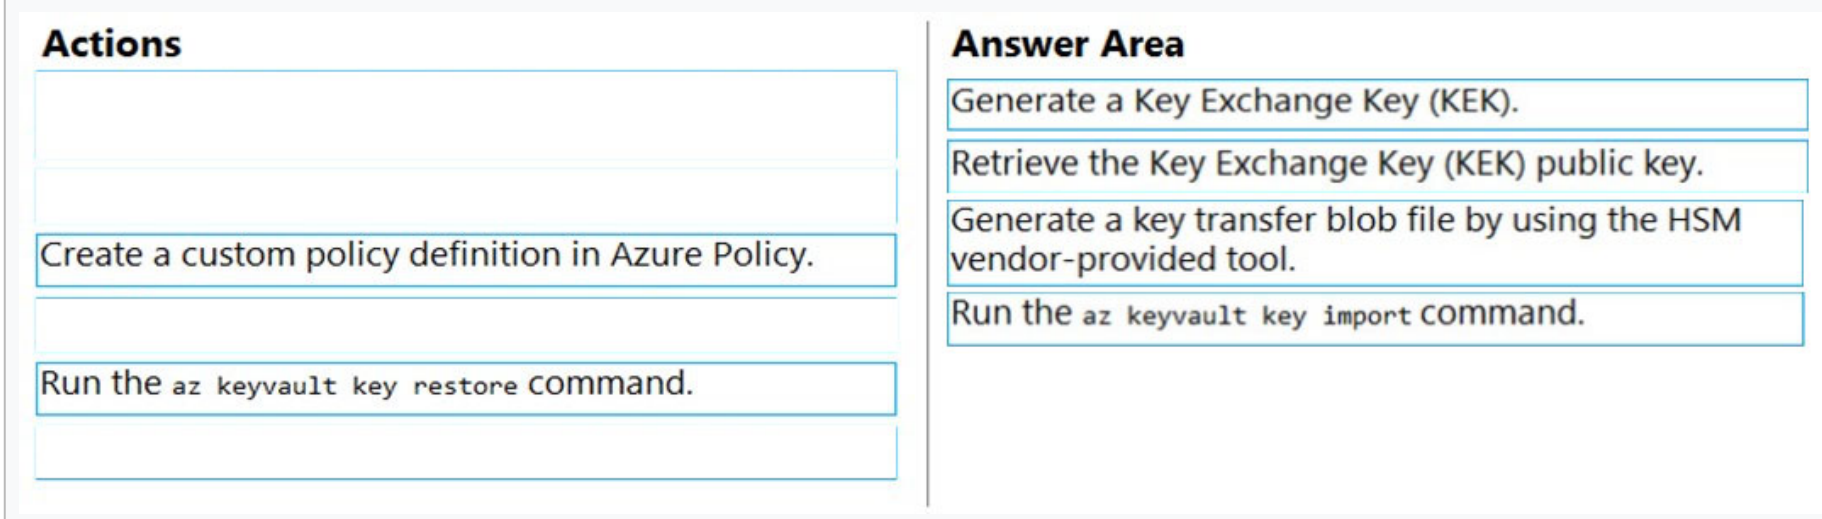

To perform a key transfer, a user performs following steps:

- Generate KEK.
- Retrieve the public key of the KEK.
- Using HSM vendor provided BYOK tool - Import the KEK into the target HSM and exports the Target Key protected by the KEK.
- Import the protected Target Key to Azure Key Vault.

Step 1: Generate a Key Exchange Key (KEK).

Step 2: Retrieve the Key Exchange Key (KEK) public key.

Step 3: Generate a key transfer blob file by using the HSM vendor-provided tool.

Generate key transfer blob using HSM vendor provided BYOK tool

Step 4: Run the az keyvault key import command

Upload key transfer blob to import HSM-key.

Customer will transfer the Key Transfer Blob (".byok" file) to an online workstation and then run a az keyvault key import command to import this blob as a new

HSM-backed key into Key Vault.

To import an RSA key use this command:

az keyvault key import

Reference:

https://docs.microsoft.com/en-us/azure/key-vault/keys/byok-specification

aradice Highly Voted 2 years, 4 months ago

correct: https://docs.microsoft.com/en-us/azure/key-vault/keys/byok-specification#user-steps

upvoted 46 times

Esward 9 months, 1 week ago

Agreed

upvoted 1 times

Esward 9 months, 1 week ago

Given answers are correct

upvoted 1 times

TakumaK 2 years, 3 months ago

agree with you!

upvoted 4 times

MasterQuestMaster Highly Voted 1 year, 9 months ago

Got this on the exam. :)
top kek

upvoted 20 times

mandynotmandy 1 year, 7 months ago

this is still showing up in feb 2022 exams, the kek definitely made this question unforgettable for me
top kek

upvoted 6 times

mewan Most Recent 5 months, 1 week ago

is it enough to pass the exam to read up to 22 pages

upvoted 6 times

Videira 8 months ago

On my exam 2023-02-25

upvoted 2 times

NombreFalso 8 months, 2 weeks ago

based kek question

upvoted 2 times

Jeff8888 9 months ago

Agreed

upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 4 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with most voted

upvoted 6 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with originally proposed solution

upvoted 6 times

danila16030 1 year, 9 months ago

Ho many questions from you test you find on this site? Just interesting)

upvoted 1 times

mcbc 2 years, 2 months ago

Generate KEK.
Retrieve the public key of the KEK.
Using HSM vendor provided BYOK tool - Import the KEK into the target HSM and exports the Target Key protected by the KEK.
Import the protected Target Key to Azure Key Vault.

upvoted 2 times

Question #31 *Topic 4*

You develop and deploy an Azure Logic app that calls an Azure Function app. The Azure Function app includes an OpenAPI (Swagger) definition and uses an
Azure Blob storage account. All resources are secured by using Azure Active Directory (Azure AD).
The Azure Logic app must securely access the Azure Blob storage account. Azure AD resources must remain if the Azure Logic app is deleted.
You need to secure the Azure Logic app.
What should you do?

A. Create a user-assigned managed identity and assign role-based access controls.

B. Create an Azure AD custom role and assign the role to the Azure Blob storage account.

C. Create an Azure Key Vault and issue a client certificate.

D. Create a system-assigned managed identity and issue a client certificate.

E. Create an Azure AD custom role and assign role-based access controls.

Correct Answer: *A*
To give a managed identity access to an Azure resource, you need to add a role to the target resource for that identity.
Note: To easily authenticate access to other resources that are protected by Azure Active Directory (Azure AD) without having to sign in and provide credentials or secrets, your logic app can use a managed identity (formerly known as Managed Service Identity or MSI). Azure manages this identity for you and helps secure your credentials because you don't have to provide or rotate secrets.
If you set up your logic app to use the system-assigned identity or a manually created, user-assigned identity, the function in your logic app can also use that same identity for authentication.
Reference:
https://docs.microsoft.com/en-us/azure/logic-apps/create-managed-service-identity https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates-for-clients

Community vote distribution
A (100%)

aradice Highly Voted 2 years, 4 months ago
correct "Azure AD resources must remain if the Azure Logic app is deleted."
upvoted 45 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000.
upvoted 6 times

debanjan10 Highly Voted 2 years ago
User-assigned keys are individual components. Even if the logic apps are deleted, the keys remain. But in case of system-assigned keys, those are auto generated and are deleted when the Azure resources themselves are deleted.
upvoted 21 times

applepie Most Recent 3 months ago
got this question today, answer A - 7/30/2023, score 895/1000
upvoted 3 times

databasejamdown 5 months, 2 weeks ago
Selected Answer: A
User assigned will persist after resource is removed
upvoted 4 times

RaghavMGupta 6 months, 3 weeks ago
Why are other options wrong?
upvoted 1 times

Sulzirsha 8 months, 1 week ago
Was on exam feb21
upvoted 1 times

rotimislaw 10 months ago
Why not E. Create Azure AD custom role?
upvoted 1 times

r3verse 11 months ago

"You need to secure the Azure Logic app.", lol, but we aren't securing the logic app, we are providing a mechanism for the logic app to securely access other resources. (a key not a lock). The logic app itself can still be accessed anonymously.

upvoted 2 times

EmnaDa 11 months ago

correct "Azure AD resources must remain if the Azure Logic app is deleted."

upvoted 1 times

sam5678 1 year ago

correct

upvoted 1 times

iamstudying 1 year, 7 months ago

Selected Answer: A

A, buddies. Agree with @dejanban10

upvoted 2 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose A. Create a user-assigned managed identity and assign role-based access controls.

upvoted 3 times

mattvasc 1 year, 8 months ago

Selected Answer: A

In here https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview

We can found:
"Common use cases: Workloads where resources are recycled frequently, but permissions should stay consistent."

upvoted 1 times

AJ309 2 years, 1 month ago

logic app is out of course right?

upvoted 8 times

mariodarken 1 year, 8 months ago

Yes, but this is more a Azure AD question

upvoted 1 times

Question #32 *Topic 4*

HOTSPOT -

You are developing an application that uses a premium block blob storage account. You are optimizing costs by automating Azure Blob Storage access tiers.

You apply the following policy rules to the storage account. You must determine the implications of applying the rules to the data. (Line numbers are included for reference only.)

```
01 {
02   "rules":
03     {
04       "name": "agingDataRule",
05       "enabled": true,
06       "type": "Lifecycle",
07       "definition":
08         "filters": {
09           "blobTypes": [ "blockBlob" ],
10           "prefixMatch": [ "container1/salesorders", "container2/inventory" ]
11         },
12         "actions": {
13           "baseBlob": {
14             "tierToCool": { "daysAfterModificationGreaterThan": 60 },
15             "tierToArchive": { "daysAfterModificationGreaterThan": 120 }
16           }
17         }
18       }
19     },
20     {
21       "enabled": true,
22       "name": "lastAccessedDataRule",
23       "type": "Lifecycle",
24       "definition": {
25         "actions": {
26           "baseBlob": {
27             "enableAutoTierToHotFromCool": true,
28             "tierToCool": {
29               "daysAfterLastAccessTimeGreaterThan": 30
30             }
31           }
32         },
33         "filters": {
34           "blobTypes": [ "blockBlob" ]
35         }
36       }
37     },
38     {
39       "rules": [
40         {
41           "name": "expirationDataRule",
42           "enabled": true,
43           "type": "Lifecycle",
44           "definition": {
45             "filters": {
46               "blobTypes": [ "blockBlob" ]
47             },
48             "actions": {
49               "baseBlob": {
50                 "delete": { "daysAfterModificationGreaterThan": 730 }
51               }
52             }
53           }
54         }
55       ]
56     }
57   ]
58 }
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

	Yes	No
Block blobs prefixed with container1/salesorders or container2/inventory which have not been modified in over 60 days are moved to cool storage. Blobs that have not been modified in 120 days are moved to the archive tier.	○	○
Blobs are moved to cool storage if they have not been accessed for 30 days.	○	○
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.	○	○
All block blobs older than 730 days will be deleted.	○	○

Correct Answer:

Answer Area

	Yes	No
Block blobs prefixed with container1/salesorders or container2/inventory which have not been modified in over 60 days are moved to cool storage. Blobs that have not been modified in 120 days are moved to the archive tier.	○	○
Blobs are moved to cool storage if they have not been accessed for 30 days.	○	○
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.	○	○
All block blobs older than 730 days will be deleted.	○	○

Box 1: Yes -

```
"rules":
  {
    "name": "agingDataRule",
    "enabled": true,
    "type": "Lifecycle",
    "definition":
      "filters": {
        "blobTypes": [ "blockBlob" ],
        "prefixMatch": [ "container1/salesorders", "container2/inventory" ]
      },
      "actions": {
        "baseBlob": {
          "tierToCool": { "daysAfterModificationGreaterThan": 60 },
          "tierToArchive": { "daysAfterModificationGreaterThan": 120 }
```

Box 2: Yes -

```
"enabled": true,
"name": "lastAccessedDataRule",
"type": "Lifecycle",
"definition": {
  "actions": {
    "baseBlob": {
      "enableAutoTierToHotFromCool": true,
      "tierToCool": {
        "daysAfterLastAccessTimeGreaterThan": 30
      }
    }
  },
```

Box 3: Yes -

Box 4: Yes -

```
"rules": [
  {
    "name": "expirationDataRule",
    "enabled": true,
    "type": "Lifecycle",
    "definition": {
      "filters": {
        "blobTypes": [ "blockBlob" ]
      },
      "actions": {
        "baseBlob": {
          "delete": { "daysAfterModificationGreaterThan": 730 }
```

finnishr Highly Voted 2 years, 2 months ago

Correct answer:
1. Yes
2. Yes
3. Yes
4. No

upvoted 72 times

windflower555 2 years, 1 month ago

Agree, 4 should be no. Instead of "All block blobs older than 730 days will be deleted", it should be: "all block blobs 730 days after last modified".
All info can be found at:

https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#move-aging-data-to-a-cooler-tier

upvoted 14 times

Chris2349 4 months, 2 weeks ago

Went with this answer as well. Received this question on 15th of June 2023. Scored 887

upvoted 2 times

sbsd 8 months ago

I think you missed the difference between Blobs and Block blobs.
1. Yes
2. No - configuration is for Block Blobs and not for blob
3. No - configuration is for Block Blobs and not for blob
4. Yes - because of delete property

upvoted 1 times

TOM101 8 months ago

The scenario is about premium block blobs, so blobs can only be block blobs here

upvoted 1 times

warchoon 8 months ago

So 1 must be No too because of "Blobs" for 120 days

upvoted 1 times

Esward 9 months, 1 week ago

Agree with you!

upvoted 1 times

jay158 Highly Voted 2 years, 4 months ago

Answer is correct
1. Yes
2. Yes
3. Yes
4. Yes

upvoted 38 times

jasifu3 1 year, 8 months ago

wrong - 4 is No. If a block has recently been modified then it will not be deleted despite being older than 730 days.

upvoted 14 times

Knightie 1 year, 1 month ago

sorry, you are correct... it's modified>730 not just older than 730,
4. No

upvoted 3 times

Knightie 1 year, 1 month ago

no, even it got back to warm due to "accessed", the "modification" >730 will directly delete it still regardless where the tier it lies.

upvoted 1 times

ZodiaC 2 years, 3 months ago

1000000000000% FAILS NOT GOOD PLZ STOP

upvoted 4 times

vtomy 2 years, 3 months ago

Answer is correct

upvoted 4 times

ucsdmiami2020 1 year, 11 months ago

4 is Yes. Per the Microsoft docs "Some data is expected to expire days or months after creation. You can configure a lifecycle management policy to expire data by deletion based on data age. The following example shows a policy that deletes all block blobs older than 365 days."

"actions": {
"baseBlob": {
"delete": { "daysAfterModificationGreaterThan": 365 }
}
}

https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#move-aging-data-to-a-cooler-tier

upvoted 3 times

mattvasc 1 year, 7 months ago

I have opened an PR, and now, the docs is correct:
"The following example shows a policy that deletes all block blobs that have not been modified in the last 365 days."
PR:
https://github.com/MicrosoftDocs/azure-docs/pull/89203

Docs:
https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#expire-data-based-on-age

upvoted 8 times

gmishra88 1 year ago

Bad that you cannot start a PR on this question from Microsoft :)

upvoted 1 times

pandugadu009 Most Recent 1 month ago

Got it today and went with
No
No
No
Yes
Passed the exam with 749 score.

upvoted 1 times

leviatas 1 month ago

Got this question in exam 2023.09.25. Contoso case Study

upvoted 1 times

SangeethaRamasamy 1 month, 2 weeks ago

Got this in 14/09/2023 exam

upvoted 2 times

Elbanna 1 month, 1 week ago

which answers did you choose?

upvoted 1 times

Dreamer999 4 months, 2 weeks ago

Answer should be : No, No ,No, No
The key is that this is premium storage account and moving between tiers are not supported using lifecycle management policy.
I tested it by creating a premium storage account with the blob type as block blob. In the life cycle management window for adding rules in the portal, I only see option to delete the blob. No option to move to different tiers.

upvoted 6 times

JogSte 2 months ago

I tested too and Premium Block Blob Storage Account only accept Delete lifecycle, not tierToCool neither tierToArchive.
The documentation says 1-NO, 2-NO, 3-NO
"Data stored in a premium block blob storage account cannot be tiered to hot, cool, cold or archive by using Set Blob Tier or using Azure Blob Storage lifecycle management. To move data, you must synchronously copy blobs from the block blob storage account to the hot tier in a different account using the Put Block From URL API or a version of AzCopy that supports this API."

https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview

upvoted 1 times

prakash007 5 months, 3 weeks ago

Correct answer :
1. Yes
2. Yes
3. Yes
4. No.

Reference: -
https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#move-aging-data-to-a-cooler-tier

upvoted 1 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 3 times

rasojol523 7 months, 3 weeks ago

Box 1:
No. Premium block BLOBs cannot use the archive layer.

Box2:
No.
Box2 is "if not 'accessed' for 30 days". The tireToCool rule, on the other hand, is "60-day modification" because it is a tireToCool rule with a tireToCoolGraterThan 60, so the dates and conditions are different.

Box 3:.
Yes. enableAutoTireToHotFromCool w/ dayAfterLastAccessTimeGreaterThan 30.

Box4:
Presumably Yes.
If "older than 730 days" has since been changed, then yes.
If since created, then no.

upvoted 2 times

rasojol523 7 months, 3 weeks ago

Box 2: Yes.
Oops, I missed the second TireToCool. This is daysAfterLastAccessTimeGreaterThan : 30.
NO if MS makes a clear distinction between "Greater than 30" and "just 30"...
I'll go with this.
No,Yes,Yes,Yes

upvoted 2 times

UniqueNickname 8 months, 2 weeks ago

We have to be careful here since the question mentions a "premium block blob storage account" is being used here.

Tiering is NOT supported for this account type as the documentation clearly says:
"Tiering is not yet supported in a premium block blob storage account."
(https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview)

What may be confusing is the fact that the documentation provided above also says the following:
"Lifecycle management policies are supported for block blobs and append blobs in general-purpose v2, premium block blob, and Blob Storage accounts"

From my understanding, lifecycle management can only be used for delete policies, but not for tiering. Hence the answers should all be NO.

upvoted 4 times

kgy01 5 months ago

D, is about a delete, not tiering, so it should be YES.

upvoted 1 times

Michael2023 9 months ago

what am I missing here?
I thought Lifecycle Managment isnt supported for "Premium block blob accounts"

I also created 2 storage accounts in Azure to check whats possible... the will be created with HOT or Cold Access Tier but only in Storage Account with "Standard general-purpose v2" you have the possibility of Lifecycle Management

https://learn.microsoft.com/en-us/azure/storage/blobs/storage-feature-support-in-storage-accounts#premium-block-blob-accounts

upvoted 1 times

Michael2023 9 months ago

thats odd... but okay... I found something here, which tells the following ->
"Lifecycle management policies are supported for block blobs and append blobs in general-purpose v2, premium block blob, and Blob Storage accounts. Lifecycle management doesn't affect system containers such as the $logs or $web containers."

upvoted 1 times

Trimack93 11 months ago

It is stated in the question that we are using premium block blob storage account. For these, tiering is not supported:
https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview#blob-lifecycle-management
https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview#rule-actions

So the correct answer should be:
1. NO
2. NO
3. NO
4. NO

upvoted 10 times

aruni_mishra 10 months, 2 weeks ago

Great catch!

"Setting Tiers using Lifecycle Management Policy is supported for Block Blobs only under GPv2 Storage Accounts.
Deletion using Lifecycle Management Policy is supported for Block/Append Blobs under GPv2 Storage Accounts and Block/Append Blobs under Premium Block Blob Storage Accounts.

https://github.com/MicrosoftDocs/azure-docs/issues/94587
&
https://github.com/MicrosoftDocs/azure-docs/issues/100695

upvoted 2 times

Igris 10 months, 1 week ago

Lifecycle management policies are supported for block blobs and append blobs in general-purpose v2, premium block blob, and Blob Storage accounts. https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview

upvoted 1 times

Michael2023 9 months ago

Lifecycle management policies is not supported on Premium Block Blob Accounts

https://learn.microsoft.com/en-us/azure/storage/blobs/storage-feature-support-in-storage-accounts#premium-block-blob-accounts

upvoted 1 times

Michael2023 9 months ago

but here it is saying the following ->
"Lifecycle management policies are supported for block blobs and append blobs in general-purpose v2, premium block blob, and Blob Storage accounts. Lifecycle management doesn't affect system containers such as the $logs or $web containers."...

So I am confused now... in Azure Storage Premium Block Blob Account, I couldnt find the Lifecycle Management menu hmm

https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview

upvoted 1 times

UniqueNickname 8 months, 2 weeks ago

The documentation mentions this:
"Tiering is not yet supported in a premium block blob storage account."
(https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview)

From my understanding, this does not exclude Lifecycle Management since you should be able to delete blobs via rules.

upvoted 2 times

warchoon 8 months ago

But the deletion rule is wrong

upvoted 1 times

kampatra 1 year ago

YES, YES, YES, NO
4th Option is NO, because it's daysAfterModification, it should have daysAfterCreation to check OLDER than 730 days.

upvoted 2 times

gmishra88 1 year ago

"older than 730" does it mean from "created" 730 day or after modification 730 days. I'm not sure Microsoft guy knows the difference or thought about it or whether we are expected to understand the nuance in that English language usage. Microsoft wins again

upvoted 2 times

kampatra 1 year, 1 month ago

Answer should be : YES, NO, YES, YES
2nd All blob WON'T move to cool tier - only with given prefix

upvoted 1 times

andrvelich 1 year, 6 months ago

4 is yes. You can't use "daysAfterCreationGreaterThan" for baseBlob delete. It is available only for snapshot or version blobs
https://docs.microsoft.com/en-us/answers/questions/637936/error-updating-azure-blob-lifecycle-34could-not-fi.html

upvoted 1 times

Azprep 1 year, 6 months ago

No
No,
No,
Yes

upvoted 1 times

Question #33 *Topic 4*

You are developing a solution that will use a multi-partitioned Azure Cosmos DB database. You plan to use the latest Azure Cosmos DB SDK for development.

The solution must meet the following requirements:

- Send insert and update operations to an Azure Blob storage account.
- Process changes to all partitions immediately.
- Allow parallelization of change processing.

You need to process the Azure Cosmos DB operations.

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

NOTE: Each correct selection is worth one point.

A. Create an Azure App Service API and implement the change feed estimator of the SDK. Scale the API by using multiple Azure App Service instances.

B. Create a background job in an Azure Kubernetes Service and implement the change feed feature of the SDK.

C. Create an Azure Function to use a trigger for Azure Cosmos DB. Configure the trigger to connect to the container.

D. Create an Azure Function that uses a FeedIterator object that processes the change feed by using the pull model on the container. Use a FeedRange object to parallelize the processing of the change feed across multiple functions.

Correct Answer: *AC*

Azure Functions is the simplest option if you are just getting started using the change feed. Due to its simplicity, it is also the recommended option for most change feed use cases. When you create an Azure Functions trigger for Azure Cosmos DB, you select the container to connect, and the Azure Function gets triggered whenever there is a change in the container. Because Azure Functions uses the change feed processor behind the scenes, it automatically parallelizes change processing across your container's partitions.

Note: You can work with change feed using the following options:

- Using change feed with Azure Functions
- Using change feed with change feed processor

Reference:

https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed

Community vote distribution

CD (38%) D (25%) A (19%) Other

Zulhin Highly Voted 2 years, 3 months ago

Answer C & D.
"What are two possible ways to achieve this goal?"

upvoted 40 times

Chiboy 1 year, 8 months ago

C is the correct answer. Each answer has two possible ways that in combination, gives the solution. Moreover, because of the requirement to "Process changes to all partitions immediately", D cannot be the answer or part of it since it "processes the change feed by using the pull model on the container"

upvoted 3 times

Baskman 1 year, 8 months ago

Agree with C & D
"NOTE: Each correct selection is worth one point."

upvoted 3 times

aradice Highly Voted 2 years, 4 months ago

c and d ? https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed#azure-functions
https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model#using-feedrange-for-parallelization

upvoted 16 times

Jurgen1234 2 years, 4 months ago

I agree

upvoted 3 times

Dianahu Most Recent 3 months ago

isn't the change feed estimator only for monitoring? it is the change feed processor who does the work https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-use-change-feed-estimator?tabs=dotnet so it is C&D

upvoted 1 times

adilkhan 7 months ago

C,D chatGPT

upvoted 2 times

Ciupaz 4 weeks, 1 day ago

In my case chatGPT response is A and C.

upvoted 1 times

adilkhan 8 months ago

CD (36%)
D (21%)
A (21%)
Other ...OMG :D

upvoted 4 times

mabdo 9 months ago

similar was on 2/23

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: CD

I agree with Zulhin.

upvoted 1 times

ldenouter 9 months, 2 weeks ago

Selected Answer: CD

I agree with Zulhin.

upvoted 1 times

rotimislaw 10 months ago

Selected Answer: AC

A & C - simplicity is your best friend

upvoted 1 times

AB1453 11 months, 3 weeks ago

Selected Answer: AC

A and C

upvoted 1 times

OPT_001122 12 months ago

Selected Answer: CD

c and d

upvoted 2 times

VirusZer0 1 year ago

Selected Answer: CD

C and D

upvoted 2 times

gmishra88 1 year ago

Another question created to confuse a person who knows change feed processor and did not notice the option had change feed estimator in an exam setting. When the pull model is not the recommended option I should also know that ?

upvoted 1 times

iamstudying 1 year, 7 months ago

Selected Answer: D

C & D, buddies.

C: Push model, Az Functions automatically parallelizes change processing https://docs.microsoft.com/en-us/azure/cosmos-db/sql/read-change-feed#azure-functions

D: Pull model, use FeedRange for parallelization https://docs.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-pull-model#using-feedrange-for-parallelization

upvoted 4 times

SivajiTheBoss 1 year, 7 months ago

CD seems correct answer

upvoted 3 times

Dev666 1 year, 8 months ago

Selected Answer: A

A and C ?

upvoted 1 times

Vantirup 1 year, 8 months ago

Selected Answer: C

A seems to be the answer

upvoted 1 times

Question #34 *Topic 4*

HOTSPOT -

You have an Azure Web app that uses Cosmos DB as a data store. You create a CosmosDB container by running the following PowerShell script:

```
$resourceGroupName = "testResourceGroup"
$accountName = "testCosmosAccount"
$databaseName = "testDatabase"
$containerName = "testContainer"
$partitionKeyPath = "/EmployeeId"
$autoscaleMaxThroughput = 5000

New-AzCosmosDBSqlContainer -
-ResourceGroupName $resourceGroupName
-AccountName $accountName
-DatabaseName $databaseName
-Name $containerName
-PartitionKeyKind Hash
-PartitionKeyPath $partitionKeyPath
-AutoscaleMaxThroughput $autoscaleMaxThroughput
```

You create the following queries that target the container:

SELECT * FROM c WHERE c.EmployeeId > '12345'

SELECT * FROM c WHERE c.UserID = '12345'

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

	Yes	No
The minimum throughput for the container is 400 R/Us.	○	○
The first query statement is an in-partition query.	○	○
The second query statement is a cross-partition query.	○	○

Correct Answer:

Answer Area

	Yes	No
The minimum throughput for the container is 400 R/Us.	○	○
The first query statement is an in-partition query.	○	○
The second query statement is a cross-partition query.	○	○

Box 1: No -

You set the highest, or maximum RU/s Tmax you don't want the system to exceed. The system automatically scales the throughput T such that 0.1* Tmax <= T <=

Tmax.

In this example we have autoscaleMaxThroughput = 5000, so the minimum throughput for the container is 500 R/Us.

Box 2: No -

First query: SELECT * FROM c WHERE c.EmployeeId > '12345'

Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition. In order to be an in-partition query, the query must have an equality filter that includes the partition key:

SELECT * FROM c WHERE c.DeviceId > 'XMS-0001'

Box 3: Yes -

Example of In-partition query:

Consider the below query with an equality filter on DeviceId. If we run this query on a container partitioned on DeviceId, this query will filter to a single physical partition.
SELECT * FROM c WHERE c.DeviceId = 'XMS-0001'
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-query-container

aradice Highly Voted 2 years, 4 months ago

correct:
1° no => line 6
2° no =>
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer#overview-of-provisioned-throughput-types
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-query-container#in-partition-query
"Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition. In order to be an in-partition query, the query must have an equality filter that includes the partition key: SELECT * FROM c WHERE c.DeviceId > 'XMS-0001'"
3° : yes => partition key is EmployeeId. https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-query-container#cross-partition-query

upvoted 40 times

xRiot007 1 year, 7 months ago

The documentation on the second point could be better.
They put the most important part at the end and that is that even if the query uses the partition key it MUST use an equality filter to be considered in-partition.

upvoted 3 times

finnishr 2 years, 2 months ago

100% correct!

upvoted 1 times

HimanshuNankani Highly Voted 2 years ago

1. No : Because as question specifies max RUs = 5000, so minimum RUs = 5000/10 = 500
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer#overview-of-provisioned-throughput-types
2. No : In-partition query needs equality filter that is missing here
3. Yes: Though we have equality filter here, 'UserId' is not a partition key here.
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-query-container

upvoted 25 times

leviatas Most Recent 1 month ago

Got this question in my exam 2023.09.25. Had contoso Case Study

upvoted 1 times

applepie 3 months ago

got this question today, answer NNY- 7/30/2023, score 895/1000

upvoted 2 times

NightshadeRC 3 months, 1 week ago

Had this question in today's exam: 2023-07-26

upvoted 2 times

Hogwarts 6 months, 4 weeks ago

Got this in exam today (5 April 2023)

upvoted 3 times

Saluk_DE 7 months ago

Question was on exam 2023-03-30

upvoted 1 times

proffesormuffin 8 months ago

400 looks correct to me.

"To estimate the minimum throughput required of a container with manual throughput, find the maximum of:

400 RU/s
Current storage in GB * 1 RU/s
Highest RU/s ever provisioned on the container / 100"
5000/100 = 50 not 500
so MAX(400 ,50) = 400 ?

https://learn.microsoft.com/en-us/azure/cosmos-db/concepts-limits

upvoted 2 times

Xardas 3 months, 2 weeks ago

Highest RU/s ever provisioned on the container * 0.1 = 500

upvoted 1 times

Xardas 3 months, 2 weeks ago

This link says formula is max * 0.1
https://learn.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer#overview-of-provisioned-throughput-types

But this link says it is max / 100
https://learn.microsoft.com/en-us/azure/cosmos-db/concepts-limits

Which is right????

upvoted 1 times

mabdo 9 months ago

on 2/23

upvoted 2 times

Esward 9 months, 1 week ago

Given answers are correct.

upvoted 1 times

gmishra88 1 year ago

First option is a question that punishes people who tried to set the RUs and remembered 400 RU as the minimum. If you never tried, you will find some relation with 5000 and 10% is a nice relation.

upvoted 3 times

proffesormuffin 8 months ago

it's 1% not 10% so max 400,50 is 400.
https://learn.microsoft.com/en-us/azure/cosmos-db/concepts-limits

upvoted 1 times

[Removed] 1 year, 1 month ago

Answer
No, No, Yes, got this in the cosmos dp-420 exam

upvoted 1 times

Azprep 1 year, 6 months ago

No,
No,
Yes

upvoted 2 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose as below:
The minimum throughput for the container is 400 R/Us = No
The first query statement is an in-partition query. = No
The second query statement is a cross-partiton query. = Yes

upvoted 2 times

leonidn 1 year, 9 months ago

No
No
Yes
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-query-container#in-partition-query
"Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition. In order to be an in-partition query, the query must have an equality filter that includes the partition key"

upvoted 3 times

xortan 2 years, 1 month ago

"Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition. In order to be an in-partition query, the query must have an equality filter that includes the partition key.
SELECT * FROM c WHERE c.DeviceId > 'XMS-0001'"

This is from https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-query-container#in-partition-query

So based on this, the second cannot be an in-partition query, because it is not an equality filter.

upvoted 2 times

mcbc 2 years, 2 months ago

No->Yes->Yes

PartitionKey EmaployeeID. use common sense

upvoted 1 times

Question #35 *Topic 4*

HOTSPOT -

You are developing a web application that makes calls to the Microsoft Graph API. You register the application in the Azure portal and upload a valid X509 certificate.

You create an appsettings.json file containing the certificate name, client identifier for the application, and the tenant identifier of the Azure Active Directory (Azure

AD). You create a method named ReadCertificate to return the X509 certificate by name.

You need to implement code that acquires a token by using the certificate.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
AuthenticationConfig config = AuthenticationConfig.ReadFromJsonFile("appsettings.json");
X509Certificate2 certificate = ReadCertificate(config.CertificateName);
var app =                                                    .Create(config.ClientId)
```

- ConfidentialClientApplicationBuilder
- GetAccountAsync()
- GetAccountsAsync()
- ConfidentialClientApplication

```
  .WithCertificate(certificate)
  .WithAuthority(new Uri(config.Authority))
  .Build();
string[] scopes = new string[] { $"{config.ApiUrl}.default" };
AuthenticationResult result = await app.AcquireTokenForClient(                ).ExecuteAsync();
```

- scopes
- app
- config

Correct Answer:

Answer Area

```
AuthenticationConfig config = AuthenticationConfig.ReadFromJsonFile("appsettings.json");
X509Certificate2 certificate = ReadCertificate(config.CertificateName);
var app =                                                    .Create(config.ClientId)
```

- ConfidentialClientApplicationBuilder
- GetAccountAsync()
- GetAccountsAsync()
- ConfidentialClientApplication

```
  .WithCertificate(certificate)
  .WithAuthority(new Uri(config.Authority))
  .Build();
string[] scopes = new string[] { $"{config.ApiUrl}.default" };
AuthenticationResult result = await app.AcquireTokenForClient(                ).ExecuteAsync();
```

- scopes
- app
- config

Box 1: ConfidentialClientApplicationBuilder

Here's the code to instantiate the confidential client application with a client secret: app =

ConfidentialClientApplicationBuilder.Create(config.ClientId)

.WithClientSecret(config.ClientSecret)

.WithAuthority(new Uri(config.Authority))

.Build();

Box 2: scopes -

After you've constructed a confidential client application, you can acquire a token for the app by calling AcquireTokenForClient, passing the scope, and optionally forcing a refresh of the token.

Sample code: result = await app.AcquireTokenForClient(scopes)

.ExecuteAsync();

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-app-configuration https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-acquire-token

aradice Highly Voted 2 years, 4 months ago

correct https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-app-configuration?tabs=dotnet#instantiate-the-confidential-client-application-with-a-client-certificate
https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-acquire-token?tabs=dotnet#acquiretokenforclient-api

upvoted 50 times

Esward 9 months, 1 week ago

Given answers are correct! agreed.

upvoted 2 times

MK22 Highly Voted 2 years, 2 months ago

Provided answer is correct.

upvoted 12 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 06-29-2023

upvoted 5 times

tmuras 8 months, 2 weeks ago

This question was on exam on 17/02/23.

upvoted 4 times

bertwin 11 months, 1 week ago

Absolute correct

upvoted 1 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 5 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 9 times

Question #36 *Topic 4*

HOTSPOT -

You develop a containerized application. You plan to deploy the application to a new Azure Container instance by using a third-party continuous integration and continuous delivery (CI/CD) utility.

The deployment must be unattended and include all application assets. The third-party utility must only be able to push and pull images from the registry. The authentication must be managed by Azure Active Directory (Azure AD). The solution must use the principle of least privilege.

You need to ensure that the third-party utility can access the registry.

Which authentication options should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Authentication	Option
Registry authentication method	Service principal Individual identity Repository-scoped access token Managed identity for Azure resources
RBAC role	AcrPull Owner AcrPush Contributor

Correct Answer:

Answer Area

Authentication	Option
Registry authentication method	Service principal Individual identity Repository-scoped access token Managed identity for Azure resources
RBAC role	AcrPull Owner AcrPush Contributor

Box 1: Service principal -

Applications and container orchestrators can perform unattended, or "headless," authentication by using an Azure Active Directory (Azure AD) service principal.

Incorrect Answers:

- Individual AD identity does not support unattended push/pull
- Repository-scoped access token is not integrated with AD identity
- Managed identity for Azure resources is used to authenticate to an Azure container registry from another Azure resource.

Box 2: AcrPush -

AcrPush provides pull/push permissions only and meets the principle of least privilege.

Incorrect Answers:

AcrPull only allows pull permissions it does not allow push permissions.

☞ Owner and Contributor allow pull/push permissions but does not meet the principle of least privilege.
Reference:
https://docs.microsoft.com/en-us/azure/container-registry/container-registry-authentication?tabs=azure-cli https://docs.microsoft.com/en-us/azure/container-registry/container-registry-roles?tabs=azure-cli

dtctx Highly Voted 1 year, 6 months ago

Got this 4/29. I passed test, but scored poorly in third party integration questions.

I went with Managed Identity and Contributor, and I think both are WRONG. I had not heard of ACR Push/Pull, so went with Contributor, but upon learning that ACR Push satisfies least privilege, that is probably the best answer.

Table of privileges:
https://docs.microsoft.com/en-us/azure/container-registry/container-registry-roles?tabs=azure-cli

I suggest ACR Push is the correct permission level.

upvoted 11 times

gmishra88 Highly Voted 1 year ago

Correct. It is only service principal.
acr-token: cannot be because the AD authentication requirement. The roles are RBAC but the identity is not AD managed
Managed identity: Not an option because it is third party

AcrPush includes AcrPull. I did not know that nuance and the choice made by Microsoft to do this

upvoted 8 times

macobuzi 2 months, 1 week ago

I don't understand, Why not Managed Identity? Azure Container Apps also support Managed Identity.

upvoted 1 times

sarmaria Most Recent 7 months, 2 weeks ago

Got this on 16/03/23. Chosen Service Principal and AcrPush. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 6 times

mabdo 9 months ago

on exam 02/23

upvoted 2 times

Esward 9 months, 1 week ago

Service Principal and AcrPush are correct answers!

https://learn.microsoft.com/en-us/azure/container-registry/container-registry-authentication?tabs=azure-cli
https://learn.microsoft.com/en-us/azure/container-registry/container-registry-roles?tabs=azure-cli

upvoted 3 times

OPT_001122 10 months, 3 weeks ago

service principal
AcrPush

upvoted 2 times

hubekpeter 11 months, 1 week ago

Individual identity - no, you don't want to use somebody's username, when he left, service will stop working (saw this tens of times). Managed identity - yes that could work, but you're not sure if CI/CD is running on Azure resource. Repository scoped access token - the next question would be regarding Scope Map, but they are apparently asking about RBAC role. So the correct answer is Service Principal with AcrPush role, which will meet the least priviledge requirement.

upvoted 4 times

coffecold 1 year ago

The question is about access to 'a' registry.
Nowhere in the question is stated that Azure Container Registry (ACR) is used. It talks about "the registry", that is confusing.
Provided Azure Container Registry (ACR) is used, then the answer is correct.
If an other kind of registry is used to store the images the answer can be everything (depending on the registry used).

upvoted 2 times

finnishr 1 year, 1 month ago

The answer is correct.

upvoted 2 times

vvlads 1 year, 5 months ago

why not managed identity? Use only from select Azure services that support managed identities for Azure resources
https://docs.microsoft.com/en-us/azure/container-registry/container-registry-authentication?tabs=azure-cli resources: ...Azure Container Registry...
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/managed-identities-status#azure-services-that-support-managed-identities-for-azure-resources

upvoted 2 times

coffecold 1 year ago

The CI/CD is third party.

upvoted 3 times

sghaha 1 year, 6 months ago

https://docs.microsoft.com/ko-kr/azure/container-registry/container-registry-authentication?tabs=azure-cli

upvoted 2 times

Question #37 *Topic 4*

You deploy an Azure App Service web app. You create an app registration for the app in Azure Active Directory (Azure AD) and Twitter.
The app must authenticate users and must use SSL for all communications. The app must use Twitter as the identity provider.
You need to validate the Azure AD request in the app code.
What should you validate?

A. ID token header

B. ID token signature

C. HTTP response code

D. Tenant ID

Correct Answer: *A*
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-auth-aad-app?tabs=dotnet

Community vote distribution
B (92%) 8%

RaghavMGupta Highly Voted 6 months, 3 weeks ago

Selected Answer: B
To validate the Azure AD request in the app code when using Twitter as the identity provider, you should validate the ID token signature (option B).

The ID token is a JSON Web Token (JWT) that contains claims about the user. It is signed by Azure AD using a private key, and the signature can be verified using the corresponding public key. Validating the ID token signature ensures that the token was issued by a trusted source and that it has not been tampered with in transit.

Option A, validating the ID token header, is not sufficient for validating the entire ID token. The header only contains metadata about the token, such as the algorithm used for signing.

Option C, validating the HTTP response code, is unrelated to validating the ID token.

Option D, validating the tenant ID, is important for ensuring that the app is only accepting tokens from a trusted Azure AD tenant, but it does not ensure the integrity of the token itself.

upvoted 10 times

elequiel Highly Voted 1 year ago

Got it in exam 20/10/2022

upvoted 6 times

OPT_001122 12 months ago

Thanks for mentioning date https://www.examtopics.com/user/elequiel/

upvoted 1 times

Ciupaz Most Recent 1 month ago

Selected Answer: B
Validating the ID token signature ensures that the token is valid and hasn't been altered.

upvoted 1 times

kabbas 4 months, 2 weeks ago

Selected Answer: B
A JWT contains three segments, a header, a body, and a signature. The signature segment can be used to validate the authenticity of the token so that it can be trusted by your application

upvoted 2 times

[Removed] 7 months, 3 weeks ago

Selected Answer: B
B. ID token signature

upvoted 3 times

sbsd 8 months ago

Current answer is correct (A). You can edit the token header in order to include the source. You cannot change the signature. check it out here: jwt.io

upvoted 1 times

Esward 9 months, 1 week ago

B is the correct answer
https://docs.microsoft.com/en-us/azure/active-directory-b2c/tokens-overview#validation
upvoted 3 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. ID token signature

When validating an Azure AD request in the app code, it is important to validate the ID token signature to ensure the authenticity of the token. The ID token contains information about the authenticated user, including the user's identity and any claims or permissions associated with the user. By validating the signature, you can ensure that the token has not been tampered with and that it was indeed issued by Azure AD.
upvoted 1 times

ExamTopicsAnkur 11 months ago

Selected Answer: B

https://docs.microsoft.com/en-us/azure/active-directory-b2c/tokens-overview#validation

When your application or API receives an ID token, it must validate the signature to prove that the token is authentic.
upvoted 3 times

hubekpeter 11 months, 1 week ago

Selected Answer: B

Id token, in another words, you are validating JWT token.
upvoted 2 times

bertwin 11 months, 1 week ago

Selected Answer: B

Signature is for validation
upvoted 3 times

OPT_001122 12 months ago

Selected Answer: B

B seems the correct answer
upvoted 2 times

alany2000 1 year, 1 month ago

Selected Answer: B

B
https://docs.microsoft.com/en-us/azure/active-directory-b2c/tokens-overview#validation
upvoted 4 times

Praks13 1 year, 1 month ago

I would say it's B.
https://docs.microsoft.com/en-us/azure/active-directory-b2c/tokens-overview#validation
upvoted 1 times

finnishr 1 year, 1 month ago

Selected Answer: B

The answer is definitely b.
upvoted 3 times

qwerty112233 1 year, 1 month ago

why? any link?
upvoted 2 times

willchenxa 1 year, 1 month ago

Selected Answer: A

Correct
https://docs.microsoft.com/en-us/azure/active-directory-b2c/identity-provider-twitter?pivots=b2c-user-flow
upvoted 3 times

sghaha 1 year, 6 months ago

my choice is A
upvoted 2 times

Question #38 *Topic 4*

A development team is creating a new REST API. The API will store data in Azure Blob storage. You plan to deploy the API to Azure App Service. Developers must access the Azure Blob storage account to develop the API for the next two months. The Azure Blob storage account must not be accessible by the developers after the two-month time period.
You need to grant developers access to the Azure Blob storage account.
What should you do?

A. Generate a shared access signature (SAS) for the Azure Blob storage account and provide the SAS to all developers.

B. Create and apply a new lifecycle management policy to include a last accessed date value. Apply the policy to the Azure Blob storage account.

C. Provide all developers with the access key for the Azure Blob storage account. Update the API to include the Coordinated Universal Time (UTC) timestamp for the request header.

D. Grant all developers access to the Azure Blob storage account by assigning role-based access control (RBAC) roles.

Correct Answer: *A*
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

Community vote distribution
A (100%)

alexein74 Highly Voted 9 months, 1 week ago

Selected Answer: A
A. Generate a shared access signature (SAS) for the Azure Blob storage account and provide the SAS to all developers.

A shared access signature (SAS) is a secure token that can be used to grant temporary and revocable access to a blob container or individual blobs. You can specify an expiration time for the SAS, so it will automatically expire after the two-month time period, making the blob storage account no longer accessible to the developers.
This approach allows you to grant the developers the necessary access to the Azure Blob storage account while still maintaining control over the access, and it also allows you to revoke access easily after the two-month time period.
upvoted 8 times

richardc7 Most Recent 1 month ago
got in exam 28/09/2029
upvoted 1 times

BaoNguyen2411 3 months ago
got this question on 06-29-2023
upvoted 3 times

OPT_001122 12 months ago
Selected Answer: A
A seems to be the correct answer
upvoted 4 times

taki_ananasek 1 year ago
Seems it's B

A JWT contains three segments, a header, a body, and a signature. The signature segment can be used to validate the authenticity of the token so that it can be trusted by your application.

https://learn.microsoft.com/en-us/azure/active-directory-b2c/tokens-overview#validation
upvoted 1 times

ThaDude 10 months, 2 weeks ago
Comment applies to different question
upvoted 4 times

gmishra88 1 year ago
In practical situations, especially it is about your own developers (not outsiders) and expected to have AD roles, it is good to give RBAC roles and remove it after 2 months. But, I guess Microsoft wants to answer SAS. So, please choose SAS as the answer
upvoted 1 times

coffecold 1 year ago
SAS automatically expires. You don't have to do changes in AD and change it back after two months.

upvoted 1 times

thedevguy 1 year ago

Correct
https://learn.microsoft.com/en-us/azure/cognitive-services/translator/document-translation/create-sas-tokens?tabs=Containers

upvoted 1 times

ArturKon 1 year, 1 month ago

SAS has start and expiry date/time defined.

upvoted 2 times

vvlads 1 year, 6 months ago

https://azure.microsoft.com/de-de/blog/azure-blob-storage-lifecycle-management-now-generally-available/ lifecycle is for blobs and not for the access

upvoted 1 times

sghaha 1 year, 6 months ago

why not B?

upvoted 1 times

finnishr 1 year, 1 month ago

B has nothing to do with granting access to blob storage account.

upvoted 1 times

Question #39 *Topic 4*

DRAG DROP -

You develop a web application.

You need to register the application with an active Azure Active Directory (Azure AD) tenant.

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

Select and Place:

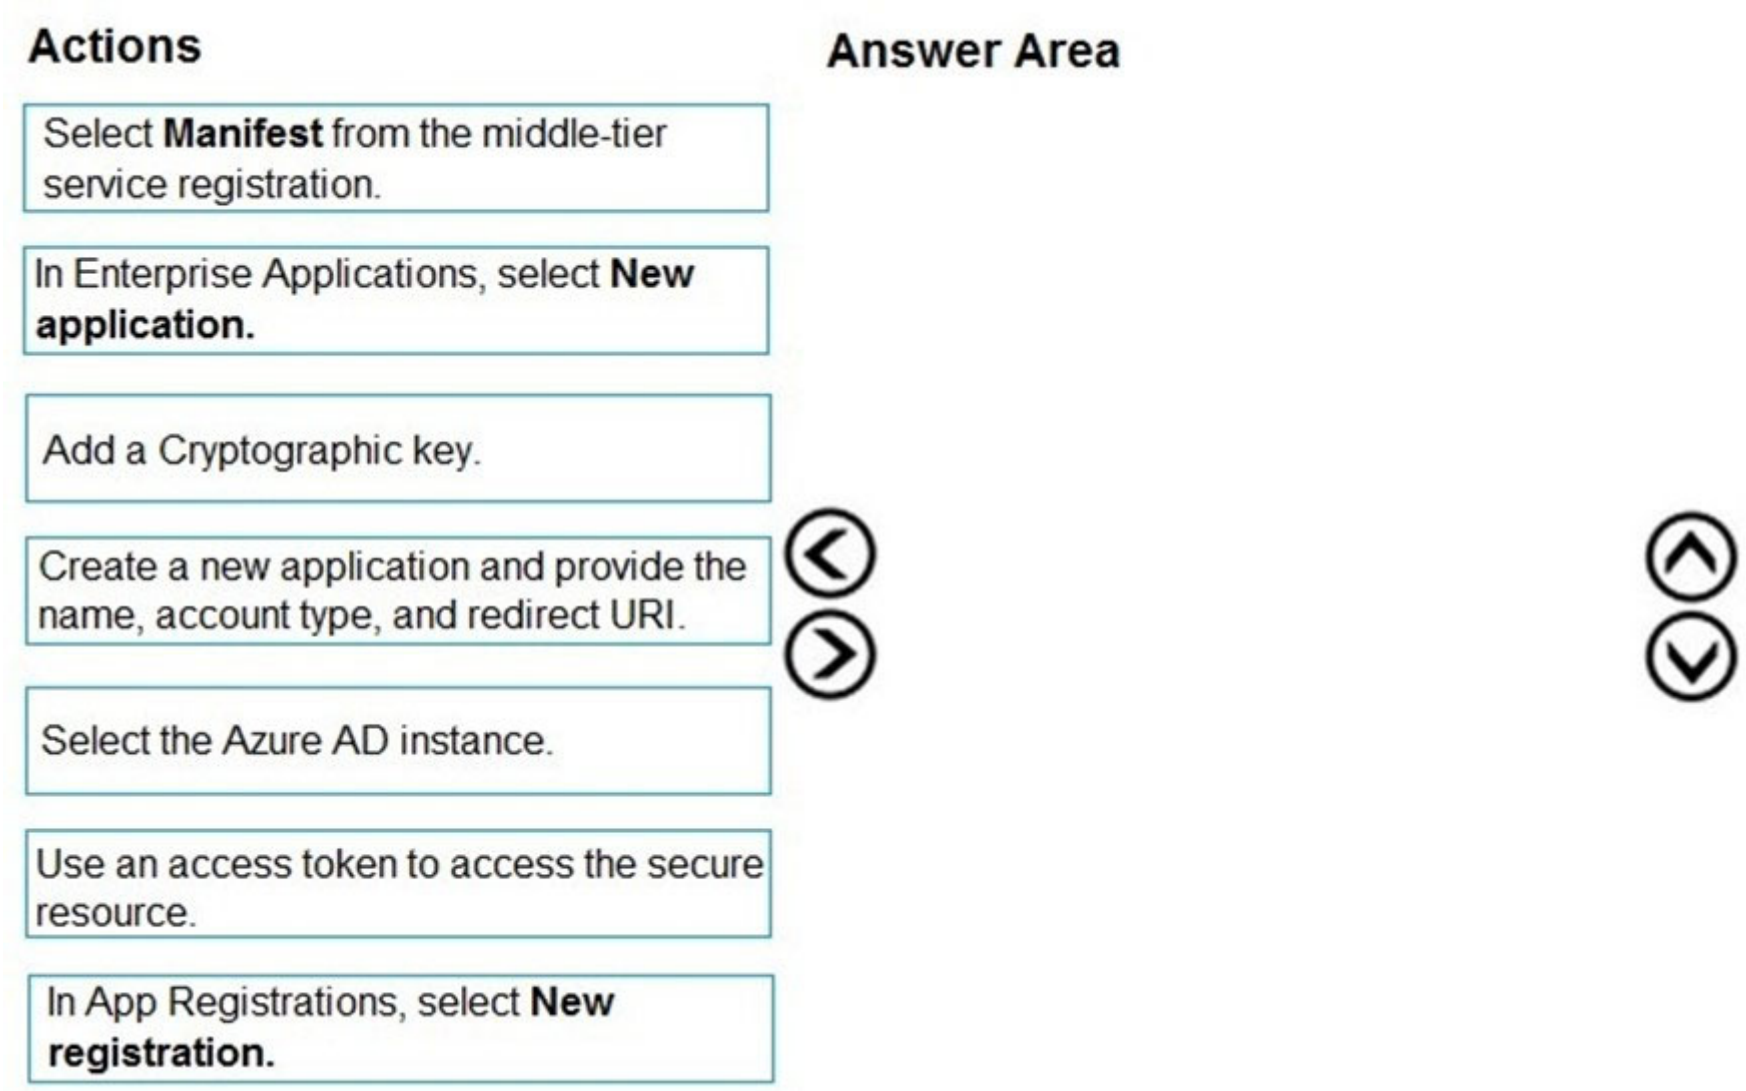

Correct Answer:

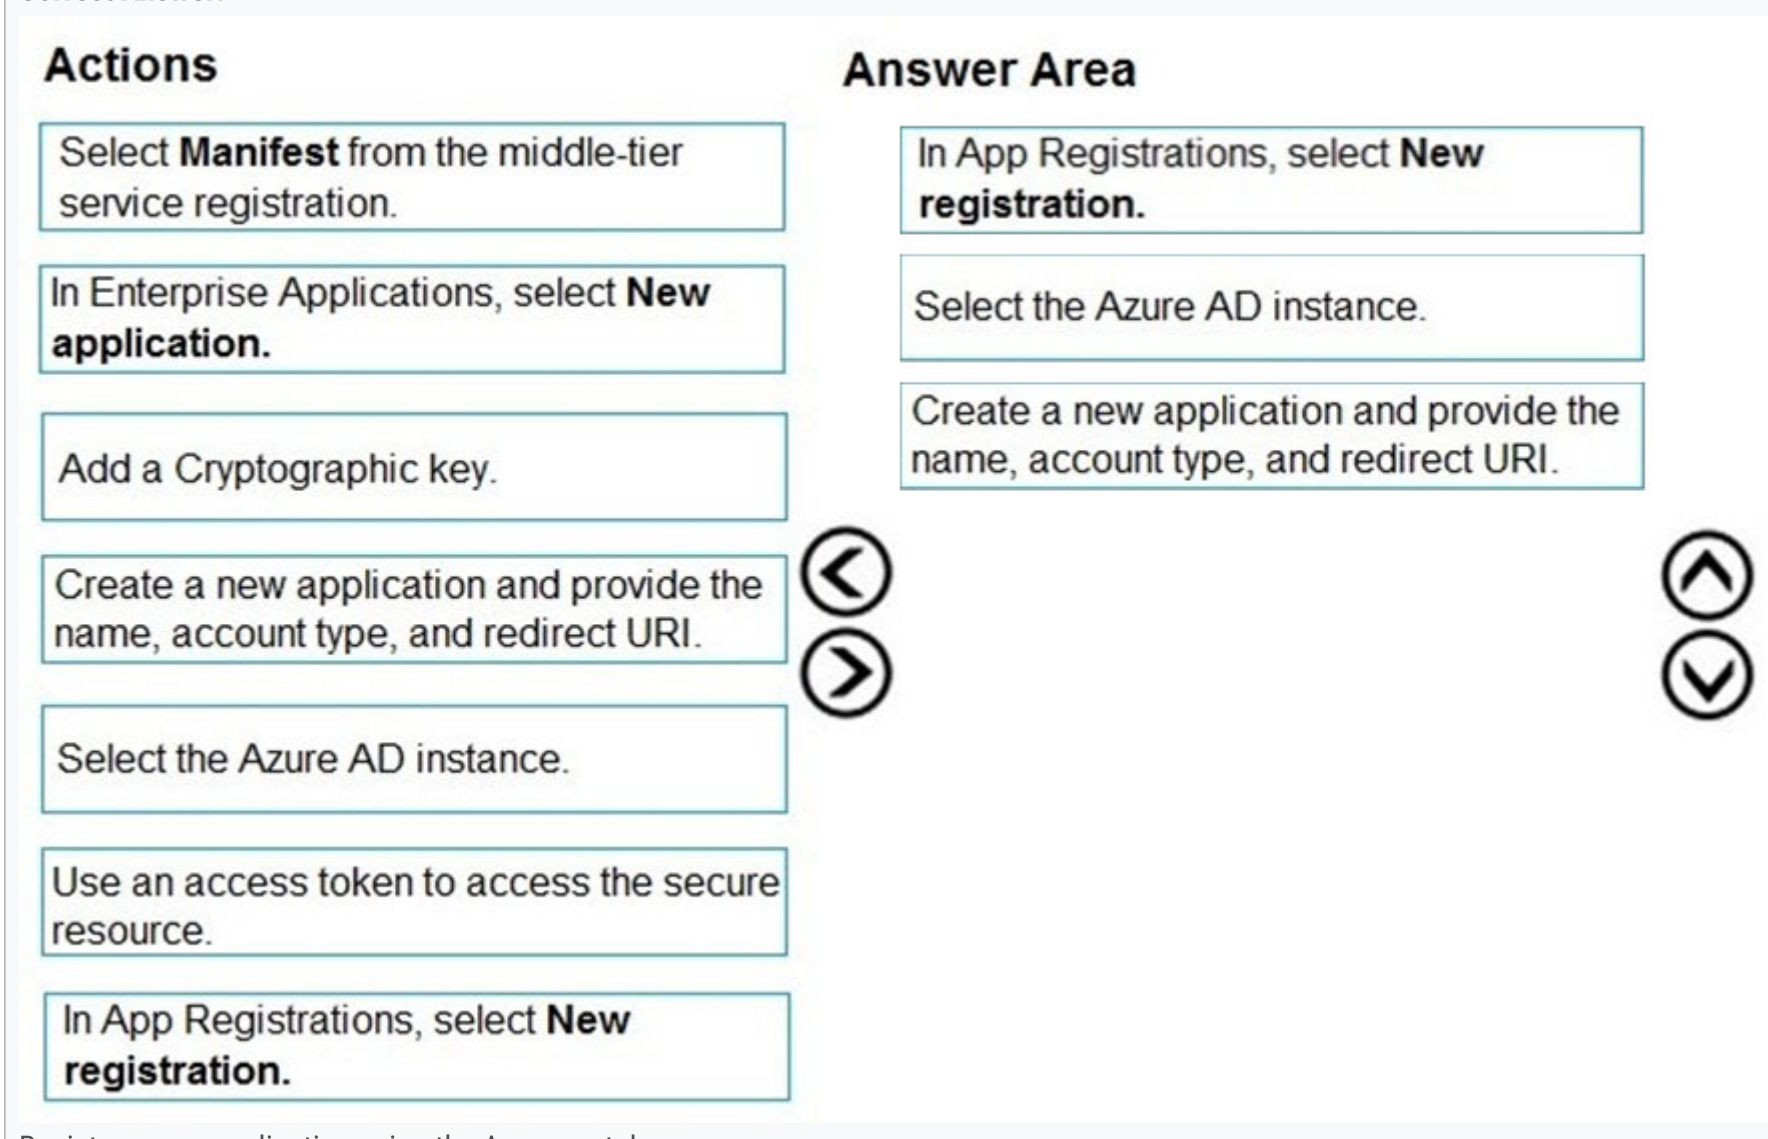

Register a new application using the Azure portal

1. Sign in to the Azure portal using either a work or school account or a personal Microsoft account.

2. If your account gives you access to more than one tenant, select your account in the upper right corner. Set your portal session to the Azure AD tenant that you want.

3. Search for and select Azure Active Directory. Under Manage, select App registrations.
4. Select New registration. (Step 1)
5. In Register an application, enter a meaningful application name to display to users.
6. Specify who can use the application. Select the Azure AD instance. (Step 2)
7. Under Redirect URI (optional), select the type of app you're building: Web or Public client (mobile & desktop). Then enter the redirect URI, or reply URL, for your application. (Step 3)
8. When finished, select Register.

Komat Highly Voted 2 years, 11 months ago

1 Azure AD instance
2 In App Registration, select new registration
3 Create a new application and provide the name
Disscussion: https://www.examtopics.com/discussions/microsoft/view/22224-exam-az-204-topic-3-question-3-discussion/

upvoted 221 times

coffecold 1 year ago

Although the right step, "Create a new application" is wrong formulated here. Either leave it out, or change it in "Register the application, by providing the name, account type, and redirect URI"

upvoted 2 times

tassisto 2 months, 1 week ago

I totally agree with @coffecold

upvoted 1 times

AnonymousJhb 2 years, 10 months ago

Correct Steps
https://docs.microsoft.com/en-us/azure/active-directory/develop/quickstart-register-app#:~:text=If%20you%20have%20access%20to,a%20Name%20for%20your%20application.

upvoted 12 times

Cornholioz 2 years, 10 months ago

It's a chicken and egg scenario for steps 1 & 2. You can go to New App Registration and Select the AAD instance for which you want to register the App... OR you can go through the AAD instance and choose the App to register under that instance.
CHICKEN and EGG!

upvoted 5 times

monniq 2 years, 8 months ago

You can't select the AAD instance on the app registration form. You can only set the supported account type which is included in 'Create new application and provide..." step

upvoted 8 times

SxyPenguin 2 years, 11 months ago

Is this the correct answer?

upvoted 2 times

matejka 2 years, 11 months ago

I would say Komat is right, I would select the same steps.

upvoted 5 times

iiiihhhh 2 years, 10 months ago

I think that the firsts two and half steps of explanation can be summarized as: "Select Azure AD instance". So the explanation is correct but the picture is not.

upvoted 2 times

Cornholioz 2 years, 11 months ago

It is inconclusive in the discussion link too. Looks like 1&2 can be done either ways. Checked in the portal too. Question is unclear.

upvoted 1 times

profesorklaus Highly Voted 2 years, 9 months ago

I agree with you that it should be in following order:
1. Select AD instance - select AD tenant in portal
2. In App Registration, select new registration - use switched tenant
3. Create a new application and provide the name

upvoted 14 times

JH81 Most Recent 4 months ago

Got this on 6/28/2023 and passed with 850. Went with answer.

upvoted 2 times

KrishTeam 2 months, 4 weeks ago

bro did you read only examtopics questions or any other materials

upvoted 1 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

Had cases:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 2 times

OPT_001122 10 months, 3 weeks ago

Komat is right

1 Azure AD instance
2 In App Registration, select new registration
3 Create a new application and provide the name

upvoted 1 times

hubekpeter 11 months, 1 week ago

In app registrationg select new reg
Create a new app and provide the name, account type, and redirect URI
Select Manifest from the middle-tier service registration

upvoted 1 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 1 times

Eltooth 1 year, 4 months ago

Select AD instance
App Registration, create new registration
Create new application, name, account type & URI

upvoted 2 times

Basant11 10 months, 1 week ago

so was your answer correct?

upvoted 2 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose as below:
1. Select the Azure AD Instance
2. In App Registration, select New registration.
3. Create a new application and provide the name, account type, and redirect URL.

upvoted 7 times

Mev4953 1 year, 9 months ago

I think provided answer is very clear

upvoted 1 times

john4p 1 year, 10 months ago

What? You register the non-existing app beforehand and then create it?

upvoted 1 times

coffecold 1 year ago

yes, registering does not check the application exists

upvoted 1 times

Lucario95 1 year, 11 months ago

First answer should be swapped with the second one:
1. Select Azure AD Instance
2. In App Registration, Selecte Create New
3. Create a new application

upvoted 1 times

mimi21212152 1 year, 11 months ago

page 12 is so perfect with the right answers throughout :)

upvoted 2 times

AJ309 2 years, 1 month ago

given answer is correct, since it has already written in question as Active Azure instance

upvoted 2 times

GigaCaster 2 years, 3 months ago

To me, the given answer is correct as it says select Azure AD instance and not create.

upvoted 3 times

DouDouQin 2 years, 3 months ago

Answer is correct!

upvoted 2 times

Kalaisuran 2 years, 3 months ago

My Answer :
1 Azure AD instance
2 In App Registration, select new registration
3 Create a new application and provide the name
Reference Link : https://docs.microsoft.com/en-us/powerapps/developer/data-platform/walkthrough-register-app-azure-active-directory

upvoted 5 times

Question #40 *Topic 4*

You have a new Azure subscription. You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (Azure
AD) for authentication.
You need to implement multifactor authentication for the website.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Configure the website to use Azure AD B2C.

B. In Azure AD, create a new conditional access policy.

C. Upgrade to Azure AD Premium.

D. In Azure AD, enable application proxy.

E. In Azure AD conditional access, enable the baseline policy.

Correct Answer: *BC*
B: 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.
C: 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
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-getstarted

Community vote distribution
BC (100%)

Codenoob Highly Voted 3 years, 1 month ago
B and C are correct.
https://docs.microsoft.com/en-us/azure/active-directory/authentication/tutorial-enable-azure-mfa
upvoted 70 times

sebainones 3 years ago
Yes, based on the link you provided:
"Prerequisites
A working Azure AD tenant with at least an Azure AD Premium P1 or trial license enabled."
upvoted 4 times

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: B and C

Azure AD Multi-Factor Authentication and Conditional Access policies give the flexibility to enable MFA for users during specific sign-in events. The recommended way to enable and use Azure AD Multi-Factor Authentication is with Conditional Access policies.

You need a working Azure AD tenant with at least an Azure AD Premium P1.

Reference:
https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-getstarted
https://docs.microsoft.com/en-us/azure/active-directory/authentication/tutorial-enable-azure-mfa
https://azure.microsoft.com/en-us/pricing/details/active-directory
upvoted 11 times

UniqueNickname Most Recent 8 months, 2 weeks ago
Selected Answer: BC
B and C are correct.
upvoted 1 times

18Marks 10 months, 3 weeks ago
Selected Answer: BC
See in exam

upvoted 1 times

OPT_001122 12 months ago

Selected Answer: BC

B and C are correct.

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: BC

B & C are the correct answers.

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

Selected Answer: BC

BC is the perfect Match

upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 3 times

sawipef270 1 year, 10 months ago

Selected Answer: BC

correct

upvoted 2 times

AhmedHamdo 1 year, 11 months ago

Selected Answer: BC

Answers are correct

upvoted 4 times

ning 2 years, 2 months ago

You can definitely create conditional access policies, in premium pricing tier. For E, the problem is "the baseline" there are several baseline policies that are offered by MS, which one is the one??? Also, I do not know any baseline policy can target a specific web app

upvoted 2 times

kondapaturi 2 years, 4 months ago

Create a new conditional access policy in Azure AD
Upgrade your licensing for Azure AD to Premium license
Hence answer is BC

upvoted 2 times

prabhjot 2 years, 5 months ago

the given ans is 100% correct

upvoted 1 times

UnknowMan 2 years, 5 months ago

The given answer is correct.

upvoted 1 times

glam 2 years, 5 months ago

correct

upvoted 2 times

daporh 2 years, 8 months ago

You can have MFA on the Free tier by enabling security defaults however to use conditional access, security defaults must be disabled
https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/concept-fundamentals-security-defaults#disabling-security-defaults
However Conditional Access is the way to setup MFA on the Premium Tier and since the question does not mention Security Defaults, then the given answer is correct.

upvoted 3 times

Albertoski 2 years, 9 months ago

Given answer is correct. Clearly in the docs says P1 is needed:
https://docs.microsoft.com/en-us/azure/active-directory/authentication/tutorial-enable-azure-mfa

upvoted 2 times

Question #41 *Topic 4*

DRAG DROP -

An organization plans to deploy Azure storage services.

You need to configure shared access signature (SAS) for granting access to Azure Storage.

Which SAS types should you use? To answer, drag the appropriate SAS types to the correct requirements. Each SAS 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:

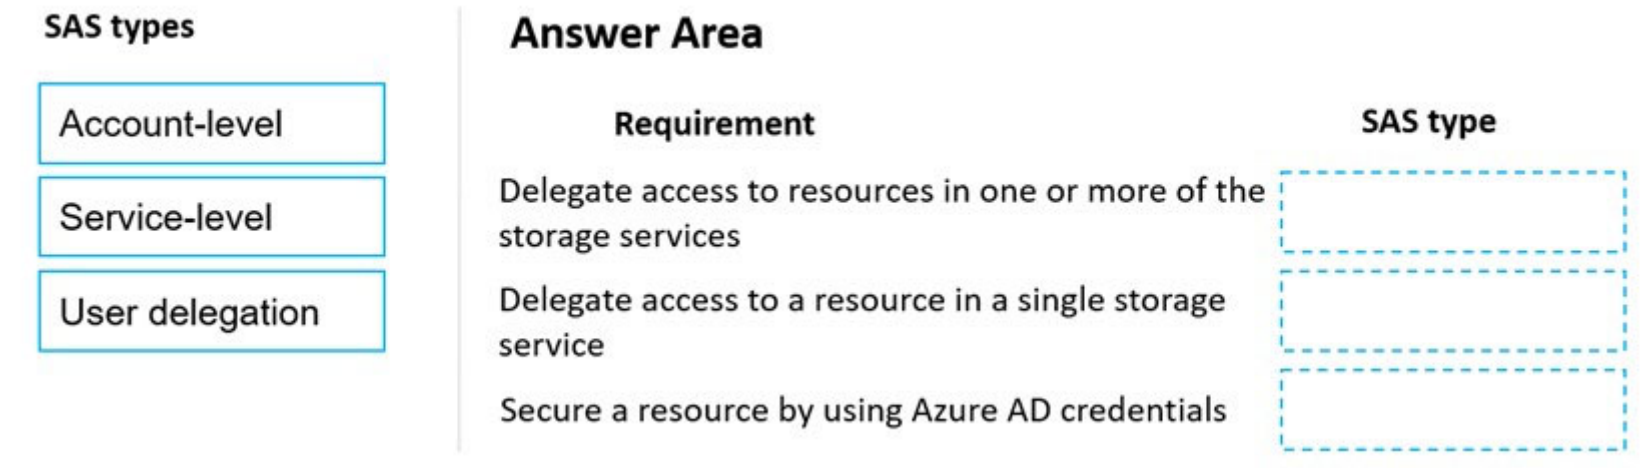

Correct Answer:

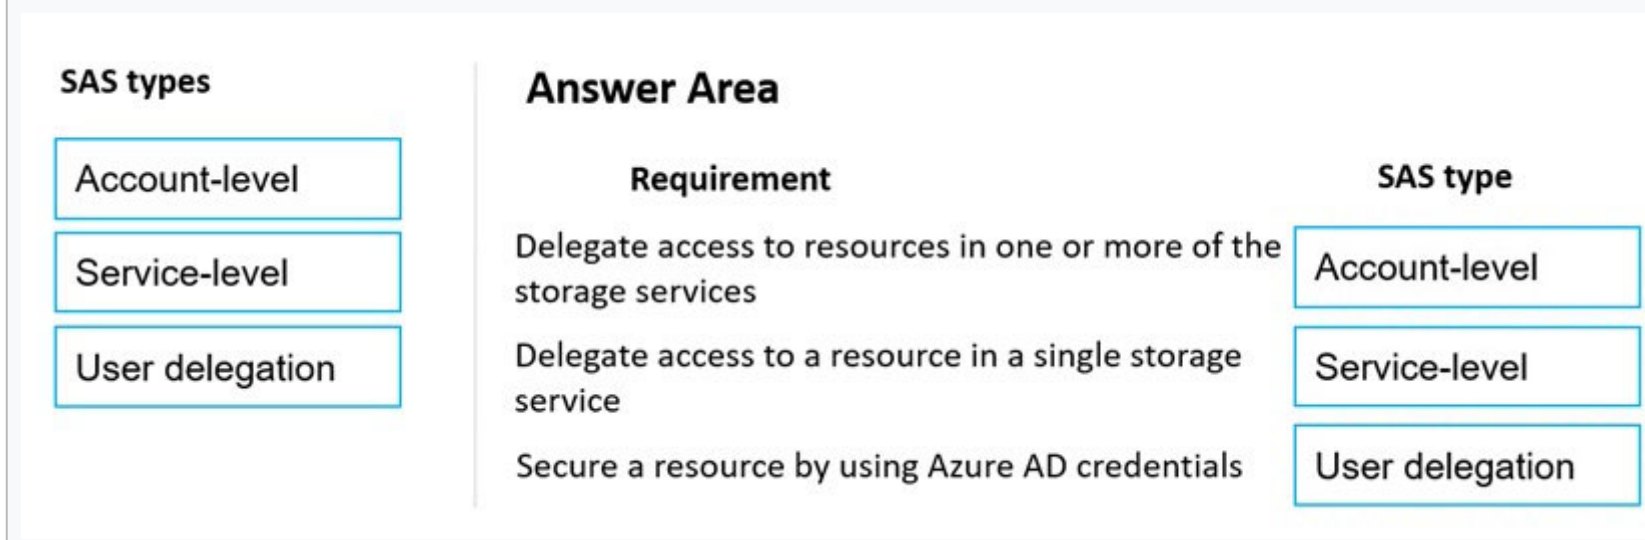

Reference:

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

SoftSol Highly Voted 1 year, 6 months ago

Correct

upvoted 17 times

karthikwarrior Highly Voted 8 months, 1 week ago

If you have reached this far, question no 160+ then you have taken paid subscription (Contributor role)

upvoted 10 times

p2006 1 month, 1 week ago

I can see it without contributor, for now :)

upvoted 2 times

macobuzi 2 months, 1 week ago

LOL, didn't realize that since I already paid since the first questions

upvoted 1 times

naivecoder786 8 months ago

Yeah indeed !

upvoted 1 times

leviatas Most Recent 1 month ago

Got his quesiton in exam 2023.09.25. Case Study Contoso

upvoted 1 times

kvtborad 2 months, 3 weeks ago

I got this question on 6th August 2023. chose highly voted. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.

upvoted 6 times

applepie 3 months ago

got this question today, go with the provided answer - 7/30/2023, score 895/1000

upvoted 3 times

NightshadeRC 3 months, 1 week ago

Had this question in today's exam: 2023-07-26

upvoted 1 times

NombreFalso 8 months, 2 weeks ago

No doubt that's correct

upvoted 1 times

Bear_Polar 11 months, 1 week ago

correct

upvoted 1 times

coffecold 1 year ago

User delegation is supported for Blob Storage only, and you can use it to grant access to containers and blobs. The question does not mention the kind of storage.

upvoted 2 times

damianadalid 5 months, 2 weeks ago

After having read User delegation SAS in the reference link, I think the mentioning of "using Azure AD credentials" is enough to select User delegation despite of not knowing the kind of storage.

upvoted 1 times

Question #42 *Topic 4*

HOTSPOT -

You are developing an ASP.NET Core app that includes feature flags which are managed by Azure App Configuration. You create an Azure App Configuration store named AppFeatureflagStore as shown in the exhibit:

Key	Label	State	Description	Last modified
Export	Export	Off On	Ability to export data.	6/11/2020, 9:13:26 ...

You must be able to use the feature in the app by using the following markup:

```
<feature name="Export">
  <li class="nav-item">
    <a class="nav-link text-dark" asp-area="" asp-controller="Home" asp-action="Export">Export Data</a>
  </li>
</feature>
```

You need to update the app to use the feature flag.

Which values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Code section	Value
Controller attribute	FeatureGate Route ServiceFilter TypeFilter
Startup method	AddAzureAppConfiguration AddControllersWithViews AddUserSecrets
AppConfig endpoint setting	https://appfeatureflagstore.azconfig.io https://appfeatureflagstore.vault.azure.net https://export.azconfig.io https://export.vault.azure.net

Correct Answer:

Answer Area

Code section	Value
Controller attribute	FeatureGate / Route / ServiceFilter / TypeFilter
Startup method	AddAzureAppConfiguration / AddControllersWithViews / AddUserSecrets
AppConfig endpoint setting	https://appfeatureflagstore.azconfig.io / https://appfeatureflagstore.vault.azure.net / https://export.azconfig.io / https://export.vault.azure.net

Box 1: FeatureGate -
You can use the FeatureGate attribute to control whether a whole controller class or a specific action is enabled.

Box 2: AddAzureAppConfiguration -
The extension method AddAzureAppConfiguration is used to add the Azure App Configuration Provider.
Box 3: https://appfeatureflagstore.azconfig.io
You need to request the access token with resource=https://<yourstorename>.azconfig.io
Reference:
https://docs.microsoft.com/en-us/azure/azure-app-configuration/use-feature-flags-dotnet-core
https://csharp.christiannagel.com/2020/05/19/azureappconfiguration/ https://stackoverflow.com/questions/61899063/how-to-use-azure-app-configuration-rest-api

gmishra88 Highly Voted 1 year ago
And Microsoft claims this is not .Net specific.
upvoted 10 times

elequiel Highly Voted 1 year ago
Got it in exam 20/10/2022
upvoted 9 times

p2006 Most Recent 1 month ago
Got on 9/25/2023
featuregare
addazureappconfiguration
appfeatureflagstore.azconfig.io
upvoted 1 times

mihailos 2 months ago
Got it in exam 28/08/23. Went with proposed answer. Scored 912
upvoted 4 times

JogSte 2 months ago
Can you say approximately the number of questions you had in the certification that we have here on the site?
upvoted 3 times

rockstar89 3 months ago

Got it Exam in 31/07/2027.
score 875

upvoted 1 times

Patrick199826 1 week, 2 days ago

2027 LOL

upvoted 1 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 6 times

nekkilodeon 6 months, 1 week ago

Feature gate is correct
https://learn.microsoft.com/en-us/azure/azure-app-configuration/quickstart-feature-flag-aspnet-core?tabs=core6x#use-a-feature-flag

upvoted 1 times

DP_Bhatt 9 months ago

Corrent

upvoted 1 times

sghaha 1 year, 6 months ago

in Korean
https://docs.microsoft.com/ko-kr/azure/azure-app-configuration/use-feature-flags-dotnet-core?tabs=core5x

upvoted 2 times

Question #43 *Topic 4*

HOTSPOT -

You have a single page application (SPA) web application that manages information based on data returned by Microsoft Graph from another company's Azure

Active Directory (Azure AD) instance.

Users must be able to authenticate and access Microsoft Graph by using their own company's Azure AD instance.

You need to configure the application manifest for the app registration.

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

NOTE: Each correct selection is worth one point.

Hot Area:

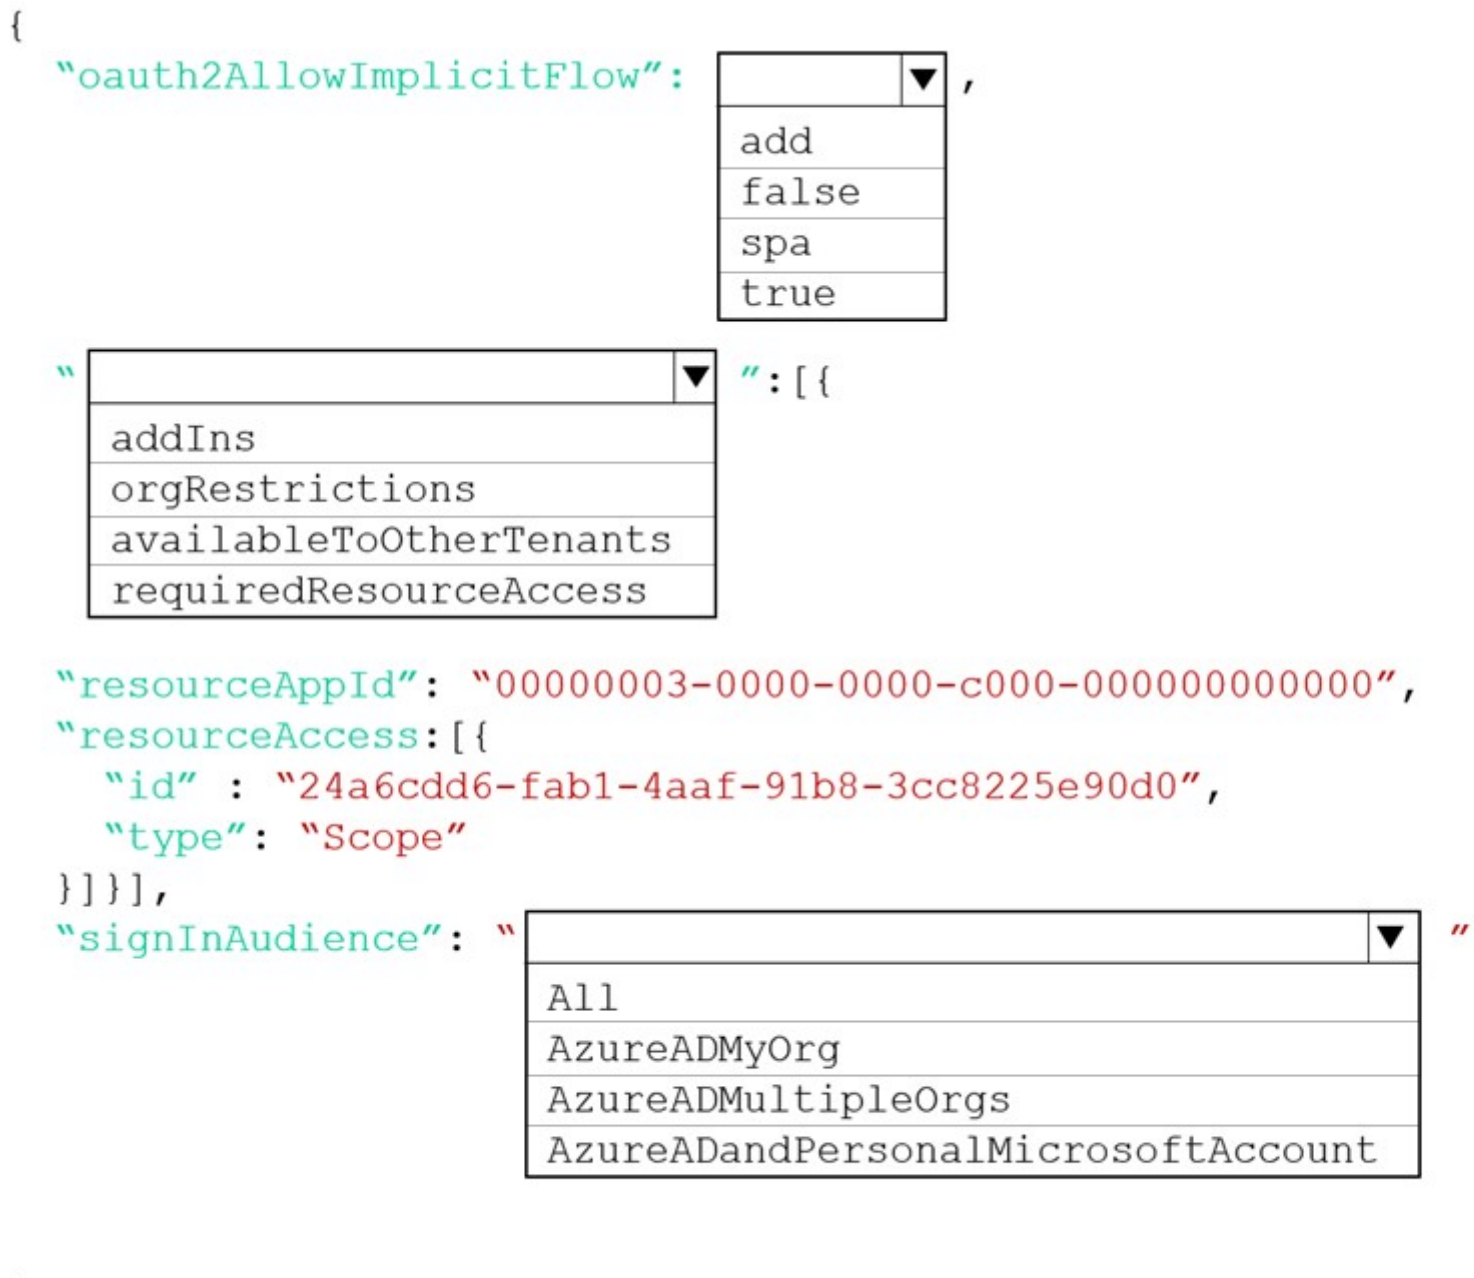

Correct Answer:

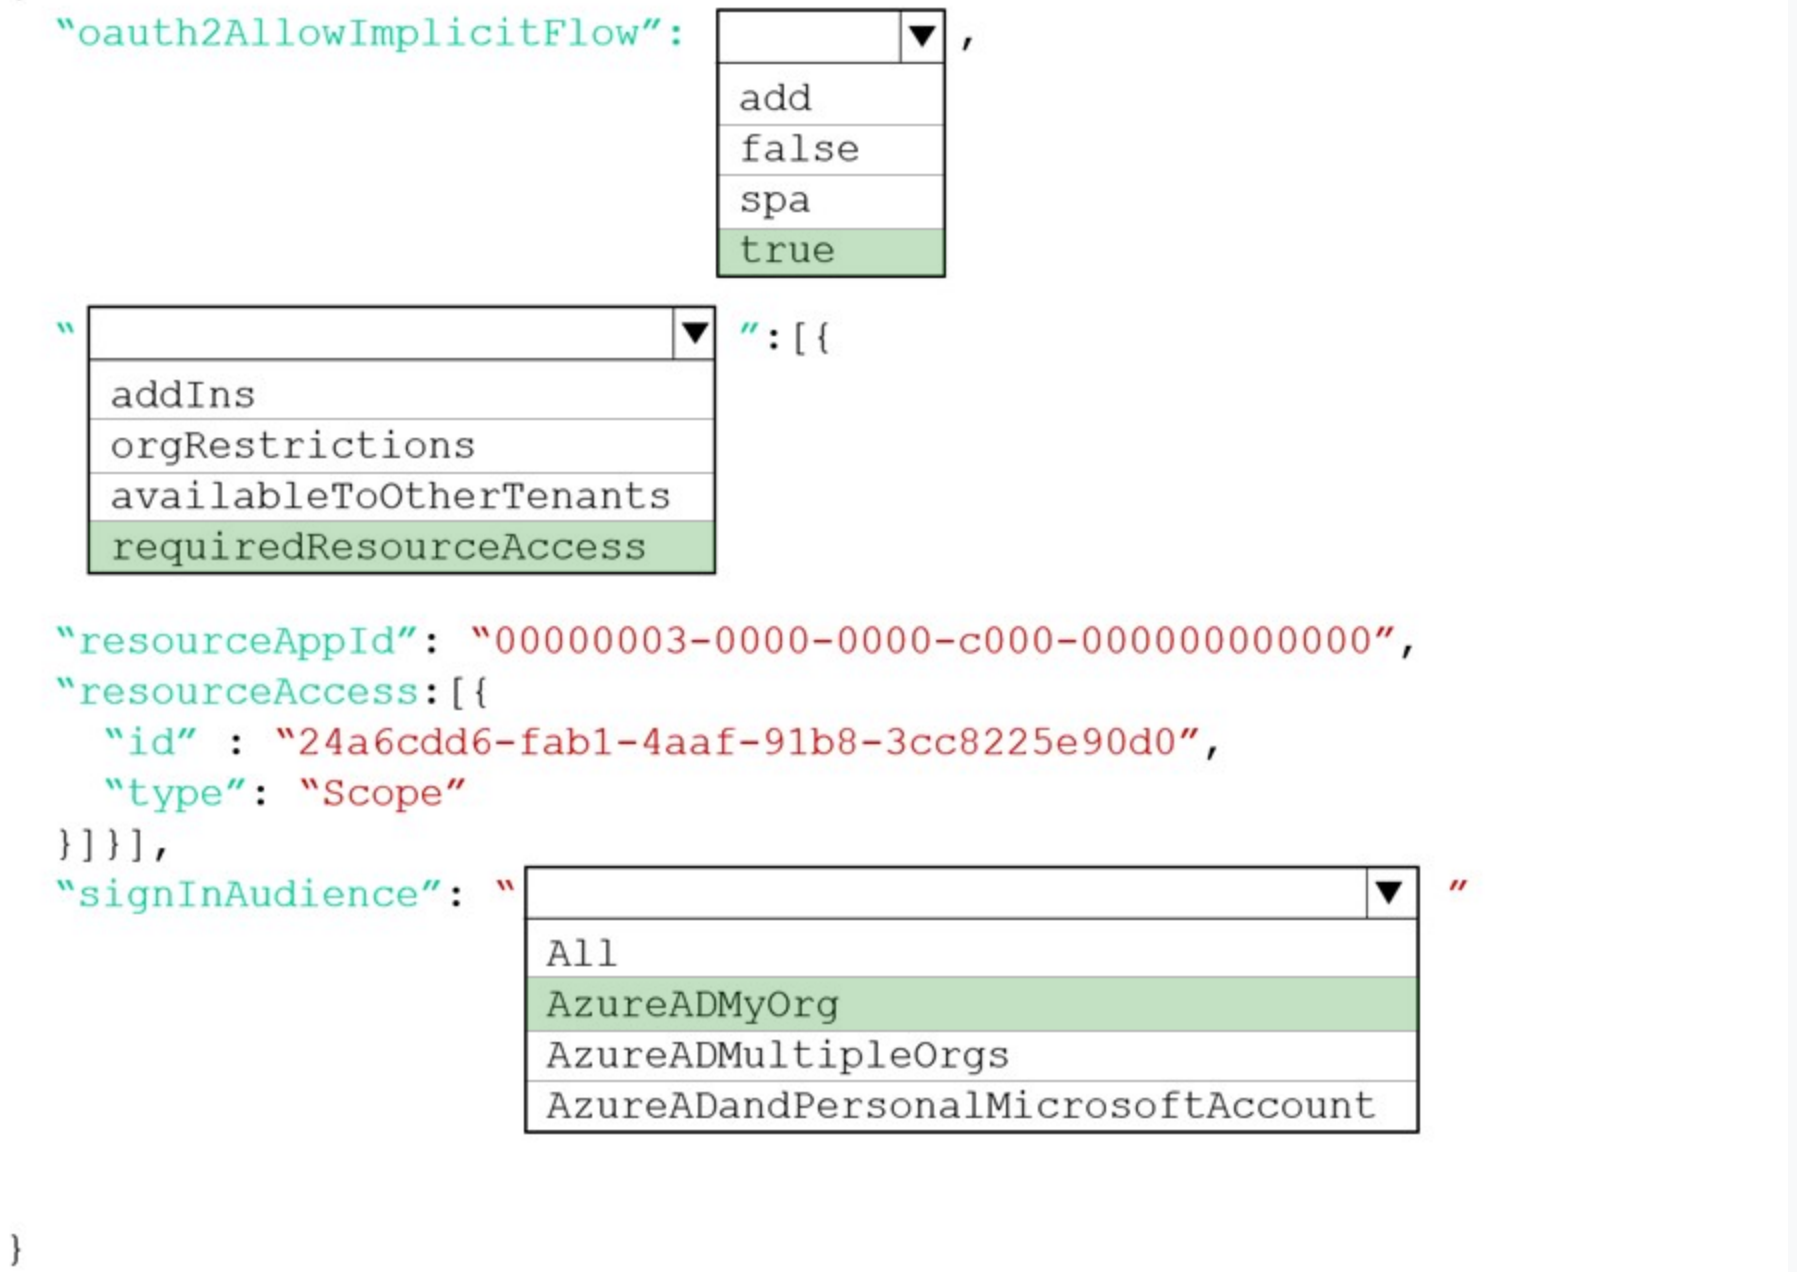

Box 1: true -

The oauth2AllowImplicitFlow attribute Specifies whether this web app can request OAuth2.0 implicit flow access tokens. The default is false. This flag is used for browser-based apps, like JavaScript single-page apps.

In implicit flow, the app receives tokens directly from the Azure Active Directory (Azure AD) authorize endpoint, without any server-to-server exchange. All authentication logic and session handling is done entirely in the JavaScript client with either a page redirect or a pop-up box.

Box 2: requiredResourceAccess -

With dynamic consent, requiredResourceAccess drives the admin consent experience and the user consent experience for users who are using static consent.

However, this parameter doesn't drive the user consent experience for the general case. resourceAppId is the unique identifier for the resource that the app requires access to. This value should be equal to the appId declared on the target resource app. resourceAccess is an array that lists the OAuth2.0 permission scopes and app roles that the app requires from the specified resource. Contains the id and type values of the specified resources.

Example:

"requiredResourceAccess": [

{

"resourceAppId": "00000002-0000-0000-c000-000000000000",

"resourceAccess": [

{

"id": "311a71cc-e848-46a1-bdf8-97ff7156d8e6",

"type": "Scope"

}

]

}

],

Incorrect Answers:

☞ The legacy attribute availableToOtherTenants is no longer supported.

☞ The addIns attribute defines custom behavior that a consuming service can use to call an app in specific contexts. For example, applications that can render file streams may set the addIns property for its "FileHandler" functionality. This parameter will let services like Microsoft 365 call the application in the context of a document the user is working on.

Example:
"addIns": [
{
"id": "968A844F-7A47-430C-9163-07AE7C31D407",
"type":" FileHandler",
"properties": [
{
"key": "version",
"value": "2"
}
]
}
],

Box 3: AzureADMyOrg -
The signInAudience attribute specifies what Microsoft accounts are supported for the current application. Supported values are:
☞ AzureADMyOrg - Users with a Microsoft work or school account in my organization's Azure AD tenant (for example, single tenant)
☞ AzureADMultipleOrgs - Users with a Microsoft work or school account in any organization's Azure AD tenant (for example, multi-tenant)
☞ AzureADandPersonalMicrosoftAccount - Users with a personal Microsoft account, or a work or school account in any organization's Azure AD tenant
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-implicit-grant-flow

vvlads Highly Voted 1 year, 6 months ago
I think the last one should be AzureADMultipleOrgs
upvoted 21 times

SaintBahamut 1 year, 6 months ago
same, question states that AD data is from external organisation
upvoted 4 times

OPT_001122 12 months ago
I think Gmishra88 is correct.
Question states that - "Users must be able to authenticate and access Microsoft Graph by using their own company's Azure AD instance."

So it should be AzureADMyOrg
upvoted 2 times

the_odd_one Most Recent 1 week, 5 days ago
Would it be an option to use Azure AD B2B collaboration to invite guest users from tenant B (Microsoft Graph tenant) to access the application in tenant A (application tenant). Would then set signInAudience to "AzureADMyOrg" to restrict sign-ins to users from the specific Azure AD tenant associated with the application (tenant A).
upvoted 1 times

plusJoyed 4 months, 1 week ago
got on my exam 6/22/2023
upvoted 4 times

apparaog99 10 months ago
On my exam 12/31
upvoted 2 times

AWAPass 10 months ago
On exam 29/12/2022
upvoted 1 times

gmishra88 1 year ago
Given answers are correct.
SigninAudience: Specifies what Microsoft accounts are supported for the current application. The users login with the current organizations AD accounts. This does not talk about the external access. And that is probably the red herring Microsoft added to the mix so that you will make a mistake and they win. Microsoft always wins
upvoted 4 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000.
upvoted 3 times

Knightie 1 year, 1 month ago

⇨ AzureADMyOrg - Users with a Microsoft work or school account in my organization's Azure AD tenant (for example, single tenant)
⇨ AzureADMultipleOrgs - Users with a Microsoft work or school account in any organization's Azure AD tenant (for example, multi-tenant)
⇨ AzureADandPersonalMicrosoftAccount - Users with a personal Microsoft account, or a work or school account in any organization's Azure AD tenant [Answer]

upvoted 3 times

Knightie 1 year, 1 month ago

first one is the current Tenant, second one is any Tenant under your Org, third one is any Tenant under Any Org... so only the 3rd make sense.

upvoted 1 times

finnishr 1 year, 1 month ago

That makes no sense. Difference between AzureADMultipleOrgs and AzureADandPersonalMicrosoftAccount is that AzureADandPersonalMicrosoftAccount also includes personal Microsoft accounts which we dont want to include.

upvoted 2 times

sghaha 1 year, 6 months ago

in Korean
https://docs.microsoft.com/ko-kr/azure/active-directory/develop/reference-app-manifest

upvoted 1 times

Question #44 *Topic 4*

You manage a data processing application that receives requests from an Azure Storage queue.

You need to manage access to the queue. You have the following requirements:

☞ Provide other applications access to the Azure queue.

☞ Ensure that you can revoke access to the queue without having to regenerate the storage account keys.

☞ Specify access at the queue level and not at the storage account level.

Which type of shared access signature (SAS) should you use?

A. Service SAS with a stored access policy

B. Account SAS

C. User Delegation SAS

D. Service SAS with ad hoc SAS

Correct Answer: *A*

A service SAS is secured with the storage account key. A service SAS delegates access to a resource in only one of the Azure Storage services:

Blob storage,

Queue storage, Table storage, or Azure Files.

Stored access policies give you the option to revoke permissions for a service SAS without having to regenerate the storage account keys.

Incorrect Answers:

Account SAS: Account SAS is specified at the account level. It is secured with the storage account key.

User Delegation SAS: A user delegation SAS applies to Blob storage only.

Reference:

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

Community vote distribution

A (100%)

sghaha Highly Voted 1 year, 6 months ago

Argee A

upvoted 15 times

gmishra88 Highly Voted 1 year ago

User delegation SAS is only for blob storage? These random restrictions that microsoft creates just makes me so happy. It is a maze of that works with this but not if that and only if those cases. Why do they expect human beings to remember these unfortunate mistakes Microsoft makes

upvoted 6 times

rolling_potato_ Most Recent 10 months ago

Selected Answer: A

Was on my exam today (03-01-2023)

upvoted 4 times

rotimislaw 10 months ago

Why D. Service SAS with ad hoc SAS is incorrect?

upvoted 1 times

tzuyichao 7 months, 4 weeks ago

"Specify access at the queue level and not at the storage account level."

upvoted 1 times

surprise0011 6 months, 2 weeks ago

there is not a valid explanation. Question says: Service SAS with ad hoc SAS
So it is at service level not account level
It is not a valid answer because ad-hoc signature cannot be invalidated. With access policy you can change access policy

upvoted 3 times

OPT_001122 12 months ago

Selected Answer: A

A - corrct

upvoted 2 times

adrian_g 1 year, 1 month ago

I would go with B. A stored access policy is defined on a resource container.
"An account SAS delegates access to resources in one or more of the storage services. All of the operations available via a service or user

delegation SAS are also available via an account SAS."
"Service SAS with stored access policy. A stored access policy is defined on a resource container, which can be a blob container, table, queue, or file share. "
from: https://learn.microsoft.com/en-us/azure/storage/common/storage-sas-overview

upvoted 1 times

ArturKon 1 year, 1 month ago

Selected Answer: A

Looks like queues can only have SAS described in A.

upvoted 1 times

Question #45 *Topic 4*

HOTSPOT -
You are developing an application to store and retrieve data in Azure Blob storage. The application will be hosted in an on-premises virtual machine (VM). The
VM is connected to Azure by using a Site-to-Site VPN gateway connection. The application is secured by using Azure Active Directory (Azure AD) credentials.
The application must be granted access to the Azure Blob storage account with a start time, expiry time, and read permissions. The Azure Blob storage account access must use the Azure AD credentials of the application to secure data access. Data access must be able to be revoked if the client application security is breached.
You need to secure the application access to Azure Blob storage.
Which security features should you use? To answer select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Component	Security Feature
Application (Client)	Storage Account Access Key System-assigned Managed Identity Shared access signature (SAS) token
Azure Storage (Server)	Stored Access Policy User-assigned Managed Identity Cross-Origin Resource Sharing (CORS)

Correct Answer:

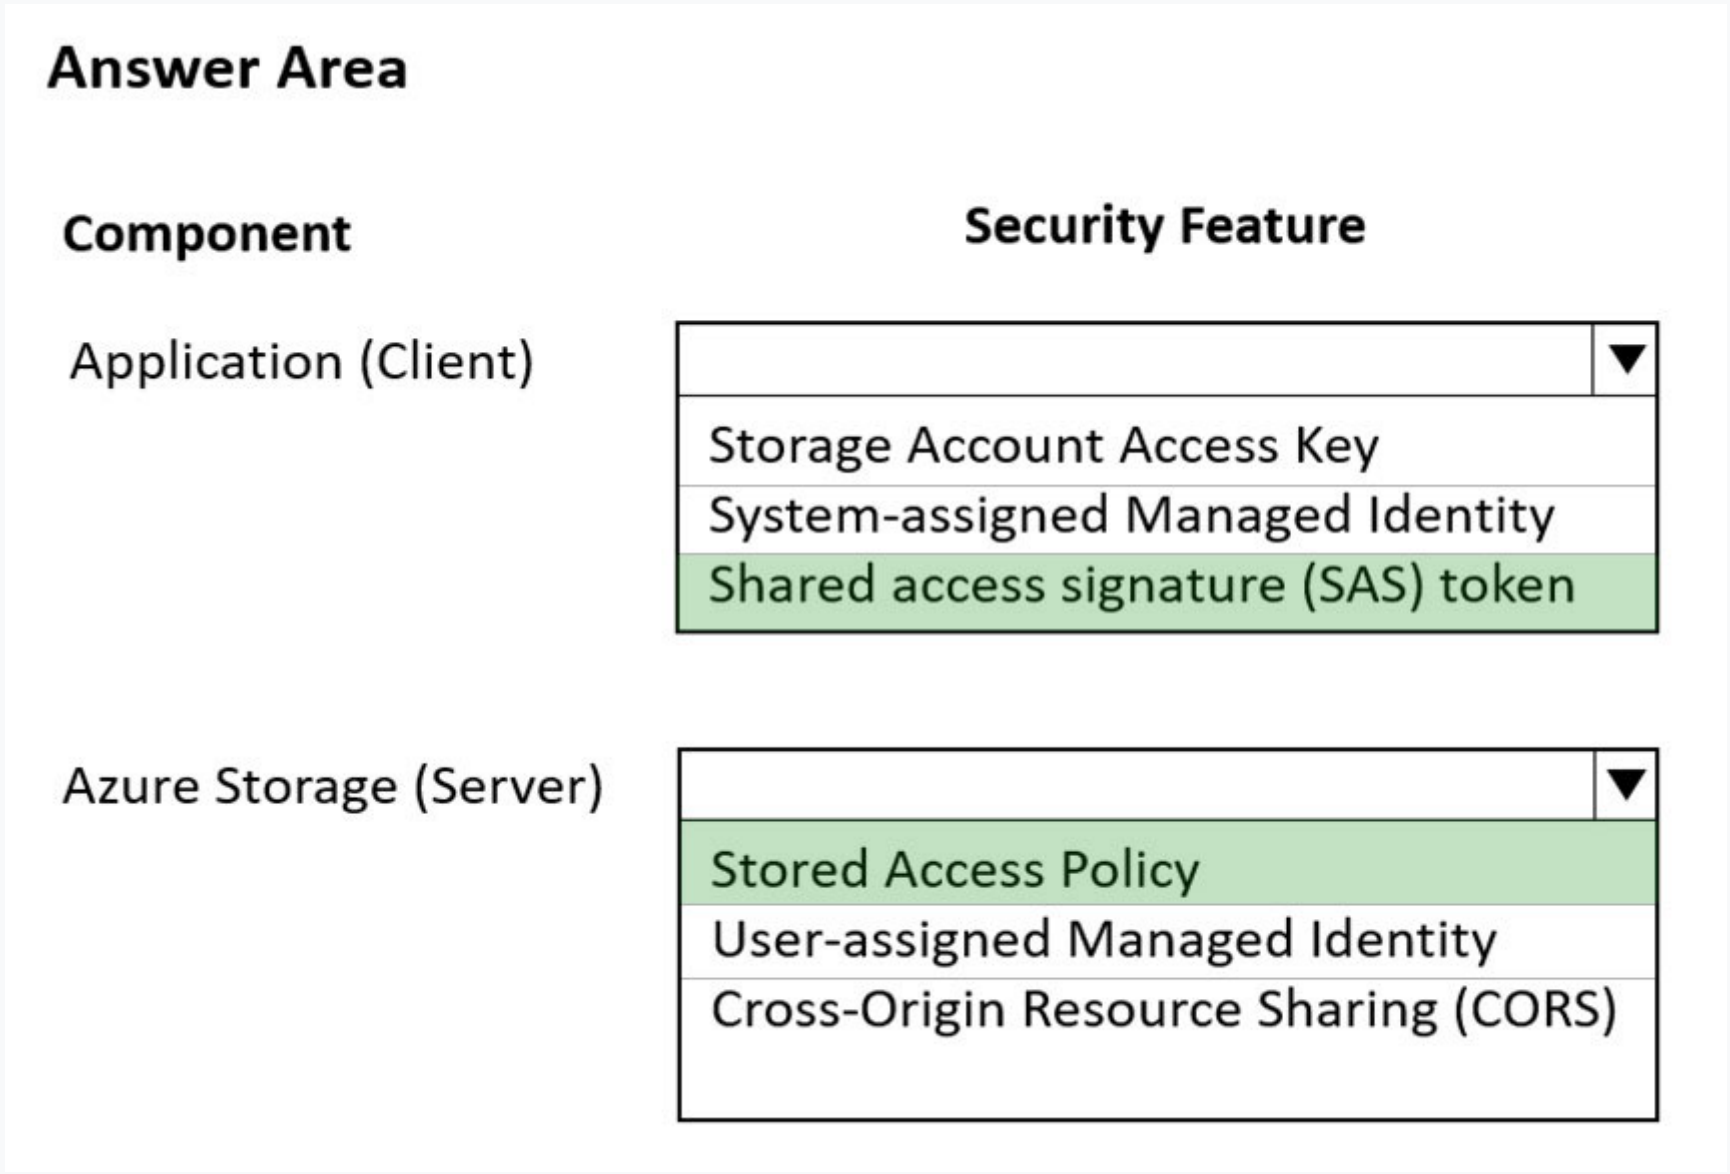

Box 1: Shared access signature (SAS) token
When your application design requires shared access signatures for access to Blob storage, use Azure AD credentials to create a user delegation SAS when possible for superior security.

Box 2: Stored access policy -
Stored access policies give you the option to revoke permissions for a service SAS without having to regenerate the storage account keys.
A shared access signature can take one of the following two forms:
☞ Service SAS with stored access policy. A stored access policy is defined on a resource container, which can be a blob container, table, queue, or file share.
The stored access policy can be used to manage constraints for one or more service shared access signatures. When you associate a service SAS with a stored access policy, the SAS inherits the constraints ـ€" the start time, expiry time, and permissions ـ€" defined for the stored access policy.
☞ Ad hoc SAS.
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

coffecold Highly Voted 1 year ago
storage account access keys : similar to a root password for your storage account, not for distrbution
system assigned managed identity : only for use within of services in azure
SAS token : right answer

Stored Access Policy : right answer
user assigned managed identity : only for use within of services in azure, this must connect to a server outside
CORS protectection : Site-to-Site VPN gateway, there is no other domain involved
upvoted 24 times

macobuzi Most Recent 2 months, 1 week ago
Why not System-Assigned Managed Identity + Stored Access Policy? I think it's also a valid option.
upvoted 1 times

vmakharashvili 8 months ago
Correct
upvoted 1 times

OPT_001122 10 months, 2 weeks ago
SAS token
Stored Access Policy

The answer is correct

upvoted 1 times

finnishr 1 year, 1 month ago

The answer is correct since none of the other possible combinations make sense. Even though service level SAS doesnt use Azure AD credentials.

upvoted 2 times

gmishra88 1 year ago

SAS with user-delegation works with AD roles

upvoted 2 times

finnishr 1 year, 1 month ago

Stored Access Policies definitely wont work. I'd pick SAS and CORS.

upvoted 1 times

ArturKon 1 year, 1 month ago

"Data access must be able to be revoked if the client application security is breached." - isn't that a clue to use access policy? Typical SAS is not revokable itself.

upvoted 2 times

sghaha 1 year, 6 months ago

in Korean
https://docs.microsoft.com/ko-kr/azure/storage/common/storage-sas-overview

upvoted 1 times

eazy_breezy_jeezy 1 year, 6 months ago

A user delegation SAS makes sense for the "Application" answer, but I don't think any of the provided options for the Azure Storage side work. Stored Access Policies only work with a Service-level SAS, so that wouldn't make sense with a user delegation SAS. Seems like the closest option there would have been something like "RBAC role assignments" or "user delegation keys", just to spit ball.
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview#user-delegation-sas
https://docs.microsoft.com/en-us/rest/api/storageservices/create-user-delegation-sas#revoke-a-user-delegation-sas

upvoted 1 times

Question #46 *Topic 4*

You are building a web application that uses the Microsoft identity platform for user authentication.
You are implementing user identification for the web application.
You need to retrieve a claim to uniquely identify a user.
Which claim type should you use?

A. aud

B. nonce

C. oid

D. idp

Correct Answer: *C*
oid -The object identifier for the user in Azure AD. This value is the immutable and non-reusable identifier of the user. Use this value, not email, as a unique identifier for users; email addresses can change. If you use the Azure AD Graph API in your app, object ID is that value used to query profile information.
Incorrect:
Not A: aud - Who the token was issued for. This will be the application's client ID.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/multitenant-identity/claims

Community vote distribution
C (100%)

serpevi Highly Voted 1 year, 1 month ago
Got this in 09/22 , went with oid, score 927.
upvoted 11 times

OPT_001122 Highly Voted 12 months ago
Selected Answer: C
oid - correct
upvoted 6 times

Ciupaz Most Recent 1 month ago
The oid claim contains a unique identifier for the user, and it is typically used to uniquely identify users across different sessions and applications
upvoted 1 times

unraval 1 month, 2 weeks ago
Determine if the token subject, such as the user (or application itself for an app-only token), is authorized.
You can either check for specific sub or oid claims.
https://learn.microsoft.com/en-us/azure/active-directory/develop/claims-validation
upvoted 1 times

devex 2 months ago
Selected Answer: C
C: oid

"The immutable identifier for the requestor, which is the verified identity of the user or service principal."

Reference: https://learn.microsoft.com/en-us/azure/active-directory/develop/access-token-claims-reference#payload-claims
upvoted 2 times

JH81 4 months ago
Got this on 6/28/2023 and passed with 850. Answer is correct.
upvoted 1 times

Shashakg 4 months, 3 weeks ago
Selected Answer: C
Oid claims will always remain same.
upvoted 1 times

elequiel 1 year ago
Got it in exam 20/10/2022

upvoted 4 times

finnishr 1 year, 1 month ago

Correct

upvoted 4 times

Question #47 *Topic 4*

You are developing an Azure Function that calls external APIs by providing an access token for the API. The access token is stored in a secret named token in an
Azure Key Vault named mykeyvault.
You need to ensure the Azure Function can access to the token. Which value should you store in the Azure Function App configuration?

A. KeyVault:mykeyvault;Secret:token

B. App:Settings:Secret:mykeyvault:token

C. AZUREKVCONNSTR_ https://mykeyveult.vault.ezure.net/secrets/token/

D. @Microsoft.KeyVault(SecretUri=https://mykeyvault.vault.azure.net/secrets/token/)

Correct Answer: *D*
Add Key Vault secrets reference in the Function App configuration.
Syntax: @Microsoft.KeyVault(SecretUri={copied identifier for the username secret})
Reference:
https://daniel-krzyczkowski.github.io/Integrate-Key-Vault-Secrets-With-Azure-Functions/

Community vote distribution
D (94%) 6%

lorenaizzo Highly Voted 1 year, 1 month ago

Selected Answer: D
It's right! D.
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references?tabs=azure-powershell
upvoted 14 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000.
upvoted 3 times

KingChuang 10 months, 1 week ago
On my exam 2022-12-26.
Chose： Given Answer
upvoted 1 times

juanckar Most Recent 3 months, 3 weeks ago
This was on the exam (July 2023). Went with proposed. Scored 917
upvoted 2 times

DarthVee 5 months, 3 weeks ago
According to ChatGPT its A:

Option A is the correct value to store in the Azure Function App configuration.

The format for referencing a secret stored in an Azure Key Vault is "KeyVault:{key vault name};Secret:{secret name}". Therefore, the correct value to store in the Azure Function App configuration to access the token stored in a secret named "token" in an Azure Key Vault named "mykeyvault" is "KeyVault:mykeyvault;Secret:token".

Option B is not a valid format for referencing an Azure Key Vault secret.

Option C is not a valid format for referencing an Azure Key Vault secret. This value appears to be an Azure Key Vault connection string.

Option D is a Key Vault reference expression, which is used in Azure Resource Manager templates to reference a Key Vault secret or certificate. This is not the correct format to store in the Azure Function App configuration.
upvoted 1 times

MikeAWS 2 weeks, 3 days ago
chatgpt changed his mind :), its saying D now
upvoted 1 times

surprise0011 6 months, 2 weeks ago
If you wonder why not A because syntax is incomplete. Should be like:
@Microsoft.KeyVault(VaultName=myvault;SecretName=mysecret)
https://learn.microsoft.com/en-us/azure/app-service/app-service-key-vault-references?tabs=azure-cli#reference-syntax
upvoted 1 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.
upvoted 3 times

OPT_001122 12 months ago

Selected Answer: D

Given answer is correct
upvoted 1 times

coffecold 1 year ago

Selected Answer: A

The Azure Function must be able to access to the token (inside the code)
All it needs is a vaultname and a secret name.

var secretClient = new SecretClient(new Uri("https://<VaultName>.vault.azure.net"),new DefaultAzureCredential());
var secret = await secretClient.GetSecretAsync("<SecretName>");
the <SecretName> and the <Vaultname> come from your app settings
upvoted 1 times

nvtienanh 1 year ago

Selected Answer: D

Given answer is correct
upvoted 2 times

Question #48 *Topic 4*

A company maintains multiple web and mobile applications. Each application uses custom in-house identity providers as well as social identity providers.
You need to implement single sign-on (SSO) for all the applications.
What should you do?

A. Use Azure Active Directory B2C (Azure AD B2C) with custom policies.

B. Use Azure Active Directory B2B (Azure AD B2B) and enable external collaboration.

C. Use Azure Active Directory B2C (Azure AD B2C) with user flows.

D. Use Azure Active Directory B2B (Azure AD B2B).

Correct Answer: *B*
You can add Google as an identity provider for B2B guest users.
Federation with SAML/WS-Fed identity providers for guest users.
Make sure your organization's external collaboration settings are configured such that you're allowed to invite guests.
Note 1: As a user who is assigned any of the limited administrator directory roles, you can use the Azure portal to invite B2B collaboration users. You can invite guest users to the directory, to a group, or to an application. After you invite a user through any of these methods, the invited user's account is added to Azure
Active Directory (Azure AD), with a user type of Guest.
Note 2: Direct federation in Azure Active Directory is now referred to as SAML/WS-Fed identity provider (IdP) federation.
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/external-identities/google-federation https://docs.microsoft.com/en-us/azure/active-directory/external-identities/add-users-administrator

Community vote distribution
A (79%) C (21%)

tcybu Highly Voted 1 year, 1 month ago
The provided answer is wrong!
Based on this article: https://docs.microsoft.com/en-us/azure/active-directory/external-identities/external-collaboration-settings-configure
External collaboration settings let you specify what roles in your organization can invite external users for B2B collaboration. These settings also include options for allowing or blocking specific domains, and options for restricting what external guest users can see in your Azure AD directory.
So, you use B2B external collaboration to invite guests into your Azure AD tenant.

I vote for Custom Policies. Both Custom Policies and User Flows support external identity providers, but because of required custom in-house providers support, I'd choose Custom Policies over the User Flows - see https://docs.microsoft.com/en-us/azure/active-directory-b2c/user-flow-overview#comparing-user-flows-and-custom-policies
upvoted 14 times

adilkhan Most Recent 7 months ago
A is correct
upvoted 3 times

JOSEEVILLASMIL 7 months, 1 week ago
Given answer is correct, remember that the question says "company maintains many aplications" and "those aplication have their own in house login providers" so there is a bussiness with other bussiness, is B2B, no B2C.

And B2B with external colaboration says:
"You can also enable self-service sign-up user flows to let external users sign up for apps or resources themselves."

https://learn.microsoft.com/en-us/azure/active-directory/external-identities/what-is-b2b
upvoted 2 times

adilkhan 9 months, 1 week ago
Should be A A company maintains multiple web and mobile applications. Each application uses custom in-house identity providers as well as social identity providers.
upvoted 2 times

mitu7 10 months ago
Selected Answer: C
"C. Use Azure Active Directory B2C (Azure AD B2C) with user flows."

Azure AD B2C is a cloud-based identity management service that enables you to customize and control how users sign up, sign in, and manage their profiles when interacting with your applications. With Azure AD B2C, you can create user flows that define the steps in the authentication and

authorization process for your applications, and you can use these user flows to enable SSO across your applications.

User flows allow you to customize the authentication experience for your users, including specifying the identity providers that you want to use, such as custom in-house identity providers or social identity providers. You can also customize the user interface, branding, and policies that apply to each user flow.

References: https://learn.microsoft.com/en-us/azure/active-directory-b2c/user-flow-overview

upvoted 3 times

guchao2000 9 months, 1 week ago

It's A
The message is "... uses custom in-house identity providers", this requires Azure AD B2C with custom policies.

upvoted 1 times

Yumi21 12 months ago

For Azure AD B2C check this: https://learn.microsoft.com/en-us/azure/active-directory-b2c/overview

upvoted 1 times

OPT_001122 12 months ago

Selected Answer: A

A - correct answer

upvoted 4 times

RochaG2 1 year ago

Selected Answer: A

A for sure

upvoted 4 times

finnishr 1 year, 1 month ago

Selected Answer: A

A is correct 100%

upvoted 3 times

Knightie 1 year, 1 month ago

The Answer seems Correct. B2C seems for external totally, so no mention of the existing AAD authetication to the app.
https://docs.microsoft.com/en-us/azure/active-directory/external-identities/external-identities-overview?bc=%2Fazure%2Factive-directory-b2c%2Fbread%2Ftoc.json&toc=%2Fazure%2Factive-directory-b2c%2FTOC.json

upvoted 1 times

jeanfmc 1 year, 1 month ago

A
https://docs.microsoft.com/en-us/azure/active-directory-b2c/custom-policy-reference-sso

upvoted 4 times

Question #49 *Topic 4*

You develop a Python application for image rendering that uses GPU resources to optimize rendering processes. You deploy the application to an Azure
Container Instances (ACI) Linux container.
The application requires a secret value to be passed when the container is started. The value must only be accessed from within the container.
You need to pass the secret value.
What are two possible ways to achieve this goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Create an environment variable Set the secureValue property to the secret value.

B. Add the secret value to the container image. Use a managed identity.

C. Add the secret value to the application code Set the container startup command.

D. Add the secret value to an Azure Blob storage account. Generate a SAS token.

E. Mount a secret volume containing the secret value in a secrets file.

Correct Answer: *AE*
A: Secure environment variables -
Another method (another than a secret volume) for providing sensitive information to containers (including Windows containers) is through the use of secure environment variables.
E: Use a secret volume to supply sensitive information to the containers in a container group. The secret volume stores your secrets in files within the volume, accessible by the containers in the container group. By storing secrets in a secret volume, you can avoid adding sensitive data like SSH keys or database credentials to your application code.
Reference:
https://docs.microsoft.com/en-us/azure/container-instances/container-instances-volume-secret

Community vote distribution
AE (100%)

osagiestar Highly Voted 1 year ago
Selected Answer: AE
A and E
upvoted 8 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000.
upvoted 2 times

elequiel Highly Voted 1 year ago
Got it in exam 20/10/2022
upvoted 7 times

BaoNguyen2411 Most Recent 3 months ago
got this question on 06-29-2023
upvoted 3 times

lorenaizzo 1 year, 1 month ago
Selected Answer: AE
Secure values
Objects with secure values are intended to hold sensitive information like passwords or keys for your application. Using secure values for environment variables is both safer and more flexible than including it in your container's image. Another option is to use secret volumes, described in Mount a secret volume in Azure Container Instances.....

https://docs.microsoft.com/en-us/azure/container-instances/container-instances-environment-variables
upvoted 5 times

Question #50 *Topic 4*

You are developing a user portal for a company.
You need to create a report for the portal that lists information about employees who are subject matter experts for a specific topic. You must ensure that administrators have full control and consent over the data.
Which technology should you use?

A. Microsoft Graph data connect

B. Microsoft Graph API

C. Microsoft Graph connectors

Correct Answer: *A*
Data Connect grants a more granular control and consent model: you can manage data, see who is accessing it, and request specific properties of an entity. This enhances the Microsoft Graph model, which grants or denies applications access to entire entities.
Microsoft Graph Data Connect augments Microsoft Graph's transactional model with an intelligent way to access rich data at scale. The data covers how workers communicate, collaborate, and manage their time across all the applications and services in Microsoft 365.
Incorrect:
Not B: The Microsoft Graph API is a RESTful web API that enables you to access Microsoft Cloud service resources. After you register your app and get authentication tokens for a user or service, you can make requests to the Microsoft Graph API.
A simplistic definition of a Graph API is an API that models the data in terms of nodes and edges (objects and relationships) and allows the client to interact with multiple nodes in a single request.
Not C: Microsoft Graph connectors, your organization can index third-party data so that it appears in Microsoft Search results.
With Microsoft Graph connectors, your organization can index third-party data so that it appears in Microsoft Search results.
Reference:
https://docs.microsoft.com/en-us/graph/data-connect-concept-overview

Community vote distribution
A (83%) B (17%)

nvtienanh **Highly Voted** 1 year ago

Selected Answer: A
https://learn.microsoft.com/en-us/graph/overview#access-microsoft-graph-data-at-scale:~:text=Find%20you%20the%20closest%20expert%20on%20a%20topic%20to%20you%20in%20your%20organization
upvoted 12 times

Morgist **Most Recent** 1 month, 1 week ago

Selected Answer: A
Microsoft Graph Data Connect provides access to many users or groups so it meets the requirements
upvoted 1 times

oreakinodidi 3 months, 2 weeks ago

Selected Answer: B
The Microsoft Graph API provides a comprehensive set of REST APIs that allows you to access and manipulate a wide range of data and resources within Microsoft 365 and Azure Active Directory. It offers powerful querying capabilities and supports various operations to retrieve and filter employee information based on specific criteria such as topic expertise.

While Microsoft Graph data connect and Microsoft Graph connectors also have their use cases, they are not directly applicable to the requirement of creating a report for the user portal. Microsoft Graph data connect is focused on data integration and export scenarios, while Microsoft Graph connectors are used for connecting external data sources to Microsoft Search.

Therefore, for the given scenario, the Microsoft Graph API is the most suitable option to retrieve and present the required employee information in the user portal report.
upvoted 1 times

DarthVee 5 months, 3 weeks ago

Selected Answer: B
In summary, Microsoft Graph API provides the necessary security mechanisms to ensure that administrators have full control and consent over the data accessed by your application.
upvoted 2 times

karthikwarrior 8 months, 1 week ago

The people who are commenting the right ans have paid subscription (Contributor role) as bec i am one of them and you can notice, comments are reduced from others who doesnt have access to this role. they cant access this content..

upvoted 4 times

macobuzi 2 months, 1 week ago

Which is sad since we might not have enough information to decide the right answer!

upvoted 1 times

coffecold 1 year ago

Big Microsoft Brother is watching you with some kind of shady tool : 'Data Connect'

upvoted 4 times

lorenaizzo 1 year, 1 month ago

Selected Answer: A

https://docs.microsoft.com/en-us/graph/overview

upvoted 2 times

Question #51 *Topic 4*

HOTSPOT -

You are a developer building a web site using a web app. The web site stores configuration data in Azure App Configuration.

Access to Azure App Configuration has been configured to use the identity of the web app for authentication. Security requirements specify that no other authentication systems must be used.

You need to load configuration data from Azure App Configuration.

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:

Answer Area

```
Host .CreateDefaultBuilder(args)
.ConfigureWebHostDefaults(wb =>
{

  wb.ConfigureAppConfiguration((hc, config) =>
  {
    var settings = config.Build();
    config.                                    (options =>
```

AddAzureKeyVault
DefaultAzureCredential
ChainedTokenCredential
ManagedIdentityCredential
AddAzureAppConfiguration

```
      options.Connect(new Uri(settings["AppConfig:Endpoint"]),
        new                                    ()));
```

AddAzureKeyVault
DefaultAzureCredential
ChainedTokenCredential
ManagedIdentityCredential
AddAzureAppConfiguration

```
  });
})
```

Correct Answer:

Answer Area

```
Host .CreateDefaultBuilder(args)
.ConfigureWebHostDefaults(wb =>
{

  wb.ConfigureAppConfiguration((hc, config) =>
  {
    var settings = config.Build();
    config.                                    (options =>
```

AddAzureKeyVault
DefaultAzureCredential
ChainedTokenCredential
ManagedIdentityCredential
AddAzureAppConfiguration

```
      options.Connect(new Uri(settings["AppConfig:Endpoint"]),
        new                                    ());
```

AddAzureKeyVault
DefaultAzureCredential
ChainedTokenCredential
ManagedIdentityCredential
AddAzureAppConfiguration

```
  });
})
```

Box 1: AddAzureAppConfiguration -

Load data from App Configuration, code example:

public static IHostBuilder CreateHostBuilder(string[] args) =>

Host.CreateDefaultBuilder(args)

.ConfigureWebHostDefaults(webBuilder =>

webBuilder.ConfigureAppConfiguration((hostingContext, config) =>

{

var settings = config.Build();

config.AddAzureAppConfiguration(options =>

{

Etc.

Box 2: ManagedIdentityCredential

Use managed identities to access App Configuration

If you want to use a user-assigned managed identity, be sure to specify the clientId when creating the ManagedIdentityCredential.

config.AddAzureAppConfiguration(options =>

{

options.Connect(new Uri(settings["AppConfig:Endpoint"]), new ManagedIdentityCredential("<your_clientId>"))

});

Full code sample:

public static IHostBuilder CreateHostBuilder(string[] args) =>

Host.CreateDefaultBuilder(args)

.ConfigureWebHostDefaults(webBuilder =>

webBuilder.ConfigureAppConfiguration((hostingContext, config) =>

{

var settings = config.Build();

config.AddAzureAppConfiguration(options =>

options.Connect(new Uri(settings["AppConfig:Endpoint"]), new ManagedIdentityCredential()));

})

.UseStartup<Startup>());

Reference:

https://docs.microsoft.com/en-us/azure/azure-app-configuration/howto-integrate-azure-managed-service-identity?tabs=core5x&pivots=framework- dotnet

coffecold Highly Voted 1 year ago

Answer is right.
"Security requirements specify that no other authentication systems must be used"
So you cannot use DefaultAzureCredential, risk of other authentication systems being used.
otherwise you could have used DefaultAzureCredential for ease

upvoted 20 times

serpevi Highly Voted 1 year, 1 month ago

Got this in 09/22 , went with the answers given, score 927.

upvoted 14 times

OPT_001122 Most Recent 10 months, 3 weeks ago

Ans is correct
AddAzureAppConfiguration
ManagedIdentityCredential

upvoted 5 times

Question #52 *Topic 4*

HOTSPOT -

You are building an application that stores sensitive customer data in Azure Blob storage. The data must be encrypted with a key that is unique for each customer.

If the encryption key has been corrupted it must not be used for encryption.

You need to ensure that the blob is encrypted.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
Uri blobUri = ... ; TokenCredential c = ...
byte[] key = ...; string verify = ...
var x = new [                                                    ];
```

- AesManaged(key)
- AsnEncodedData(key)
- CustomerProvidedKey(key)
- BlobContainerEncryptionScopeOptions { DefaultEncryptionScope = key }

```
if ( [                                                           ] ) {
```

- x.IV == verify
- x.RawData == verify
- x.Encryption == verify
- x.PreventEncryptionScopeOverride == verify

```
  var o = new BlobClientOptions()
  {
    [                                                            ] = x
```

- Version
- Transport
- EncryptionScope
- CustomerProvidedKey

```
  };
  var blobClient = new BlobClient(blobUri, c, o);
}
```

Correct Answer:

Answer Area

```
Uri blobUri = ... ; TokenCredential c = ...
byte[] key = ...; string verify = ...
var x = new [                                                    ];
```

- AesManaged(key)
- AsnEncodedData(key)
- CustomerProvidedKey(key)
- BlobContainerEncryptionScopeOptions { DefaultEncryptionScope = key }

```
if ( [                                                           ] ) {
```

- x.IV == verify
- x.RawData == verify
- x.Encryption == verify
- x.PreventEncryptionScopeOverride == verify

```
  var o = new BlobClientOptions()
  {
    [                                                            ] = x
```

- Version
- Transport
- EncryptionScope
- CustomerProvidedKey

```
  };
  var blobClient = new BlobClient(blobUri, c, o);
}
```

Box 1: CustomerProvidedKey(key)

The data must be encrypted with a key that is unique for each customer.

Sample code:

async static Task UploadBlobWithClientKey(Uri blobUri,

```
Stream data,
byte[] key,
string keySha256)
{
// Create a new customer-provided key.
// Key must be AES-256.
var cpk = new CustomerProvidedKey(key);

Box 2: Encryption -
CustomerProvidedKey.EncryptionKey Property
Sample code continued:
// Check the key's encryption hash.
if (cpk.EncryptionKeyHash != keySha256)
{
throw new InvalidOperationException("The encryption key is corrupted.");
}

Box 3: CustomerProvidedKey -
Sample code continued;
// Specify the customer-provided key on the options for the client.
BlobClientOptions options = new BlobClientOptions()
{

CustomerProvidedKey = cpk -
};
// Create the client object with options specified.
BlobClient blobClient = new BlobClient(
blobUri,
new DefaultAzureCredential(),
options);
Incorrect:
* Version - Gets the BlobClientOptions.ServiceVersion of the service API used when making requests.
Transport - The HttpPipelineTransport to be used for this client.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-customer-provided-key
```

Praks13 Highly Voted 1 year, 1 month ago

This looks correct!
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-customer-provided-key

upvoted 13 times

gmishra88 Highly Voted 1 year ago

Shouldn't it be EncryptionKeyHash based on the link provided?

upvoted 11 times

NPE_ Most Recent 5 days, 18 hours ago

AZ204 should be renamed as Azure .NET/c# Developer

upvoted 2 times

NightshadeRC 3 months, 1 week ago

Had this question in today's exam: 2023-07-26
Also the exam did have the value as EncryptedKeyHash

upvoted 6 times

Elbanna 1 month ago

Does All exam questions from here?

upvoted 1 times

willchenxa 1 year, 1 month ago

The second box probably should be EncryptionKey, CustomerProvidedKey doesn't even has Encryption property, https://docs.microsoft.com/en-us/dotnet/api/azure.storage.blobs.models.customerprovidedkey?view=azure-dotnet

upvoted 3 times

coffecold 1 year ago

Must be either EncryptionKey or EncryptionKeyHash. But why verify the EncryptionKey while it is given ? I would suggest the hash.

upvoted 2 times

coffecold 1 year ago

Base64 encoded string of the AES256 encryption key's SHA256 hash.

upvoted 1 times

HOTSPOT
-

You are developing a web application that uses the Microsoft Identity platform for user and resource authentication. The web application called several REST APIs.

You are implementing various authentication and authorization flows for the web application.

You need to validate the claims in the authentication token.

Which token type should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Token type
Identify users for the application by using a JWT token that contains claims.	Access / ID / Refresh / SAML
Identify the permissions granted to APIs by using a JWT token that contains claims.	Access / ID / Refresh / SAML
Provide the web application with long-term access to resources on behalf of users without requiring interaction with those users.	Access / ID / Refresh / SAML
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	Access / ID / Refresh / SAML

Correct Answer:

Answer Area

Requirement	Token type
Identify users for the application by using a JWT token that contains claims.	ID
Identify the permissions granted to APIs by using a JWT token that contains claims.	Access
Provide the web application with long-term access to resources on behalf of users without requiring interaction with those users.	Refresh
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	SAML

warchoon Highly Voted 8 months ago

looks correct
https://learn.microsoft.com/en-us/azure/active-directory/develop/active-directory-configurable-token-lifetimes#token-lifetime-policies-for-access-saml-and-id-tokens
https://learn.microsoft.com/en-us/azure/active-directory/develop/refresh-tokens#refresh-token-lifetime

upvoted 5 times

JH81 Most Recent 4 months ago

Got this on 6/28/2023 and passed with 850. Answer is correct.

upvoted 4 times

NombreFalso 8 months, 2 weeks ago

Suggested answer makes sense

upvoted 2 times

mabdo 8 months, 2 weeks ago

was on exam 02/23

upvoted 2 times

g2000 9 months, 2 weeks ago

correct
https://learn.microsoft.com/en-us/azure/active-directory/develop/access-tokens

upvoted 2 times

Question #54 *Topic 4*

HOTSPOT
-

You are developing a content management application for technical manuals. The application is deployed as an Azure Static Web app.

Authenticated users can view pages under/manuals but only contributors can access the page /manuals/new.html.

You need to configure the routing for the web app.

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

NOTE: Each correct selection is worth one point.

Answer Area

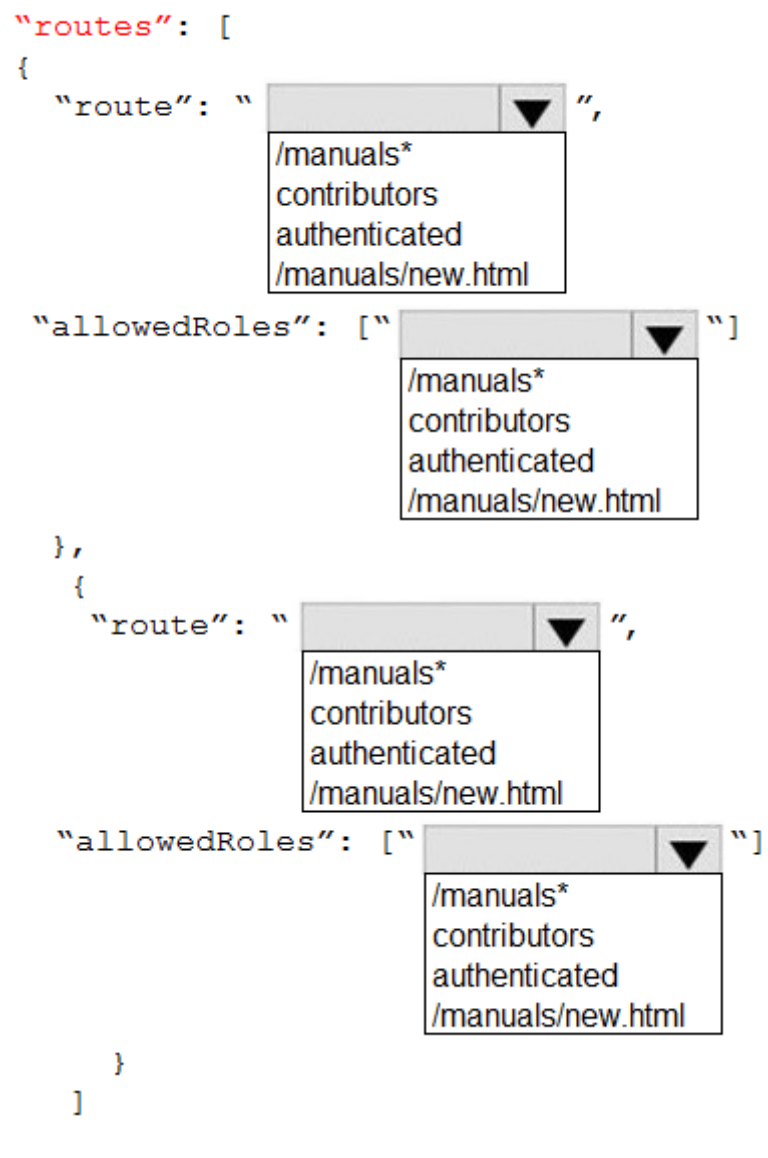

Correct Answer:

Answer Area

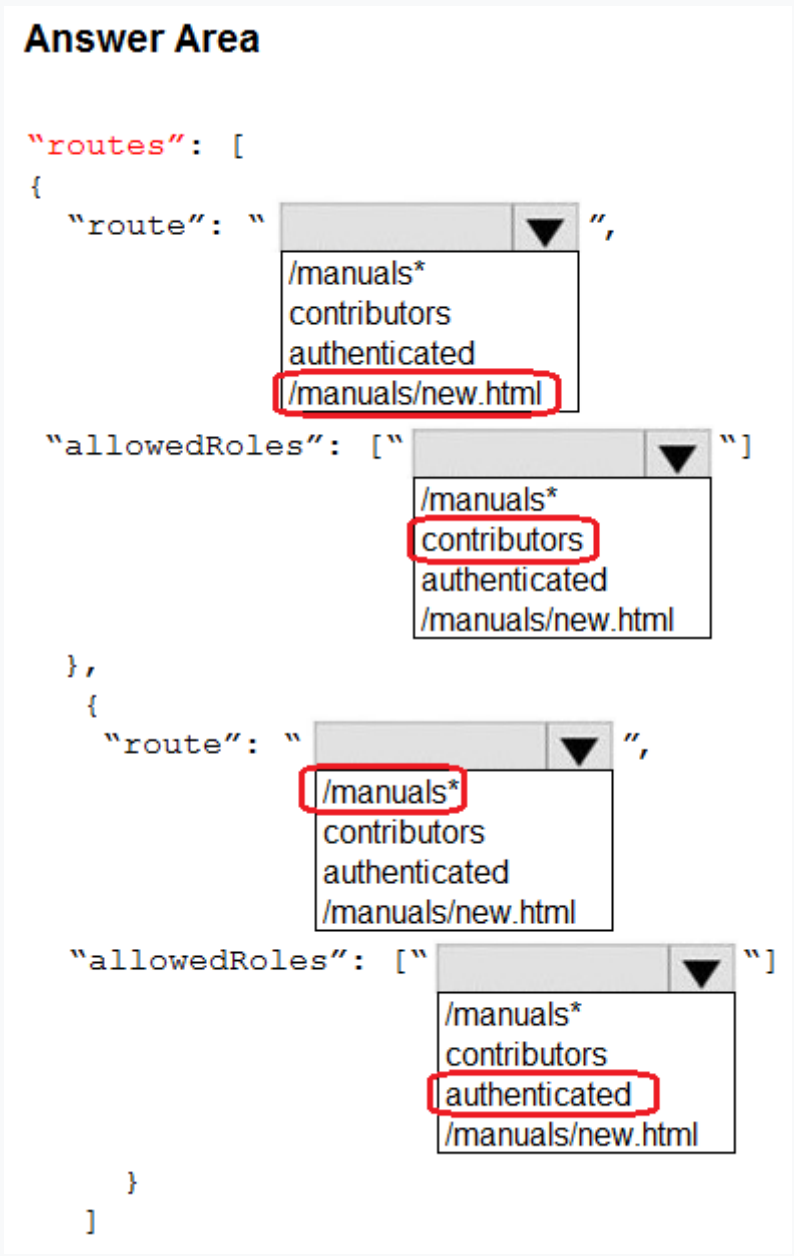

2fa Highly Voted 9 months, 3 weeks ago

Correct.
"Rule evaluation stops at the first match. A match occurs when the route property and a value in the methods array (if specified) match the request. Each request can match at most one rule." So the first rule must be for the /manuals/new.html page.

https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#routes

upvoted 22 times

unraval Most Recent 1 month, 2 weeks ago

Rule evaluation stops at the first match. If authenticated came first all users including contributors would have access to /manuals/new.html page which goes against the requirement

upvoted 1 times

rockstar89 3 months ago

got it exam 31/07/2023. correct

upvoted 1 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 4 times

plusJoyed 5 months, 2 weeks ago

couldn't it just have easily have been manuals/authenticated then html/contributers?

upvoted 4 times

Question #55 *Topic 4*

You are developing a web application that uses the Microsoft identity platform for user and resource authentication. The web application calls several REST APIs.

A REST API call must read the user's calendar. The web application requires permission to send an email as the user.

You need to authorize the web application and the API.

Which parameter should you use?

A. tenant

B. code_challenge

C. state

D. client_id

E. scope

Correct Answer: *E*

Community vote distribution

E (100%)

alexein74 Highly Voted 9 months, 1 week ago

Selected Answer: E

E. scope

Scope is the parameter used to specify the permissions that the web application is requesting to access the user's resources, such as the calendar and the ability to send an email as the user. The scope parameter is typically included in the API call or in the authorization request sent to the Microsoft identity platform.

upvoted 9 times

g2000 Highly Voted 9 months, 2 weeks ago

right
https://learn.microsoft.com/en-us/azure/active-directory/develop/scopes-oidc

upvoted 8 times

NightshadeRC Most Recent 3 months, 1 week ago

Had this question in today's exam: 2023-07-26

upvoted 3 times

Question #56 *Topic 4*

HOTSPOT
-

You develop and deploy a web app to Azure App service. The web app allows users to authenticate by using social identity providers through the Azure B2C service. All user profile information is stored in Azure B2C.

You must update the web app to display common user properties from Azure B2C to include the following information:

- Email address
- Job title
- First name
- Last name
- Office location

You need to implement the user properties in the web app.

Which code library and API should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

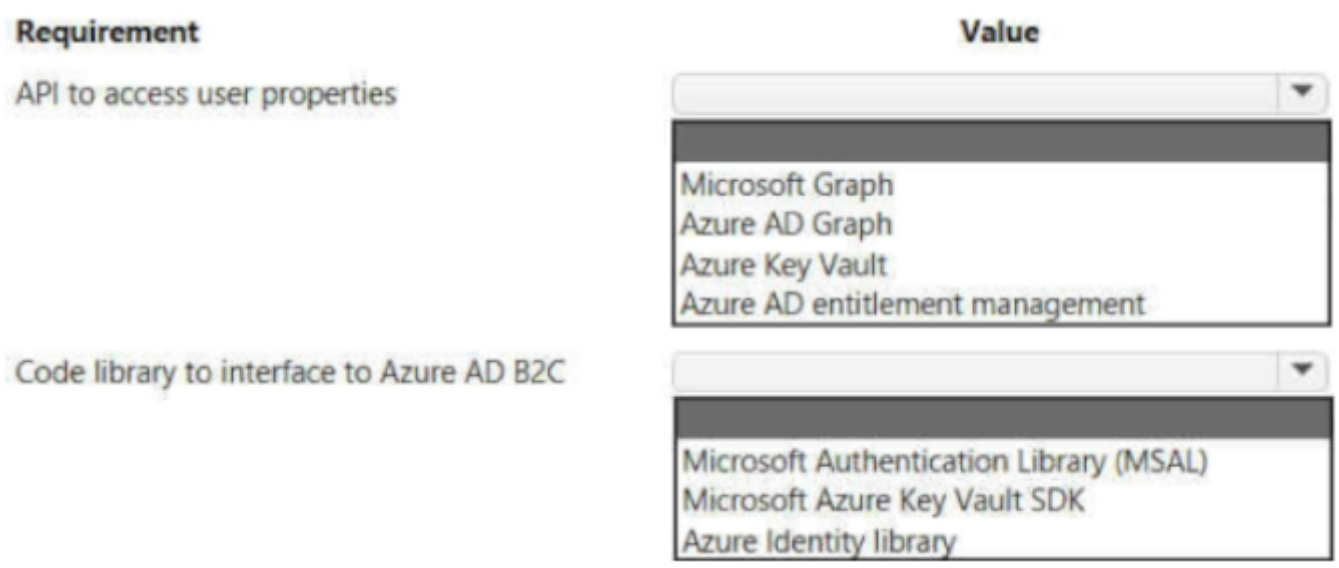

Correct Answer:

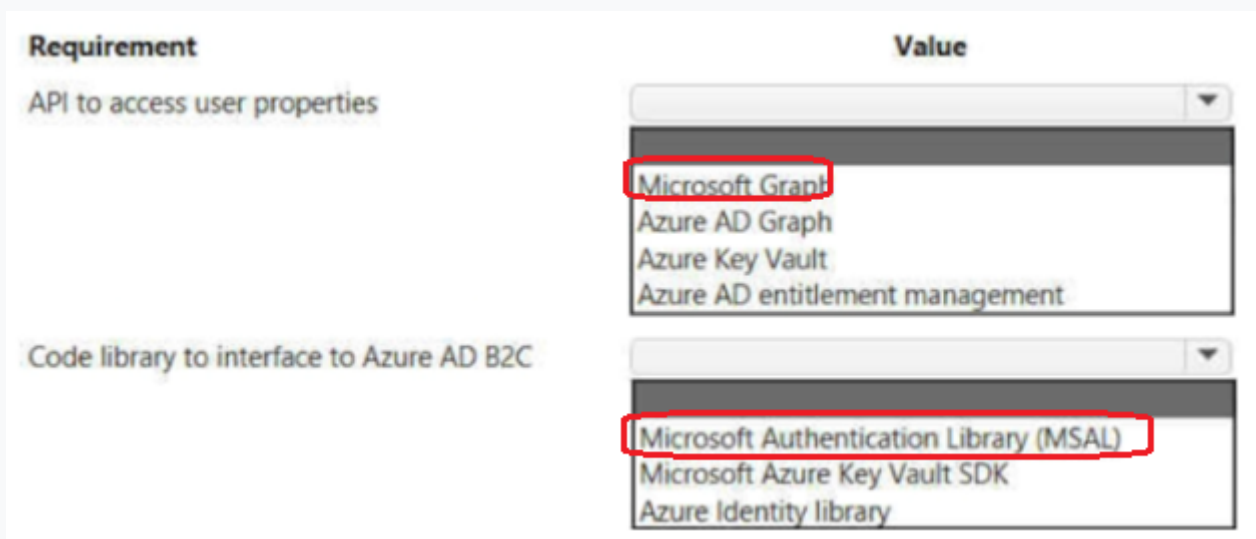

g2000 Highly Voted 9 months, 2 weeks ago

right..
https://learn.microsoft.com/en-us/azure/active-directory-b2c/microsoft-graph-operations
https://learn.microsoft.com/en-us/azure/active-directory/develop/msal-net-aad-b2c-considerations

upvoted 9 times

rockstar89 Most Recent 3 months ago

Received on 31/07/2023

upvoted 1 times

JH81 4 months ago

Got this on 6/28/2023 and passed with 850. Graph, MASL is correct.

upvoted 4 times

uncledana 4 months, 1 week ago

Correct
- Microsoft Graph
- MSAL

upvoted 1 times

Vmwarevirtual 5 months ago

Received on 27 May 2023
- Microsoft Graph
- MSAL

upvoted 2 times

HOTSPOT

-

You develop and deploy the following staticwebapp.config.json file to the app_location value specified in the workflow file of an Azure Static Web app:

```
{
  "routes": [
    {
      "route": "/api/*",
      "methods": ["GET"],
      "allowedRoles": ["registeredusers"]
    },
    {
      "route": "/api/*",
      "methods": ["PUT", "POST", "PATCH", "DELETE"],
      "allowedRoles": ["administrator"]
    },
    {
      "route": "/login",
      "rewrite": "/.auth/login/github"
    },
    {
      "route": "/.auth/login/twitter",
      "statusCode": 404
    },
    {
      "route": "/logout",
      "redirect": "/.auth/logout"
    }
  ],
  "navigationFallback": {
    "rewrite": "index.html",
    "exclude": ["/images/*.{png,jpg,gif}", "/css/*"]
  },
  "responseOverrides": {
    "400": {
      "rewrite": "/invalid-invitation-error.html"
    },
    "401": {
      "redirect": "/.auth/login/aad",
      "statusCode": 302
    },
    "403": {
      "rewrite": "/forbidden.html"
    },
    "404": {
      "rewrite": "/404.html"
    }
  },
  "mimeTypes": {
    ".json": "text/json"
  }
}
```

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Unauthenticated users are challenged to authenticate with GitHub.	○	○
A non-existent file in the `/images/` folder will generate a 404 response code.	○	○
HTTP GET method requests from authenticated users in the role named **registeredusers** are sent to the API folder.	○	○
Authenticated users that are not in the role named **registeredusers** and unauthenticated users are served a 401 HTTP error when accessing the API folder.	○	○

Correct Answer:

Answer Area

Statements	Yes	No
Unauthenticated users are challenged to authenticate with GitHub.	○	●
A non-existent file in the `/images/` folder will generate a 404 response code.	○	●
HTTP GET method requests from authenticated users in the role named **registeredusers** are sent to the API folder.	●	○
Authenticated users that are not in the role named **registeredusers** and unauthenticated users are served a 401 HTTP error when accessing the API folder.	○	●

mabdo Highly Voted 8 months, 2 weeks ago

1) Y. login with github
2) Y. true because it is excluded from returning index.html in navigation fallback, so the other option is a 404 not found
3) Y. straight forward from first routes object
4) N. from https://learn.microsoft.com/en-us/azure/static-web-apps/configuration
401 is for requesting when unauthenticated but here are you are authenticated and so the error should be 403, because you are authenticated but don't have authorization (due to not having the registeredusers role)

upvoted 13 times

aboezebal 8 months ago

1 should be no.
Unauthenticated users get a 401 - Unauthorized response and are redirected by the responseOverrides to "/.auth/login/aad". I believe that is not Github. You will only be redirected to GitHub if the an unauthenticated user routes to "/login".

upvoted 17 times

oCaio 2 months, 2 weeks ago

Initially I thought YNYN (since only 3 types of images are excluded in /images/) because:

https://learn.microsoft.com/en-us/azure/static-web-apps/configuration
"In the following example, requests for certain routes in the /images folder and all files in the /css folder are excluded from returning the fallback file."

JSON
{
"navigationFallback": {
"rewrite": "/index.html",
"exclude": ["/images/*.{png,jpg,gif}", "/css/*"]
}
}

However I changed to YYYN, because of
https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#example-configuration-file

So this depends on which Microsoft article they are trying to test us.

upvoted 1 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with N,Y,Y,N , score 926
i have doubts about second one - I don't now that status code will persist or not. We set new page 404.html without status code like in example with 401 where status code id set to 302. Needs to be verified, good luck

upvoted 8 times

adilkhan 6 months, 2 weeks ago

did you get most questions from exam topics?

upvoted 1 times

drkpnsh Highly Voted 2 months, 2 weeks ago

It should be:
1. N
In the example written here: https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#example-configuration-file
- If you go to the login page, you redirect to Github, just like in the question.
- If the request got a 401 error (Unauthorized), it redirects the user to the login route (with 302 response), which then redirects to the GitHub login.

But in this question, if you check what happens when the request got a 401 error (Unauthorized), it redirects to the AAD route (with 302 response), not to Github. So this part of the question and the example are different.

2. Y
That part of the question seems the same as the example above.
A non-existent file in the /images/ folder -> A 404 error.

3. Y
That part of the question seems the same as the example above.
GET requests from authenticated users in the registeredusers role are sent to the API. Authenticated users not in the registeredusers role and unauthenticated users are served a 401 error.

4. N
It overrides the 401 error to the 302 and redirects the user to the AAD URL.

upvoted 6 times

JH81 Most Recent 4 months ago

Got this on 6/28/2023 and passed with 850. I went with Y,N,Y,N. No idea if that is correct.

upvoted 1 times

BikashKatwal 4 months, 2 weeks ago

https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#example-configuration-file
I think based on the above config
1. Y - Unauthenticated users are challenged to authenticate with GitHub.
2. Y: An non-existent file in the /images/ folder results in a 404 error
3. Y: GET requests from authenticated users in the registeredusers role are sent to the API.
4. N: Authenticated users not in the administrator role and unauthenticated users are served a 401 error.

upvoted 1 times

nekkilodeon 6 months, 1 week ago

N,Y,Y,N
https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#example-configuration-file

upvoted 5 times

AndySmith 5 days, 1 hour ago

Almost identical example from Microsoft with good explanation. No doubts about N Y Y N.

upvoted 1 times

adilkhan 6 months, 2 weeks ago

N,Y,Y,N IS CORRECT

upvoted 3 times

DP_Bhatt 6 months, 4 weeks ago

No - Unauthorized users receive a 401 response code, which subsequently redirects to "/.auth/login/aad" with a status code 302 (see responseOverrides). You will only be redirected to GitHub if the unauthenticated user routes to "/login".

Yes - The files in the path "/images/*" with extensions png, jpg and gif are excluded from the fallback which means if the files with the mentioned extensions do not exist in the path, an 404 will be returned. The index.html fallback will be returned only if the file has an extension other than the excluded extensions. For example, if the file with name /images/xyz.svg will return index.html because the .svg extension is not excluded. But if the file is /images/unknown.jpg and the file is not exist on the path will return 404.

Yes -

No - In this case, if the user does not have registeredusers role, the 403 status code will be returned for the authenticated users.

upvoted 3 times

DP_Bhatt 6 months, 4 weeks ago

No - Unauthorized users receive a 401 response code, which subsequently redirects to "/.auth/login/aad" with a status code 302 (see responseOverrides). You will only be redirected to GitHub if the unauthenticated user routes to "/login".Yes - The files in the path "/images/*" with extensions png, jpg and gif are excluded from the fallback which means if the files with the mentioned extensions do not exist in the path, an 404 will be returned. The index.html fallback will be returned only if the file has an extension other than the excluded extensions. For example, if the file with name /images/xyz.svg will return index.html because the .svg extension is not excluded. But if the file is /images/unknown.jpg and the file is not exist on the path will return 404.
Yes -
No - In this case, if the user does not have registeredusers role, the 403 status code will be returned for the authenticated users.

upvoted 2 times

Kluk_Kluk 7 months, 2 weeks ago

I saw the same configuration in this link. All answers are Yes.
https://learn.microsoft.com/en-us/azure/static-web-apps/configuration

upvoted 1 times

courtguard 7 months, 1 week ago

the last one is No, as per your link.
If they are not in role registeredusers they are served 403 error, and if they are unauthenticated they are redirected to /login

upvoted 1 times

erieri 7 months, 3 weeks ago

Haha, really different opinions about this one... really tricky, but I will giv it a try:
1) N . Unauthticated = 401 wich is redirected to .../login/aad (with 302 status). See aboezebal's comment.
2) Y. See rpatsev's comment below. Always returns 404 even though some are rewritten with index.html
3) Y.
4) N. Authenticated users but not in correct role = 403. Unautheticated user, see question.

upvoted 1 times

lixd 7 months, 3 weeks ago

Not sure 100% the answers, but all the cases are described and covered in https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#example-configuration-file
So, the answers with comments from the web-site
1. (Y) /login : Unauthenticated users are challenged to authenticate with GitHub.
2. (Y) An non-existent file in the /images/ folder : A 404 error.
3. (Y) /api/admin : GET requests from authenticated users in the registeredusers role are sent to the API.
4. (Y) /api/admin : POST, PUT, PATCH, and DELETE requests from authenticated users in the administrator role are sent to the API. Authenticated users not in the administrator role and unauthenticated users are served a 401 error.

upvoted 1 times

adilkhan 8 months, 1 week ago

Y Y Y N is the correct answer

upvoted 2 times

baroo1 8 months, 2 weeks ago

The option 2 is tricky.

If file does not exist in /images/. Considering below files does not exist:
- /images/Image.png > /images/*.png > Excluded > 404.
- /images/Image.csv > /images/*.csv > Included > Redirect to indext.html.

As I can easily find at least 1 scenrio where it does not return a 404 > Answer is NO!

upvoted 2 times

rpatsev 8 months, 1 week ago

- /images/Image.csv > /images/*.csv > Included > Rewrite (not Redirect) to indext.html.

"navigationFallback" mode is "rewrite" (not "redirect") it means the body will be rewritten, but the statusCode will still be 404.

upvoted 3 times

mabdo 8 months, 2 weeks ago

was on exam 02/23

upvoted 1 times

sbubettino 8 months, 1 week ago

can u tell us which answers are right?

upvoted 1 times

Osip21 8 months, 3 weeks ago

Y, N, Y, Y

upvoted 2 times

TaoLu 9 months ago

Y
N
N
Y

upvoted 1 times

eugene1234 8 months, 2 weeks ago

1. there is no redirect to /login route in response overrides, 4. 401 is redirected to aad, so no 302 response code; IMHO the answer to this specifice JSON is all NO

upvoted 1 times

whiteblack 9 months, 3 weeks ago

Y,Y,Y,Y
https://learn.microsoft.com/en-us/azure/static-web-apps/configuration

upvoted 3 times

JustHereToLearn 9 months, 1 week ago

I do not understand how the last option is Y. User without role "registeredusers" role will get 401 error unauthorized. 401 will redirect user to "/.auth/login/aad" with status code 302. So how the answer is yes? I would go with no due to this reason.

upvoted 6 times

Question #58 *Topic 4*

You develop and deploy an Azure App Service web app named App1. You create a new Azure Key Vault named Vault1. You import several API keys, passwords, certificates, and cryptographic keys into Vault1.

You need to grant App1 access to Vault1 and automatically rotate credentials. Credentials must not be stored in code.

What should you do?

A. Enable App Service authentication for Appl. Assign a custom RBAC role to Vault1.

B. Add a TLS/SSL binding to App1.

C. Upload a self-signed client certificate to Vault1. Update App1 to use the client certificate.

D. Assign a managed identity to App1.

Correct Answer: *D*

Community vote distribution

D (100%)

alexein74 Highly Voted 9 months, 1 week ago

Selected Answer: D

D. Assign a managed identity to App1.

To grant App1 access to Vault1 and automatically rotate credentials without storing them in code, you should assign a managed identity to App1. Managed identities for Azure resources enable Azure services to authenticate to other Azure resources without needing to manage the authentication details.
After you enable a managed identity for App1, you can grant the identity access to Vault1 and use Azure Key Vault's built-in rotation feature to automatically rotate the credentials.

Additionally, you can use Azure Key Vault's built-in rotation feature to automatically rotate the credentials.

upvoted 15 times

Ciupaz 1 month ago

In this case ChatGPT got the right answer.

upvoted 1 times

Iyaq 1 month ago

Not a reliable source. Not accurate 100%.

upvoted 1 times

warchoon 8 months ago

Don't put ChatGPT shit here, please

upvoted 8 times

adilkhan Highly Voted 9 months, 1 week ago

D is correct https://learn.microsoft.com/en-us/azure/active-directory/fundamentals/service-accounts-managed-identities

upvoted 9 times

PKUSYD Most Recent 1 month, 2 weeks ago

D is correct

upvoted 1 times

Question #59 *Topic 4*

You are developing a Java application to be deployed in Azure. The application stores sensitive data in Azure Cosmos DB.

You need to configure Always Encrypted to encrypt the sensitive data inside the application.

What should you do first?

A. Create a new container to include an encryption policy with the JSON properties to be encrypted.

B. Create a customer-managed key (CMK) and store the key in a new Azure Key Vault instance.

C. Create a data encryption key (DEK) by using the Azure Cosmos DB SDK and store the key in Azure Cosmos DB.

D. Create an Azure AD managed identity and assign the identity to a new Azure Key Vault instance.

Correct Answer: *B*

Community vote distribution

B (100%)

alexein74 Highly Voted 9 months, 1 week ago

Selected Answer: B

B. Create a customer-managed key (CMK) and store the key in a new Azure Key Vault instance.

In order to configure Always Encrypted for the Java application, you need to first create a customer-managed key (CMK) using Azure Key Vault. This key will be used to encrypt and decrypt the sensitive data stored in Cosmos DB.
After creating the key, you should store it in a new Azure Key Vault instance, which will be used to manage and secure the key.
Once you have the key stored in Key vault, you can use Azure Cosmos DB SDK to encrypt the sensitive data and store it in Cosmos DB.

It's important to note that Always Encrypted in Azure Cosmos DB is a client-side encryption feature that encrypts and decrypts sensitive data inside the application and not in the Azure Cosmos DB service.

upvoted 8 times

warchoon 8 months ago

Don't post ChatGPT answers here. They are not always correct, even if they look logical.

upvoted 5 times

warchoon 8 months ago

CMK does not encrypt the data. It wraps the DEK. But CMK is produced first.
So the answer is correct, but the explanation is usual ChatGPT shit.

upvoted 1 times

Samueleghagha Highly Voted 9 months, 3 weeks ago

Selected Answer: B

https://learn.microsoft.com/en-us/azure/cosmos-db/how-to-always-encrypted?tabs=dotnet

upvoted 8 times

warchoon 8 months ago

=> https://learn.microsoft.com/en-us/azure/cosmos-db/how-to-always-encrypted?tabs=dotnet#setup-azure-key-vault

upvoted 2 times

dom271219 Most Recent 4 weeks, 1 day ago

A of course. Why do you talk about Azure Key Vault ? The statement of the question does not address that.

upvoted 1 times

JH81 4 months ago

Selected Answer: B

Got this on 6/28/2023 and passed with 850. Answer is correct.

upvoted 2 times

Tom111 7 months, 4 weeks ago

As the question asks what to do first I would select answer A.
"This policy must be provided when the container is created and it is immutable. In the current release, you can't update the encryption policy."
=> https://learn.microsoft.com/en-us/azure/cosmos-db/how-to-always-encrypted?tabs=dotnet#encryption-policy

upvoted 1 times

Question #60 *Topic 4*

HOTSPOT
-

You develop a web app that interacts with Azure Active Directory (Azure AD) groups by using Microsoft Graph.

You build a web page that shows all Azure AD groups that are not of the type 'Unified'.

You need to build the Microsoft Graph query for the page.

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

NOTE: Each correct selection is worth one point.

Answer Area

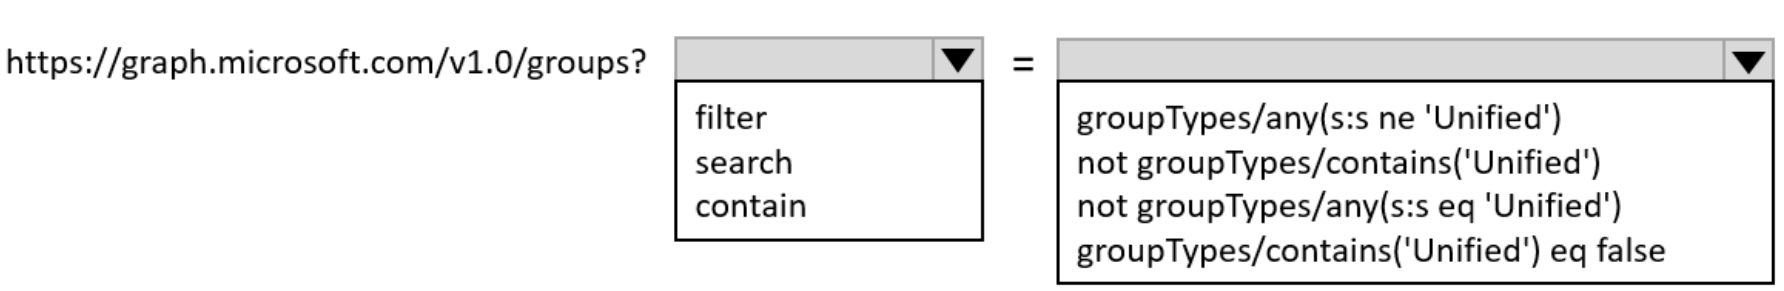

Correct Answer:

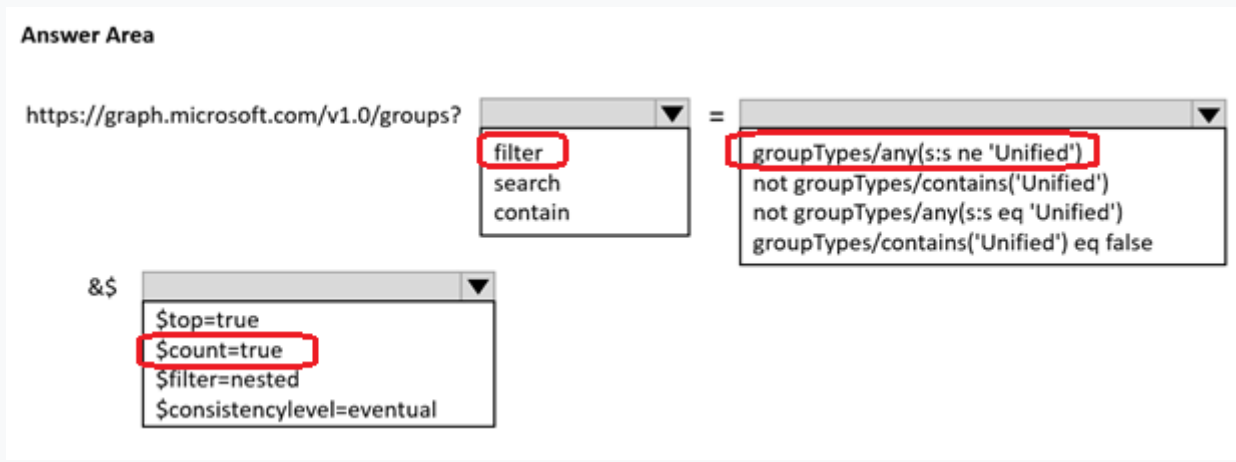

mmdex Highly Voted 8 months, 1 week ago

~/groups?$filter=NOT groupTypes/any(c:c eq 'Unified')&$count=true

This exact example is mentioned in Microsoft documentation:
https://learn.microsoft.com/en-us/graph/filter-query-parameter?tabs=http#examples-using-the-filter-query-operator

upvoted 27 times

c_mazzotta 6 months, 2 weeks ago

Agree. if you try:"$filter=groupTypes/any(s:s ne 'Unified')&$count=true", you get: "Using 'ne' inside 'any' clause is invalid. Suggestion: consider using the 'not' operator instead."

upvoted 1 times

warchoon 8 months ago

=> https://learn.microsoft.com/en-us/graph/filter-query-parameter?tabs=http#examples-using-the-filter-query-operator:~:text=~/groups%3F%24filter%3DNOT%20groupTypes/any(c%3Ac%20eq%20%27Unified%27)%26%24count%3Dtrue

upvoted 3 times

g2000 Highly Voted 9 months, 2 weeks ago

it is correct
https://learn.microsoft.com/en-us/azure/cosmos-db/how-to-always-encrypted?tabs=dotnet

upvoted 7 times

tcybu 9 months, 2 weeks ago

Agreed, it's correct, but here're the links with explanation
https://learn.microsoft.com/en-us/graph/filter-query-parameter?tabs=http#any-operator
https://learn.microsoft.com/en-us/graph/aad-advanced-queries
upvoted 2 times

Acelexout Most Recent 1 month ago
From documentation (https://learn.microsoft.com/en-us/graph/filter-query-parameter?tabs=http#examples-using-the-filter-query-operator):
List all non-Microsoft 365 groups in an organization.
GET ~/groups?$filter=NOT groupTypes/any(c:c eq 'Unified')&$count=true*
upvoted 1 times

BaoNguyen2411 3 months ago
got this question on 06/29/2023
upvoted 4 times

BaoNguyen2411 3 months ago
Got this answer on 06-29-2023
upvoted 2 times

Dianahu 3 months ago
Just try it yourself here: https://developer.microsoft.com/en-us/graph/graph-explorer
https://graph.microsoft.com/v1.0/groups?$filter=groupTypes/any(s:s ne 'Unified')&$count=true -> 400
https://graph.microsoft.com/v1.0/groups?$filter=not groupTypes/contains('Unified')&$count=true -> 400
https://graph.microsoft.com/v1.0/groups?$filter=not groupTypes/any(s:s eq 'Unified')&$count=true -> 200 with ConsistencyLevel: eventual
https://graph.microsoft.com/v1.0/groups?$filter=groupTypes/contains('Unified') eq false&$count=true -> 400
upvoted 1 times

nekkilodeon 6 months, 1 week ago
From Microsoft documentation https://learn.microsoft.com/en-us/graph/filter-query-parameter?tabs=http#examples-using-the-filter-query-operator
List all non-Microsoft 365 groups in an organization. GET ~/groups?$filter=NOT groupTypes/any(c:c eq 'Unified')&$count=true*
upvoted 1 times

Heighte 6 months, 3 weeks ago
any idea why $count=true though?
upvoted 3 times

Tomn82 7 months, 1 week ago
Answer ~/groups?$filter=NOT groupTypes/any(c:c eq 'Unified')&$count=true

Option 2: Use the $filter query parameter with the ne operator. This request is not supported by default because the ne operator is only supported in advanced queries. Therefore, you must add the ConsistencyLevel header set to eventual and use the $count=true query string.

https://learn.microsoft.com/en-us/graph/aad-advanced-queries?tabs=http
upvoted 1 times

baroo1 8 months, 2 weeks ago
Answer: /groups?$filter=not groupTypes/any(s:s eq 'Unified')&$count=true
It also requires an extra header
ConsistencyLevel: eventual

"ne" is not supported.
message": "Using 'ne' inside 'any' clause is invalid. Suggestion: consider using the 'not' operator instead. For example: $filter=not groupTypes/any(x:x eq 'Unified')
upvoted 2 times

vmakharashvili 8 months ago
I don't think so. I see if you want to use negative this way, you have to set parenthesis after not. like this: /groups$filter=not(groupTypes/any(s:s eq 'Unified'))
upvoted 1 times

Net_IT 8 months, 2 weeks ago
IMO it should be filter=NOT groupTypes/any(s:s eq 'Unified')&$count=true.
Source: https://learn.microsoft.com/en-us/azure/cosmos-db/how-to-always-encrypted?tabs=dotnet

In advanced querries you can also use the ne operator but then the consistencylevel has to be put to eventual and I'm not sure that's the case here?
upvoted 1 times

eugene1234 8 months, 2 weeks ago
The answer is wrong! 1. filter should be $filter, and the expression should be NOT groupTypes/any(s:s ne 'Unified'), we want records that are NOT unified, see examples at https://learn.microsoft.com/en-us/graph/filter-query-parameter?tabs=http
upvoted 2 times

Question #61 *Topic 4*

DRAG DROP

-

You are developing an Azure solution.

You need to develop code to access a secret stored in Azure Key Vault.

How should you complete the code segment? To answer, drag the appropriate code segments to the correct location. Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Code segments

- `DefaultAzureCredential`
- `ClientSecretCredential`
- `CloudClients`
- `SecretClient`

Answer area

```
def get_secret():
    var1 = os.environ.get("KEY_VAULT_URI")
    var2 = [          ] ( var1 , [          ]())
    var3 = var2.get_secret("secret" )
    return "secret value {}" .format(var3.value)
```

Correct Answer:

Answer area

```
def get_secret():
    var1 = os.environ.get("KEY_VAULT_URI")
    var2 = SecretClient ( var1 , DefaultAzureCredential ())
    var3 = var2.get_secret("secret" )
    return "secret value {}" .format(var3.value)
```

g2000 Highly Voted 9 months, 2 weeks ago

it's correct
https://learn.microsoft.com/en-us/azure/cosmos-db/how-to-always-encrypted?tabs=dotnet

upvoted 11 times

JustHereToLearn 9 months, 1 week ago

Looks good but I think this URL should be more accurate https://learn.microsoft.com/en-us/python/api/azure-keyvault-secrets/azure.keyvault.secrets.secretclient?view=azure-python#examples

upvoted 15 times

JH81 Highly Voted 4 months ago

Got this on 6/28/2023 and passed with 850. Correct answer.

upvoted 5 times

Question #62 *Topic 4*

HOTSPOT
-

You are a developer building a web site using a web app. The web site stores configuration data in Azure App Configuration.

Access to Azure App Configuration has been configured to use the identity of the web app for authentication. Security requirements specify that no other authentication systems must be used.

You need to load configuration data from Azure App Configuration.

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

NOTE: Each correct selection is worth one point.

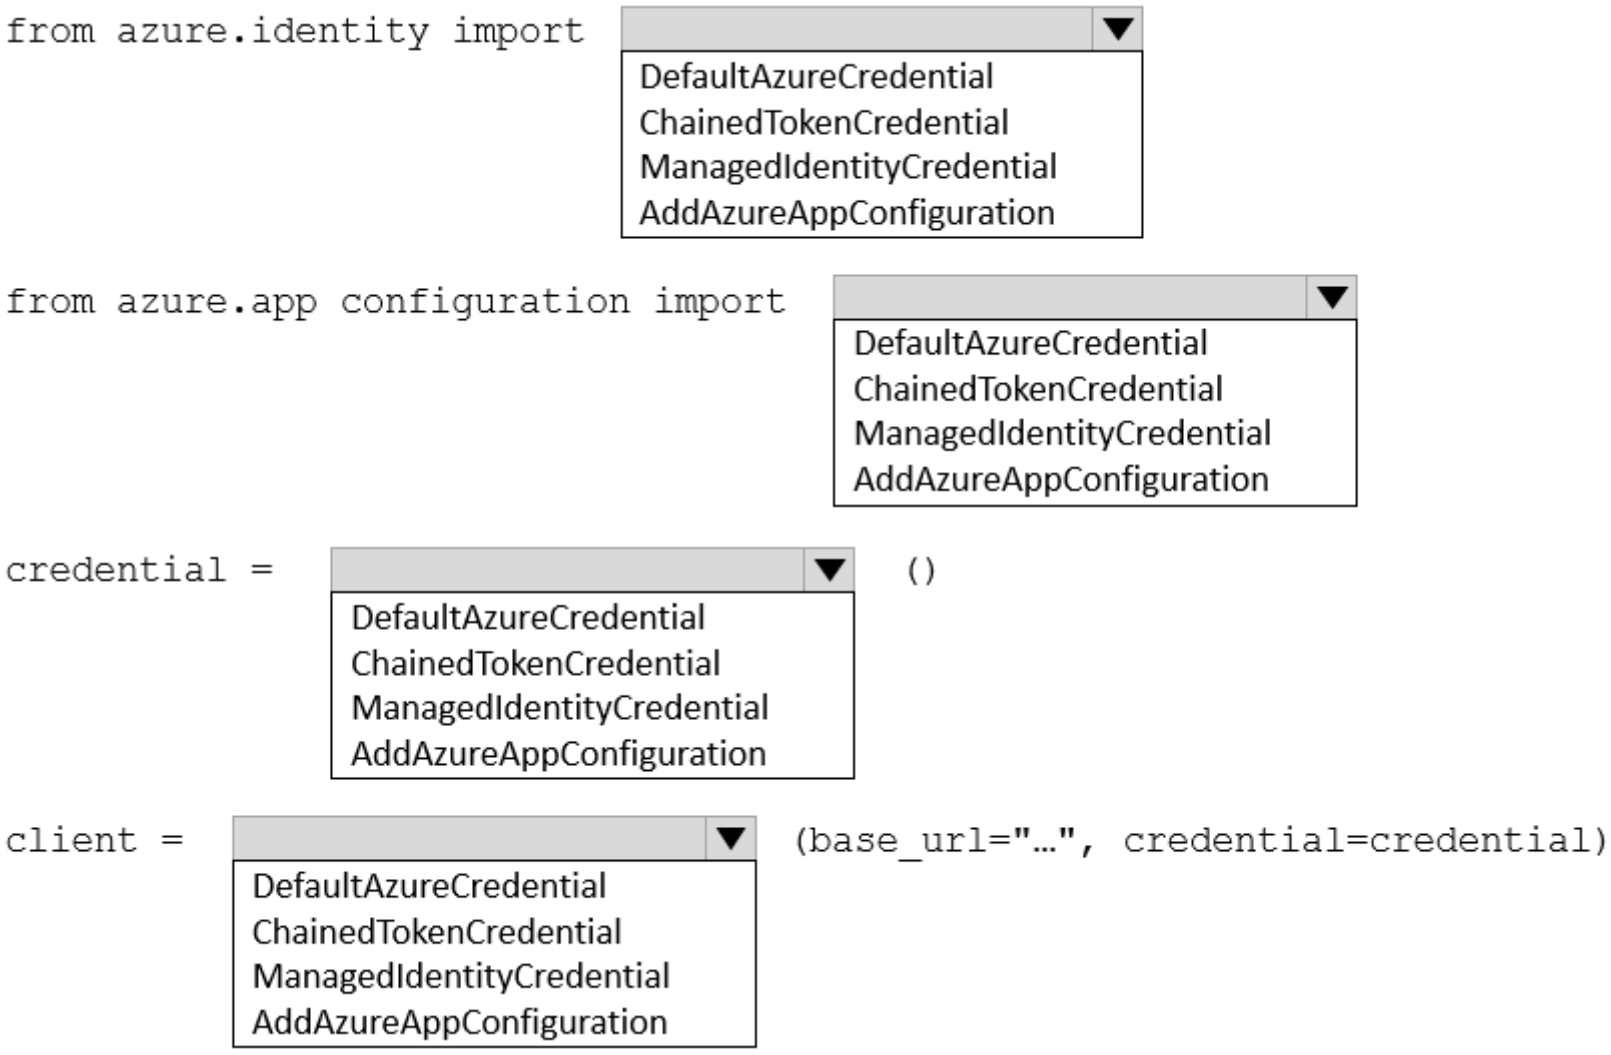

Correct Answer:

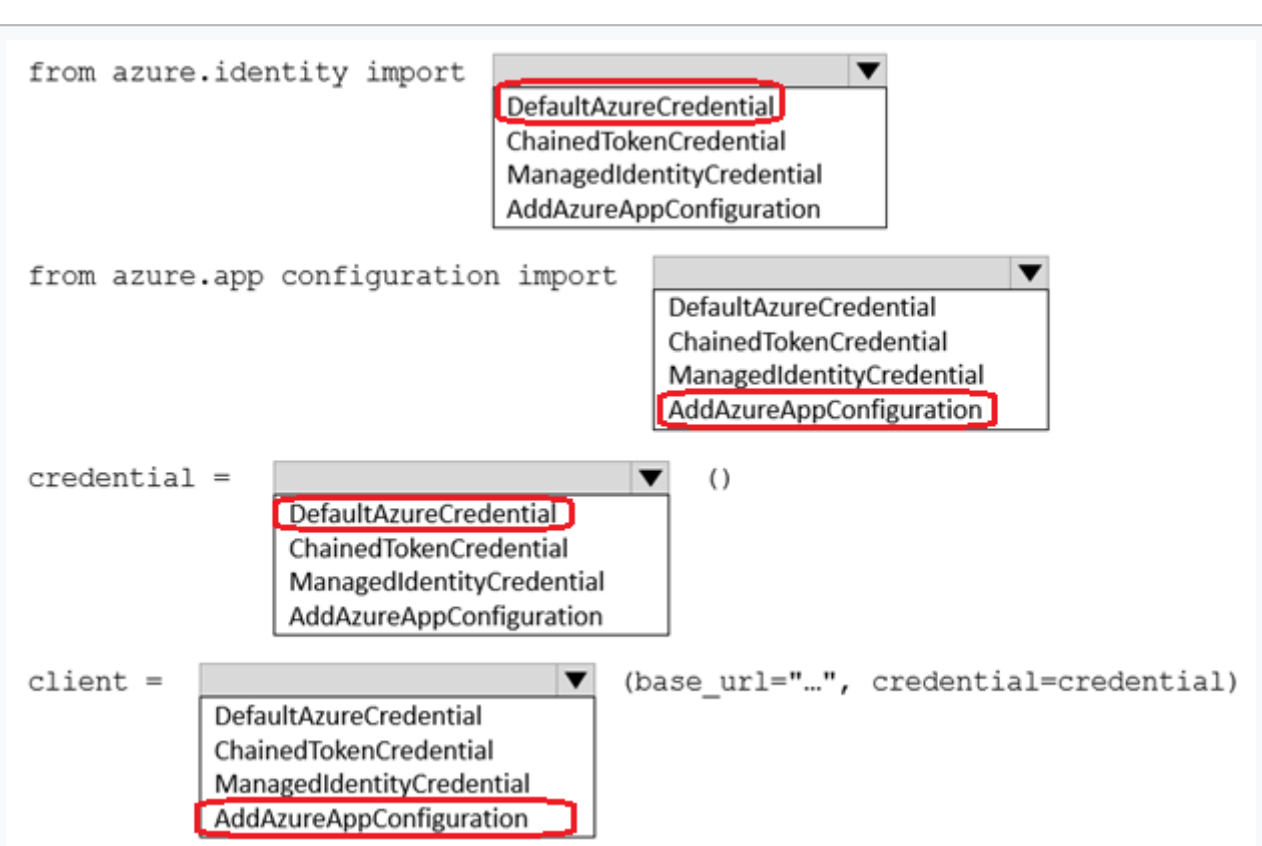

Taschiloge Highly Voted 9 months, 2 weeks ago

I think it should be ManagedIdentityCredential:
ManagedIdentityCredential: attempts to authenticate using a managed identity that is assigned to the deployment environment (if any).

DefaultAzureCredential can not be, since only one type of authentication is allowed.
ChainedTokenCredential could work if there is only one authentication type is specified in code.
Source: https://yourazurecoach.com/2020/08/13/managed-identity-simplified-with-the-new-azure-net-sdks/
upvoted 14 times

adilkhan 9 months, 1 week ago
from azure.identity import DefaultAzureCredential
from azure.appconfiguration import AzureAppConfigurationClient

credential = DefaultAzureCredential()

client = AzureAppConfigurationClient(base_url="your_endpoint_url", credential=credential)
upvoted 5 times

SharpZx 7 months, 2 weeks ago
I think it should be DefaultAzureCredential. Check DefaultAzureCredential section at
https://learn.microsoft.com/en-us/python/api/overview/azure/identity-readme?view=azure-python

DefaultAzureCredential is appropriate for most applications which will run in the Azure Cloud because it combines common production credentials with development credentials

Managed Identity - If the application is deployed to an Azure host with Managed Identity enabled, DefaultAzureCredential will authenticate with it.
upvoted 2 times

JamieS 7 months, 2 weeks ago
I disagree, if they used options (https://docs.microsoft.com/en-us/dotnet/api/azure.identity.defaultazurecredentialoptions?view=azure-dotnet-preview) they could exclude unwanted auth methods, but it's not in the code. Otherwise "It is not always clear what authentication method got executed".

Question states that "configured to use the identity" so I think it's ManagedIdentityCredential in place of azure default.
upvoted 3 times

SharpZx 7 months, 1 week ago
JamieS you are right. It should be ManagedIdentityCredential. I overlooked .. "Security requirements specify that no other authentication systems must be used".
upvoted 3 times

tarek0811 3 months ago
chatgpt generated code with ManagedIdentityCredential
upvoted 1 times

EricPerezVillar Highly Voted 5 months, 3 weeks ago
Answer is correct. It is here:
https://learn.microsoft.com/en-us/python/api/overview/azure/appconfiguration-readme?view=azure-python#create-a-client
upvoted 9 times

MikeAWS 2 weeks ago
great catch! thanks
upvoted 1 times

Ciupaz Most Recent 1 month ago
This question is not updated with the latest Microsoft docs.
The 2nd and 4th answers are not correct, and are not present in DDL.
Here the right Python code:

from azure.identity import DefaultAzureCredential
from azure.appconfiguration import AzureAppConfigurationClient

credential = DefaultAzureCredential()

client = AzureAppConfigurationClient(base_url="your_endpoint_url", credential=credential)

https://learn.microsoft.com/en-us/python/api/overview/azure/appconfiguration-readme?view=azure-python#create-a-client
upvoted 2 times

macobuzi 2 months, 1 week ago
And why the hell did Microsoft decide to include Python code in the exam?
upvoted 2 times

EliteAllen 2 months, 3 weeks ago
To load configuration data from Azure App Configuration using the identity of the web app for authentication, you would typically use the ManagedIdentityCredential class for authentication. Here's how you should complete the code:

from azure.identity import ManagedIdentityCredential

from azure.appconfiguration import AzureAppConfigurationClient

credential = ManagedIdentityCredential()

client = AzureAppConfigurationClient(base_url="your_endpoint_url", credential=credential)

upvoted 4 times

nickk6425 4 months ago

Should be ManagedIdentity because "...configured to use the identity of the web app for authentication. Security requirements specify that no other authentication systems must be used."

upvoted 2 times

nickk6425 4 months ago

DefaultAzureCredentials will try different authentication, ie. it will use AZ, Visual Studio, ManagedIdentity, and so on until one succeeds

upvoted 1 times

Dats1987 6 months, 2 weeks ago

Note that the DefaultAzureCredential is used to authenticate using the identity of the web app. This credential provider tries multiple authentication methods (e.g. environment variables, Azure Managed Identity, Azure CLI) until it finds a suitable one. In this way, you can ensure that no other authentication systems are used
Above ans are correct.

upvoted 1 times

Dilmurod 7 months ago

It should be ManagedIdentityCredential.

Managed identity authentication is supported via either the DefaultAzureCredential or the ManagedIdentityCredential directly for the following Azure services:
Example:
from azure.identity import ManagedIdentityCredential
from azure.keyvault.secrets import SecretClient

credential = ManagedIdentityCredential()
client = SecretClient("https://my-vault.vault.azure.net", credential)

Reference: https://github.com/Azure/azure-sdk-for-python/tree/azure-appconfiguration_1.4.0/sdk/identity/azure-identity

upvoted 2 times

naivecoder786 7 months, 4 weeks ago

Correct one is ManagedIdentityCredential 110% Sure !

upvoted 3 times

[Removed] 7 months, 2 weeks ago

Can you explain please?

upvoted 1 times

karthikwarrior 8 months, 1 week ago

What is the right answer then? everyone thrown some kitchen sink.. pls keep it simple.

upvoted 4 times

MikeAWS 2 weeks ago

from azure.identity import DefaultAzureCredential
from azure.appconfiguration import AzureAppConfigurationClient

credential = DefaultAzureCredential()

client = AzureAppConfigurationClient(base_url="your_endpoint_url", credential=credential)

https://learn.microsoft.com/en-us/python/api/overview/azure/appconfiguration-readme?view=azure-python#create-a-client

upvoted 1 times

adilkhan 9 months, 1 week ago

The options are wrong https://learn.microsoft.com/en-us/python/api/overview/azure/appconfiguration-readme?view=azure-python

from azure.identity import DefaultAzureCredential
from azure.appconfiguration import AzureAppConfigurationClient

credential = DefaultAzureCredential()

client = AzureAppConfigurationClient(base_url="your_endpoint_url", credential=credential)

upvoted 3 times

ajayasa 9 months ago

@adilkhan
can you please provide right answers. your comment starts with answers being wrong and you are providing the same answers mentioned

upvoted 6 times

maiwufsiorn 8 months, 3 weeks ago

No, he is not. Read it again: AzureAppConfigurationClient is different from AddAzureAppConfiguration.
upvoted 2 times

warchoon 8 months ago
It looks like typo. AddAzureAppConfiguration is a method, not object.
upvoted 1 times

g2000 9 months, 2 weeks ago
DefaultAzureCredential is wrong because it tries different authentication methods.
https://yourazurecoach.com/2020/08/13/managed-identity-simplified-with-the-new-azure-net-sdks/
upvoted 4 times

Question #63 *Topic 4*

You are developing several microservices to deploy to a new Azure Kubernetes Service cluster. The microservices manage data stored in Azure Cosmos DB and Azure Blob storage. The data is secured by using customer-managed keys stored in Azure Key Vault.

You must automate key rotation for all Azure Key Vault keys and allow for manual key rotation. Keys must rotate every three months. Notifications of expiring keys must be sent before key expiry.

You need to configure key rotation and enable key expiry notifications.

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

NOTE: Each correct selection is worth one point.

A. Create and configure a new Azure Event Grid instance.

B. Configure Azure Key Vault alerts.

C. Create and assign an Azure Key Vault access policy.

D. Create and configure a key rotation policy during key creation.

Correct Answer: *AC*

Community vote distribution

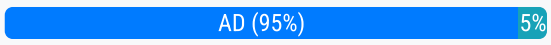

abcdx Highly Voted 6 months, 1 week ago

Selected Answer: AD

A&D for sure!
https://learn.microsoft.com/en-us/azure/key-vault/keys/how-to-configure-key-rotation

upvoted 14 times

halfway Highly Voted 6 months, 2 weeks ago

Selected Answer: AD

Key rotation policy and Event Grid notification

upvoted 6 times

macobuzi Most Recent 2 months, 1 week ago

Selected Answer: AD

We can use the Key Rotation Policy in Azure Key Vault combined with Event Grid to trigger sending notification when a secret in the key vault is about to expire.
https://learn.microsoft.com/en-us/azure/key-vault/general/event-grid-tutorial

upvoted 1 times

nekkilodeon 6 months, 1 week ago

Selected Answer: BC

B & C are correct
Alerts for notifications and Access policies for storage acess to existing keys

upvoted 1 times

paunski7 6 months, 1 week ago

B & D

To configure key rotation and enable key expiry notifications for Azure Key Vault, you should perform the following two actions:

B. Configure Azure Key Vault alerts: Configure alerts for when a key is expiring, so that you can receive notifications before the key expiry. You can configure these alerts in Azure Key Vault using Azure Monitor, which sends an email or a webhook notification to a recipient or service when the key is about to expire. You can specify the notification threshold in terms of days, so you can receive alerts, for example, seven days before the key expiry.

C. Create and assign an Azure Key Vault access policy: Create and assign an access policy for your Azure Key Vault that allows your microservices to perform key rotations manually, and automated key rotations using a key rotation script. You can create and assign access policies for Azure Key Vault through the Azure portal or the Azure CLI.

These actions enable you to configure key rotation and notifications for key expiry.

upvoted 3 times

dy0917 5 months, 2 weeks ago

Key vault alerts use to monitor health issues.

After you start to use Azure Key Vault to store your production secrets, it's important to monitor the health of your key vault to make sure that your service operates as intended.

upvoted 1 times

paunski7 6 months, 1 week ago

Options A and D are not relevant to the solution. Azure Event Grid is a service that allows you to react to events in Azure services by routing them to different endpoints, but it is not required for key rotation or notifications for Azure Key Vault. Key rotation policies are not created during key creation, but rather they are created and applied to existing keys.

upvoted 2 times

smariussorin 2 months, 3 weeks ago

https://learn.microsoft.com/en-us/azure/key-vault/general/event-grid-tutorial

check documentation:". When one of the secrets in the key vault is about to expire (defined as 30 days before expiration date), Event Grid is notified of the status change and makes an HTTP POST to the endpoint. "
is A & D

upvoted 1 times

Question #64 *Topic 4*

You are developing a web application that uses the Microsoft identity platform to authenticate users and resources. The web application calls several REST APIs.

The APIs require an access token from the Microsoft identity platform.

You need to request a token.

Which three properties should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Redirect URI/URL

B. Application ID

C. Application name

D. Application secret

E. Supported account type

Correct Answer: *ABD*

Community vote distribution

ABD (100%)

AbidooKing Highly Voted 3 months, 3 weeks ago

Answer seems correct!

https://learn.microsoft.com/en-us/graph/auth-v2-user?tabs=http#3-request-an-access-token

upvoted 9 times

nardk Most Recent 1 month, 3 weeks ago

Selected Answer: ABD

correct

upvoted 2 times

Roel1 3 months, 3 weeks ago

Seems plausible

upvoted 3 times

Steffexx 3 months, 3 weeks ago

Answer A B D is correct

upvoted 1 times

Question #65 *Topic 4*

HOTSPOT
-

You are developing an application that uses Azure Storage to store customer data. The data must only be decrypted by the customer and the customer must be provided a script to rotate keys.

You need to provide a script to rotate keys to the customer.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
$h = $(az keyvault show --hsm-name … --query "properties.hsmUri"
$x = az keyvault [ ▾ ] list-versions --name ""
                   key
                   secret
                   recover
                   certificate
  --vault-name "" --query "[0].kid"
az storage account update
  --name … \
  --resource-group … \
  --encryption-key-name … \
  --encryption-key-version $x \
  --encryption-key-source [ ▾ ]
                          Microsoft.Secret
                          Microsoft.Storage
                          Microsoft.Keyvault
                          Microsoft.Certificate
  --encryption-key-vault $h
```

Correct Answer:

Answer Area

```
$h = $(az keyvault show --hsm-name … --query "properties.hsmUri"
$x = az keyvault [key ▾] list-versions --name ""
                   key
                   secret
                   recover
                   certificate
  --vault-name "" --query "[0].kid"
az storage account update
  --name … \
  --resource-group … \
  --encryption-key-name … \
  --encryption-key-version $x \
  --encryption-key-source [ ▾ ]
                          Microsoft.Secret
                          Microsoft.Storage
                          Microsoft.Keyvault
                          Microsoft.Certificate
  --encryption-key-vault $h
```

FakerJani 1 month, 3 weeks ago

Answers is correct.
https://learn.microsoft.com/en-us/cli/azure/storage/account?view=azure-cli-latest#az-storage-account-update
https://learn.microsoft.com/en-us/cli/azure/keyvault/key?view=azure-cli-latest

upvoted 2 times

Momook 1 month, 3 weeks ago

maybe the answer should be , secret and Microsoft.Keyvault ??

upvoted 2 times

BaoNguyen2411 3 months ago

got this question on 29/06/2023

upvoted 2 times

CloudlessCloud 3 months, 2 weeks ago

Ty! (^///^)

upvoted 1 times

Question #66 *Topic 4*

You are developing several Azure API Management (APIM) hosted APIs.

You must transform the APIs to hide private backend information and obscure the technology stack used to implement the backend processing.

You need to protect all APIs.

What should you do?

A. Configure and apply a new inbound policy scoped to a product.

B. Configure and apply a new outbound policy scoped to the operation.

C. Configure and apply a new outbound policy scoped to global.

D. Configure and apply a new backend policy scoped to global.

Correct Answer: *B*

Community vote distribution

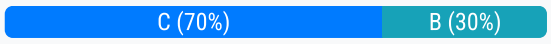

kakajos 1 month, 1 week ago

Selected Answer C:
Global scope is for All APIs in your API Management Instance. And since we need to protect all the APIs the, we want an outbound policy with a global scope. Ans

upvoted 2 times

katrang 1 month, 2 weeks ago

Selected Answer: C

Add outbound policy -> Set headers and the scope should be global (transform the APIs - no mention of specific operation)

upvoted 2 times

devex 1 month, 3 weeks ago

Selected Answer: C

C: You need to protect all APIs ("All operations")

upvoted 2 times

tempacc4nk 2 months ago

Selected Answer: C

To safeguard private backend information and mask the underlying technology stack utilized in the backend processes of your Azure API Management (APIM) hosted APIs, implement and apply a fresh global-scoped outbound policy. This policy will uniformly apply the necessary transformations and concealments to the responses of all APIs within your APIM instance, ensuring a consistent approach to securing backend details and technology stack disclosure. So C

upvoted 3 times

namlus 2 months ago

Its 'C'. Check https://learn.microsoft.com/en-us/azure/api-management/set-edit-policies?tabs=form

upvoted 4 times

lyggwtwtczxnhznebw 2 months, 2 weeks ago

It should be an outbound policy. The articles says you shoyld select 'All operations', so either you take that as 'global' and go for C, or there is a typo in B and it should have said 'All operations'

upvoted 3 times

smariussorin 2 months, 3 weeks ago

C, is a global policy. Check https://learn.microsoft.com/en-us/azure/api-management/transform-api.

"Select Demo Conference API > Design > All operations."

upvoted 2 times

Firo 3 months ago

Selected Answer: B

B is correct refer to below link
https://learn.microsoft.com/en-us/azure/api-management/transform-api

upvoted 3 times

macobuzi 2 months, 1 week ago

Why not C?

upvoted 1 times

Topic 5 - Question Set 5

Question #1 *Topic 5*

DRAG DROP -

You develop a web app that uses the tier D1 app service plan by using the Web Apps feature of Microsoft Azure App Service.

Spikes in traffic have caused increases in page load times.

You need to ensure that the web app automatically scales when CPU load is about 85 percent and minimize costs.

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

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

Select and Place:

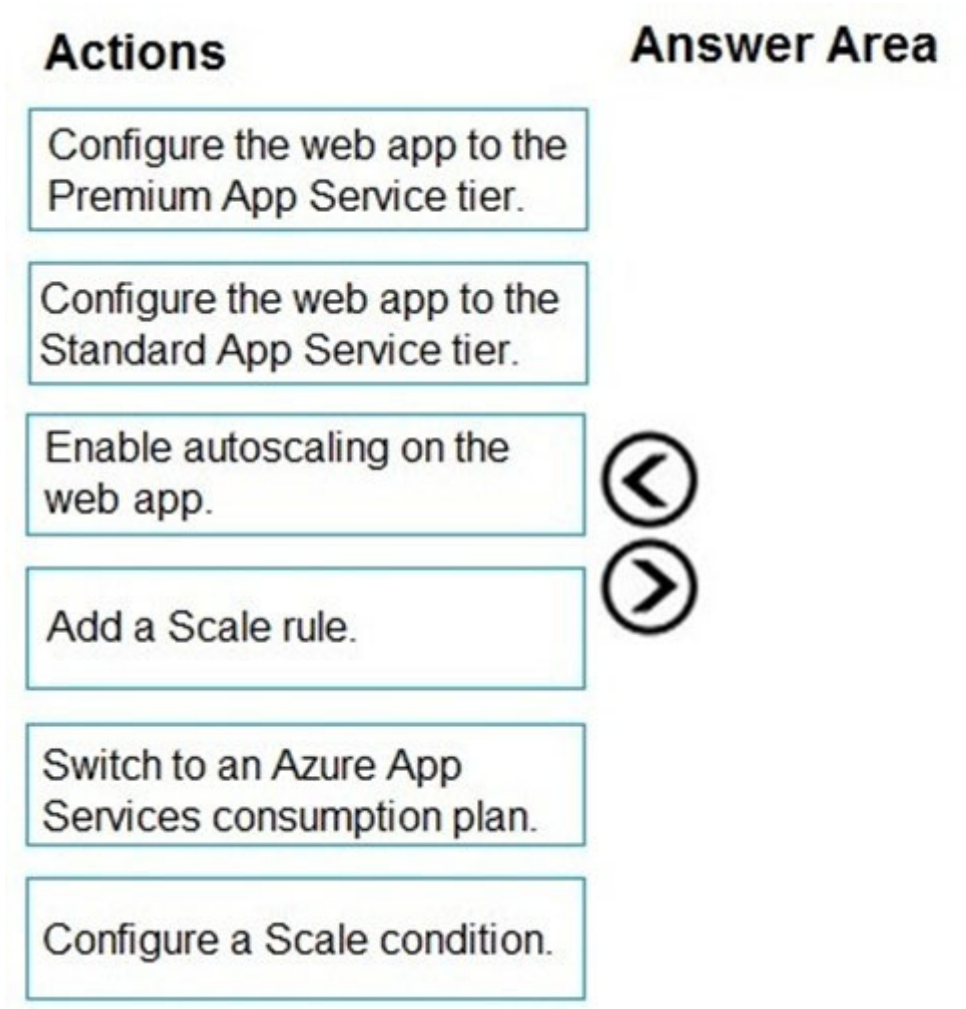

Correct Answer:

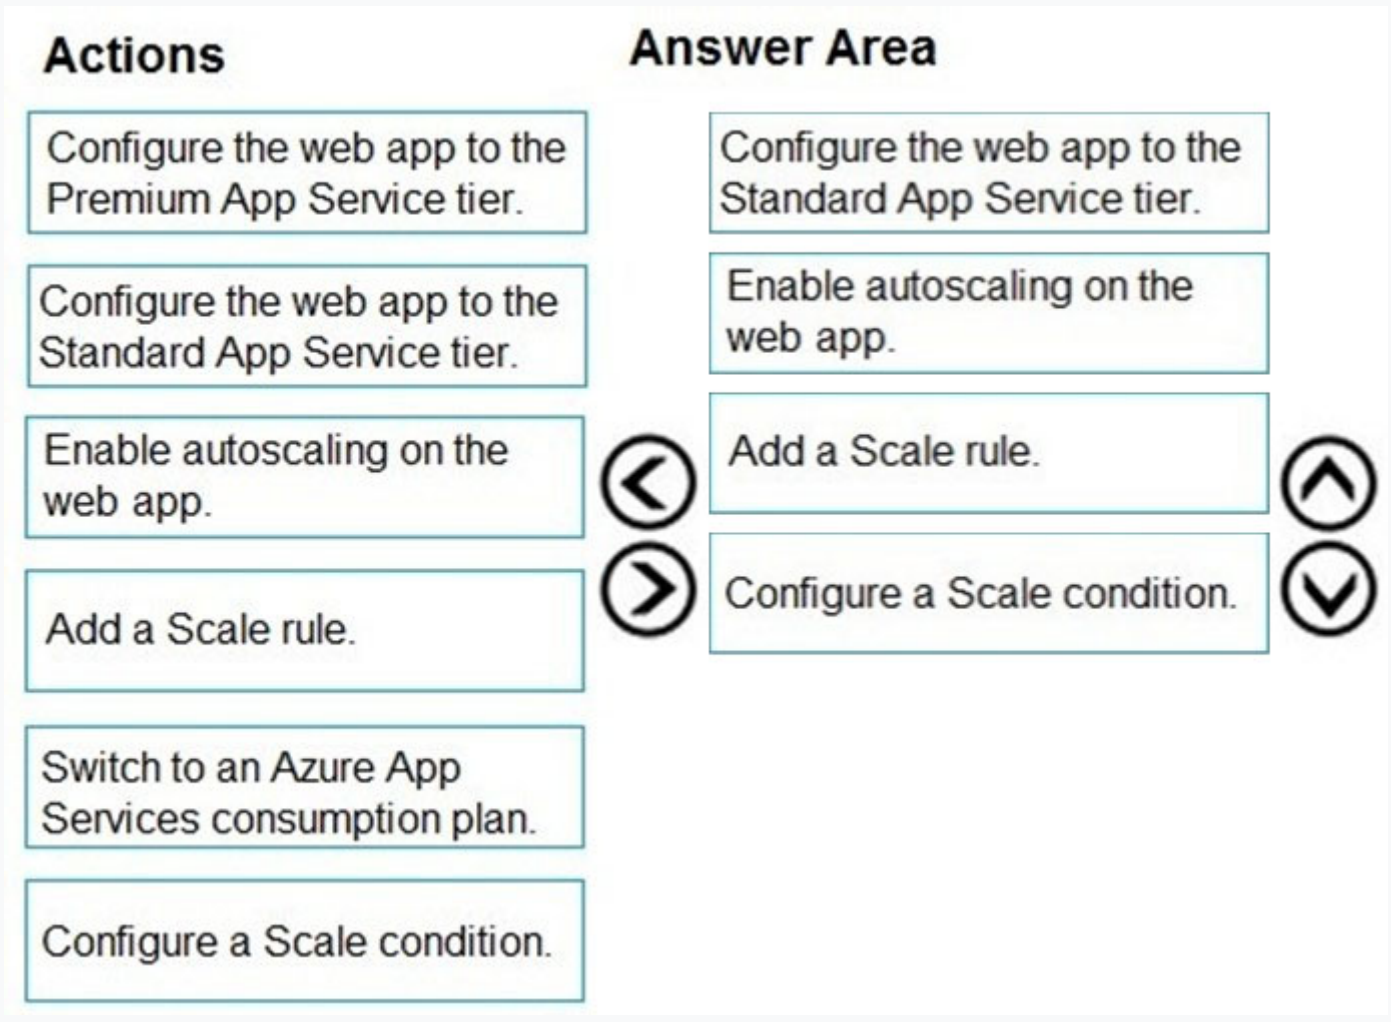

Step 1: Configure the web app to the Standard App Service Tier

The Standard tier supports auto-scaling, and we should minimize the cost.

Step 2: Enable autoscaling on the web app

First enable autoscale -

Step 3: Add a scale rule -

Step 4: Add a Scale condition -
Reference:
https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-autoscale-get-started

AmirCSUDH Highly Voted 2 years, 11 months ago
The Provided Answer is Correct
upvoted 93 times

dhishkiyaau 2 years, 2 months ago
the correct answer is :-
Configure the web app to the Standard App Service Tier
Enable autoscaling on the web app
Add a Scale condition
Add a scale rule
Please refer - https://docs.microsoft.com/en-us/learn/modules/app-service-autoscale-rules/6-autoscale-a-web-app
upvoted 42 times

altafpatel1984 1 year, 11 months ago
No. This is for Monitor auto scale. Question is for Web App auto scale, and for that steps would be as follow:

Configure the web app to the Standard App Service Tier
Enable autoscaling on the web app
Add a Scale condition
Add a scale rule

Reference:
https://docs.microsoft.com/en-us/learn/modules/app-service-autoscale-rules/6-autoscale-a-web-app
upvoted 37 times

AbidooKing 3 months, 2 weeks ago
Got this 2023July13. Score 917 and went with
Configure the web app to the Standard App Service Tier
Enable autoscaling on the web app
Add a Scale condition
Add a scale rule
upvoted 4 times

Mev4953 1 year, 9 months ago
There is no "Add a scale condition". Do you mean "Configure a scale condition" ?
upvoted 2 times

Chiboy 1 year, 8 months ago
It actually says " Add Scale Conditions" on the web app
upvoted 2 times

rkuifje 1 year, 9 months ago
Seems you are right Altafpatel based on your link. Thanks.
upvoted 2 times

rkuifje 1 year, 9 months ago
BTW, may be both answer are correct:
See question note: More than one order of answer choices is correct
upvoted 2 times

azurelearner666 2 years, 4 months ago
No, you must create a condition (a default one is created for you) and then add scale rules.
upvoted 18 times

sgsvve 2 years, 4 months ago
NO, NO, NO.... You can find it out yourself bij going to an appservice-> Scale Out -> Click Manual Scale (Now you have a default SCALE CONDITION) -> Click Add Scale Rule (Now you have a SCALE RULE)
upvoted 12 times

TEMPKAKAM Highly Voted 2 years, 11 months ago
Configure the web app to the Standard App Service Tier
Enable autoscaling on the web app
Add a Scale condition
Add a scale rule
upvoted 40 times

profesorklaus 2 years, 9 months ago

There is nothing about adding scale condition. You have default scale condition so first configure it then add a rule (85%)

upvoted 5 times

kakajos Most Recent 1 month, 1 week ago

Upgrade to Standard App Service tier (it supports autoscaling)
Enable autoscaling
Add a scale condition
Create scale rule
Monitor autoscaling activity

upvoted 1 times

NombreFalso 8 months, 1 week ago

"NOTE: More than one order of answer choices is correct." Makes me think the order in which you add the rule or scale condition does not matter, nothing else could be in different order IMO.

upvoted 1 times

macobuzi 2 months ago

Nice observation!

upvoted 1 times

OPT_001122 12 months ago

Configure the web app to the Standard App Service Tier
Enable autoscaling on the web app
Add a Scale condition
Add a scale rule

upvoted 2 times

Eltooth 1 year, 4 months ago

Configure the web app to the Standard App Service Tier
Enable autoscaling on the web app
Add a Scale condition
Add a scale rule

https://docs.microsoft.com/en-gb/learn/modules/scale-apps-app-service/4-autoscale-app-service

upvoted 2 times

adilkhan 9 months, 1 week ago

100% correct !

upvoted 1 times

Evo_Morales 1 year, 6 months ago

Sloppy answer set - this needs cleanup

upvoted 3 times

iamstudying 1 year, 7 months ago

answer is correct. les gettit, buddies. 6 figure here i come

upvoted 3 times

massnonn 1 year, 8 months ago

Once you enable autoscaling, you can edit the automatically created default scale condition, after you add scale rule

upvoted 1 times

sozturk88 1 year, 8 months ago

Configure the web app to the Standard App Service Tier
Enable autoscaling on the web app
Add a Scale condition
Add a scale rule

Reference:
https://docs.microsoft.com/en-us/learn/modules/app-service-autoscale-rules/6-autoscale-a-web-app

@altafpatel1984 is right

upvoted 1 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with most voted

upvoted 5 times

altafpatel1984 1 year, 11 months ago

Given answer is for Monitor auto scale. But question is for Web App auto scale and for this sequence will be as follow:

Configure the web app to the Standard App Service Tier
Enable autoscaling on the web app
Add a Scale condition
Add a scale rule

Reference: https://docs.microsoft.com/en-us/learn/modules/app-service-autoscale-rules/6-autoscale-a-web-app

upvoted 6 times

Pankaj78 1 year, 11 months ago

Step 1: Configure the web app to the Standard App Service Tier
The Standard tier supports auto-scaling, and we should minimize the cost.

Step 2: Enable autoscaling on the web app

Step 4: Configure a Scale condition

Step 3: Add a scale rule
Scale rule is within the Scale condition.

upvoted 4 times

Jan91 1 year, 12 months ago

Had this question in my exam this morning. It stated that there are 2 possible orders to answer. Most likely the scale condition and scale rule part is meant with this.

upvoted 3 times

catalene 2 years ago

I think, you only create "autoscaling rules app" with condition like ">85%" with Premium Plan. IT's Correct?

upvoted 1 times

AOE 2 years, 2 months ago

why not use consumption plan ?

upvoted 1 times

the_odd_one 1 week, 3 days ago

Consumption plans are only for Azure functions

upvoted 1 times

MontyKL 2 years, 1 month ago

I have the same question. I think there is no such a thing as consumption plan for App service. Consumption plan is for Azure Functions. I might be wrong.

upvoted 4 times

mlantonis 2 years, 5 months ago

Step 1: Configure the web app to the Standard App Service Tier
The Standard tier supports auto-scaling, and we should minimize the cost.

Step 2: Enable autoscaling on the web app

Step 3: Add a scale rule

Step 4: Add a Scale condition

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

upvoted 11 times

mlantonis 2 years, 5 months ago

Maybe...

Step 1: Configure the web app to the Standard App Service Tier
The Standard tier supports auto-scaling, and we should minimize the cost.

Step 2: Enable autoscaling on the web app

Step 4: Configure a Scale condition

Step 3: Add a scale rule
Scale rule is within the Scale condition.

upvoted 10 times

Question #2 *Topic 5*

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: Enable Application Request Routing (ARR).

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Instead deploy and configure Azure Cache for Redis. Update the web applications.

Reference:

https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching#managing-concurrency-in-a-cache

Community vote distribution

B (100%)

gematsaljoa Highly Voted 2 years, 11 months ago

B. NoNoNoNo

upvoted 30 times

Archimedes Highly Voted 2 years, 9 months ago

The correct answer is NO. Application Request Routing is for controlling internet traffic in IIS using a proxy server. It doesn't deal with managing state information.

upvoted 16 times

clarionprogrammer 2 years, 6 months ago

As stated, "No" is the correct answer. Application Request Routing is for scaling and load balancing.
https://www.iis.net/downloads/microsoft/application-request-routing

upvoted 5 times

OPT_001122 Most Recent 12 months ago

Selected Answer: B

No - redis cache is the answer

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer. NO

upvoted 2 times

Eche_lr73 1 year, 6 months ago

Selected Answer: B

Es Correcto, es NO

upvoted 3 times

AzureDJ 1 year, 7 months ago

B. No. ARR is not for session state. We typically use ARR module to host a reverse proxy on IIS.

upvoted 2 times

Freidrich 1 year, 8 months ago

Selected Answer: B

The correct answer is B: No.

upvoted 2 times

Baskman 1 year, 8 months ago

Correct answer: NO

Explanation: "Application Request Routing (ARR) [...] is a proxy-based routing module that forwards HTTP requests to content servers based on HTTP headers"

SOURCE: https://docs.microsoft.com/en-us/iis/extensions/planning-for-arr/using-the-application-request-routing-module

upvoted 1 times

glam 2 years, 5 months ago

B. No...

upvoted 4 times

AzureAz204Fan 2 years, 5 months ago

In Azure, you can store ASP.NET Session State in a SQL database, Azure Table Storage or an in-memory distributed cache such as REDIS.
https://docs.microsoft.com/en-us/archive/msdn-magazine/2014/august/microsoft-azure-use-distributed-cache-in-microsoft-azure

upvoted 6 times

pac1311 2 years, 9 months ago

correct!

upvoted 3 times

Question #3 *Topic 5*

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

Correct Answer: *B*

Instead deploy and configure Azure Cache for Redis. Update the web applications.

Reference:

https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching#managing-concurrency-in-a-cache

Community vote distribution

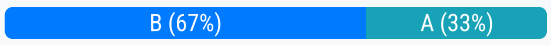

profesorklaus Highly Voted 2 years, 9 months ago

The worst solution from a performance and scalability standpoint is to use a database backed session state provider.

https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-aspnet-session-state-provider

upvoted 24 times

warchoon 8 months ago

maybe worst but exists
https://learn.microsoft.com/en-us/previous-versions/aspnet/ms178587(v=vs.100)

upvoted 2 times

ferut 2 years, 5 months ago

Although Postgres will lack the performance compared to Redis, the solution will work. Any thought?

upvoted 8 times

lugospod 1 year, 9 months ago

https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-aspnet-session-state-provider

Postgresql is not listed as viable option on MS site. So I guess this is why it is not acceptable.

upvoted 8 times

venki21 Highly Voted 2 years, 11 months ago

But the question says, does it meet the goal? So using a postgres db to save the session does meet the goal. Any thoughts on this

upvoted 22 times

Bartimaeus 1 year, 7 months ago

" Each question in the series contains a unique solution" - Redis is a better option.

upvoted 19 times

AshT1 2 years, 9 months ago

"Save full HTTP responses" so Postgre may will not be a good choice.

upvoted 5 times

Cornholioz 2 years, 11 months ago

It will meet the goal, but my take is that it is a huge Overkill to introduce PostgreSQL.

upvoted 3 times

matejka 2 years, 9 months ago

I think that it would be better to provide four options and you should pick the best one rather than four questions with YES/NO answers as there is no chance to review the futher options (further questions).

upvoted 3 times

Pooochie 2 years, 7 months ago

Got a simmilar question on udemy, where both soultions were marked as correct. Redis and Postgress

upvoted 8 times

dom271219 Most Recent 4 weeks, 1 day ago

B. No
why they would choose this DB while Microsoft sells Sql DB?

upvoted 1 times

75daltonic 3 months, 1 week ago

Selected Answer: B

I think No.

The question says: Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution.

There are three questions in examtopics about this:

1-Enable Application Request Routing (ARR).
2-Deploy and configure an Azure Database for PostgreSQL. Update the web applications.
3-Deploy and configure Azure Cache for Redis. Update the web applications.

And only one is the solution in order to the note at the beginning, so the best answer is Redis.

upvoted 6 times

ay_m 2 months, 2 weeks ago

Thanks for summing them up!

upvoted 2 times

JH81 4 months, 1 week ago

Selected Answer: B

MS is looking for Redis as the answer.

upvoted 1 times

thebarber87 6 months ago

Selected Answer: B

No - not best for the goal

upvoted 2 times

JamieS 7 months, 2 weeks ago

I'd do A, but B would work. Potentially poorly worded question if they want you to pick the best option.

upvoted 1 times

warchoon 8 months ago

ASP.NET session state is designed to enable you to store user session data in different sources. By default, session state values and information are stored in memory ... Another alternative is to store session data in a SQL Server database, where it can be shared by multiple Web servers.
Reference:
https://learn.microsoft.com/en-us/previous-versions/aspnet/ms178587(v=vs.100)

upvoted 1 times

warchoon 8 months ago

=> https://www.codeproject.com/Articles/18456/Sample-Session-Providers-with-PostgreSQL

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. No

An Azure Database for PostgreSQL is a fully-managed relational database service based on the PostgreSQL open-source database engine, it can be used to store data, but it's not designed to store session state information or HTML output, it will not meet the requirements of sharing 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, or save full HTTP responses for concurrent requests.

upvoted 3 times

rotimislaw 10 months ago

Selected Answer: B

The question description says "unique solution" so I'd say Redis is better and here the answer is B. NO

upvoted 2 times

mitu7 10 months ago

Selected Answer: B

Azure Database for PostgreSQL is a managed database service that can be used to store and retrieve data, but it is not designed for storing and managing session state information or HTML output. It does not provide the necessary features and capabilities to share session state across multiple web applications or support concurrent access to session state data.

upvoted 2 times

hubekpeter 11 months, 1 week ago

Selected Answer: A

Yes, we used to do it that way when redis/memcached key value in memory store wasn't available.

upvoted 1 times

OPT_001122 12 months ago

Selected Answer: B

No - redis cache is the correct ans

upvoted 4 times

gmishra88 1 year ago

Providing both session state and output-caching is Redis. It is fine to keep session state in a DB or in CosmosDB, but the output caching is not a good idea. So, wait for Redis as the answer. But yes, all of this is a guess on what Microsoft thinks

upvoted 1 times

Knightie 1 year, 2 months ago

Selected Answer: B

Deploy and configure an Azure Database for PostgreSQL. Update the web applications.
It didn't say storing the state or session or whatever else...

upvoted 3 times

Knightie 1 year, 1 month ago

Instead deploy and configure Azure Cache for Redis. Update the web applications.

Compare to the Redis Default answer, it also didn't say saving state.

upvoted 1 times

Ami_Nou 1 year, 5 months ago

Selected Answer: B

It says "Each question in the series contains a unique solution", so while this may work, deploying and configuring Azure Cache for Redis and updating the web applications is a better answer

upvoted 3 times

Amy_ant 1 year, 2 months ago

"Each question in the series contains a unique solution" not "Each series of questions contains a unique (valid) solution" all they are stating is that that each question in the series is unique you are not answering the same question. A postgres db can be written to and read concurrently it can store session data.

Redis would be a faster solution, postgres would be more persistent.

upvoted 1 times

Eche_lr73 1 year, 6 months ago

Selected Answer: B

Como dice: "Cada pregunta de la serie contiene una solucion unica" y Redis genera una solucion mas optima.

upvoted 2 times

Question #4 *Topic 5*

HOTSPOT -

A company is developing a gaming platform. Users can join teams to play online and see leaderboards that include player statistics. The solution includes an entity named Team.

You plan to implement an Azure Redis Cache instance to improve the efficiency of data operations for entities that rarely change.

You need to invalidate the cache when team data is changed.

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:

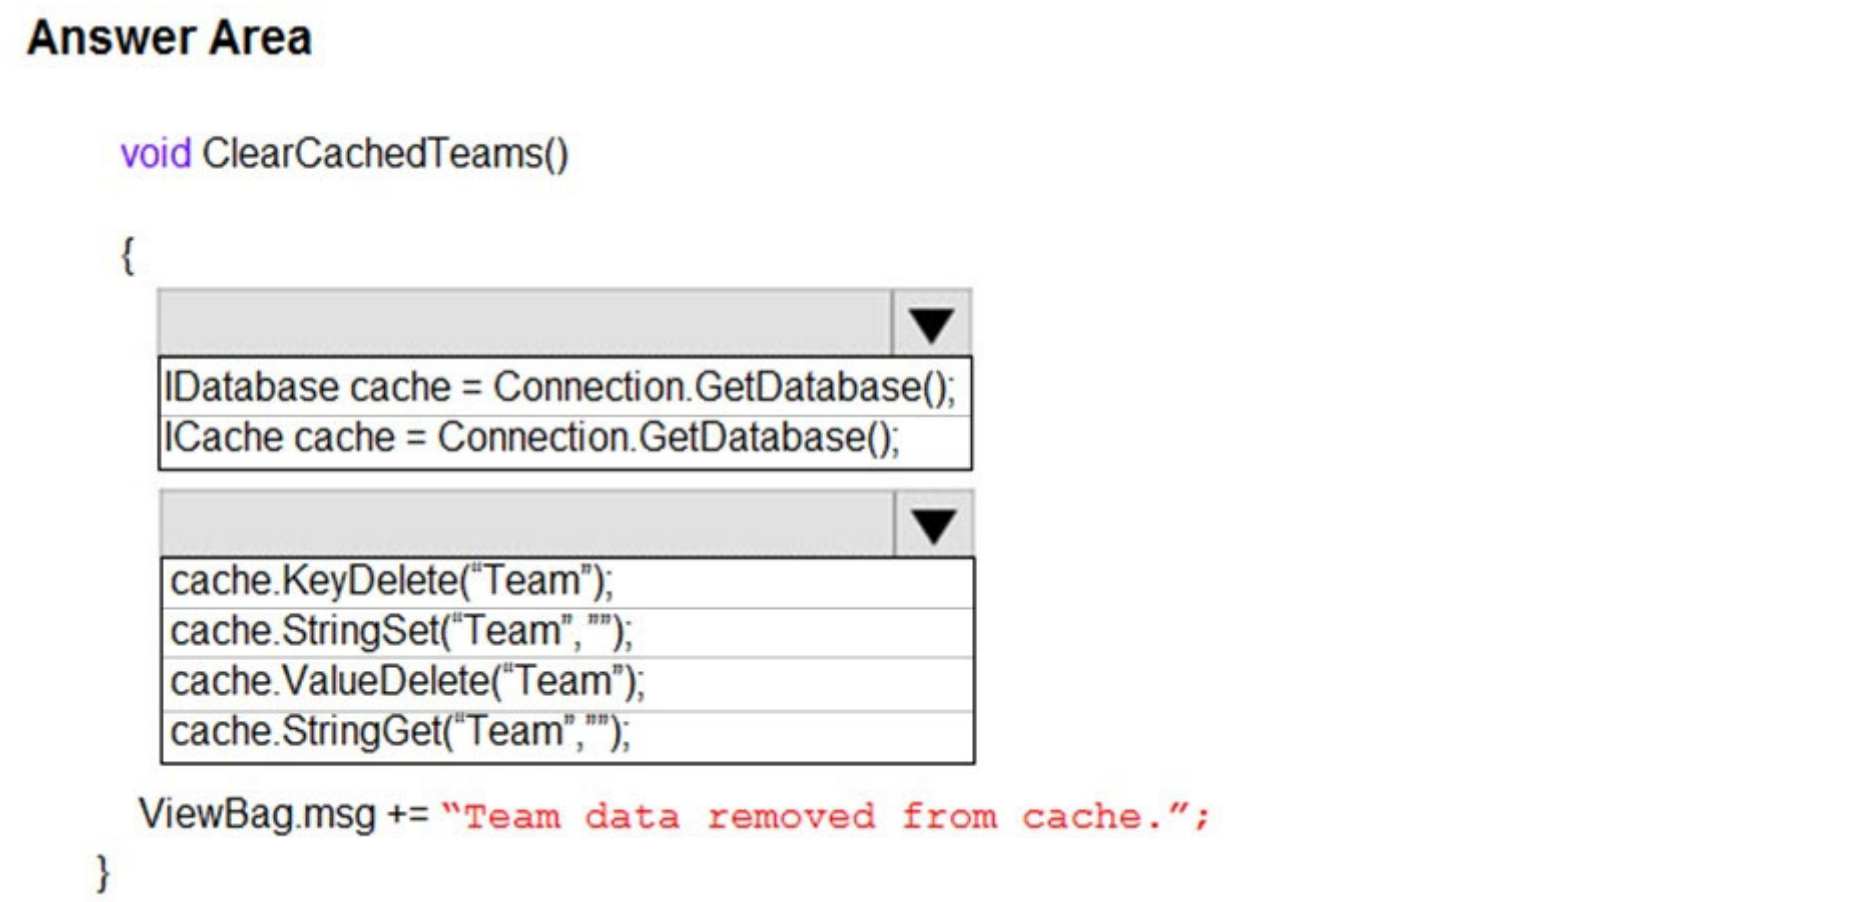

Correct Answer:

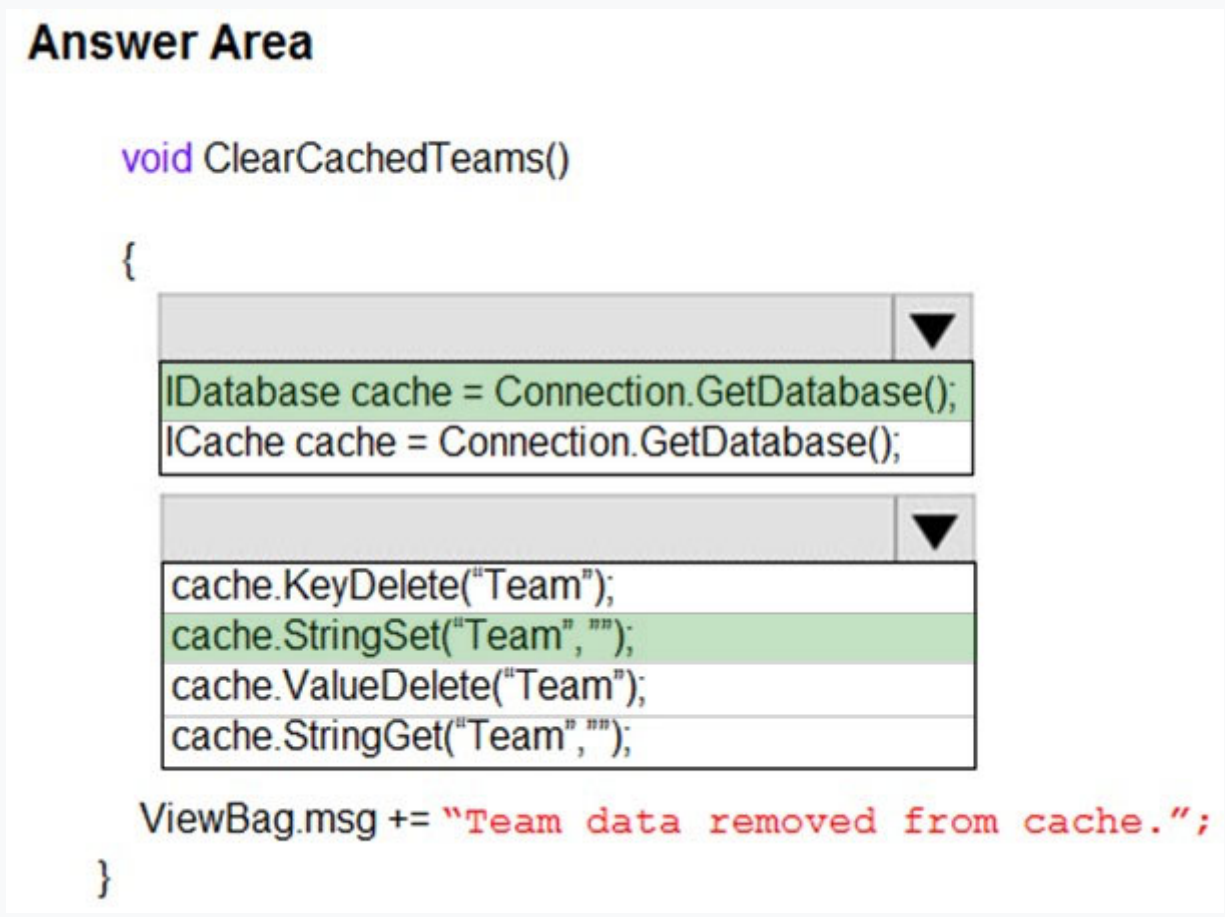

Box 1: IDatabase cache = connection.GetDatabase();

Connection refers to a previously configured ConnectionMultiplexer.

Box 2: cache.StringSet("teams","")

To specify the expiration of an item in the cache, use the TimeSpan parameter of StringSet. cache.StringSet("key1", "value1", TimeSpan.FromMinutes(90));

Reference:

https://azure.microsoft.com/sv-se/blog/lap-around-azure-redis-cache-preview/ https://docs.microsoft.com/en-us/cli/azure/webapp/config/container

agueda Highly Voted 2 years, 7 months ago

Same question on AZ203
Answer is:
IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("teams")
https://www.examtopics.com/discussions/microsoft/view/12706-exam-az-203-topic-5-question-9-discussion/

upvoted 112 times

MonkeyKing1024 8 months ago

https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-cache-aside-leaderboard#:~:text=void%20ClearCachedTeams()%0A%7B%0A%20%20%20%20IDatabase%20cache%20%3D%20Connection.GetDatabase()%3B%0A%20%20%20%20cache.KeyDelete(%22teamsList%22)%3B%0A%20%20%20%20cache.KeyDelete(%22teamsSortedSet%22)%3B%0A%20%20%20%20ViewBag.msg%20%2B%3D%20%22Team%20data%20removed%20from%20cache.%20%22%3B%0A%7D

upvoted 1 times

BrettusMaximus 2 years, 6 months ago

It is KeyDelete. If you use StringSet you will set Teams to the value of Empty string and a future request to Teams will return an empty string and not null

upvoted 18 times

coffecold 1 year ago

Please see the GetFromList() in the tutorial.
if (String.IsNullOrEmpty(serializedTeams)) then the DB will be read instead of the cache.
Implementation with an empty string returned from cache if invalidated is no problem.

upvoted 2 times

hobob Highly Voted 2 years, 7 months ago

Answer is incorrect (for the 2nd part).
cache.KeyDelete() is the correct method for removing a key from the redis cache.

upvoted 32 times

Sukon_Desknot 2 years, 2 months ago

Kindly provide references if you believe your knowledge is better than exam topics', we are not in the faith business 😁

upvoted 6 times

sas12321 2 years, 2 months ago

https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-cache-aside-leaderboard#update-the-teamscontroller-to-read-from-the-cache-or-the-database

upvoted 19 times

Firo Most Recent 3 months ago

Answer is:
IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("teams"),
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-cache-aside-leaderboard

upvoted 2 times

Kluk_Kluk 7 months, 2 weeks ago

https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-cache-aside-leaderboard
void ClearCachedTeams()
{
IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("teamsList");
cache.KeyDelete("teamsSortedSet");
ViewBag.msg += "Team data removed from cache. ";
}

upvoted 2 times

NeelParghi 8 months, 2 weeks ago

KeyDelete Delete the key/value.
Ref: https://learn.microsoft.com/en-us/training/modules/develop-for-azure-cache-for-redis/4-interact-redis-api

upvoted 2 times

rolling_potato_ 10 months ago

Was on my exam today (03-01-2023) I went with
IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("teams")

upvoted 2 times

st0rmtrooperx 10 months, 2 weeks ago

Got this on Dec 16th, 2022. I scored 921 and used given answer.

upvoted 1 times

coffecold 1 year ago

It is stringset, you want to have the key with the team name kept in your session.
Please see the GetFromList() code in the tutorial (https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-cache-aside-leaderboard)

if (String.IsNullOrEmpty(serializedTeams)) then the DB will be read instead of the cache and the cache is set again with string set.
Implementation with an empty string returned from cache if invalidated is no problem.

upvoted 2 times

serpevi 1 year, 1 month ago

Got it on 09/2022, went with: IDatabase cache = connection.GetDatabase(); cache.KeyDelete("teams",''). Score 927

upvoted 2 times

Eltooth 1 year, 4 months ago

IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("Team");

upvoted 1 times

lllleroiv 1 year, 4 months ago

IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("Team");

upvoted 1 times

liuliangzhou 1 year, 5 months ago

Answer is:
IDatabase cache = Connection.GetDatabase();
cache.StringSet("Team")
Per the subject mentioned: You need to invalidate the cache when team data is changed.
It is changing value, not delete the object. we can update value for next time with no need re-create "Team" object.
https://azure.microsoft.com/sv-se/blog/lap-around-azure-redis-cache-preview/

upvoted 5 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose as below:
1. IDatabase cache = connection.GetDatabase();
2. cache.KeyDelete("teams",'')

upvoted 3 times

leonidn 1 year, 9 months ago

IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("teamsList");

https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-cache-aside-leaderboard#update-the-teamscontroller-to-read-from-the-cache-or-the-database

upvoted 3 times

BeshoyRomany 2 years, 1 month ago

The Correct answer is
IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("teams")

We don't need to use setTeams with “TimeSpan.FromMinutes(90)” becuse the requirement is:
--> You need to invalidate the cache when "team data is changed"
When data is changed only
So we have to cache.KeyDelete("teams") to clear the cach once the "team data is changed"

upvoted 3 times

myuv1131 2 years, 1 month ago

The ClearCachedTeams method removes any cached team statistics from the cache.
void ClearCachedTeams()
{
IDatabase cache = Connection.GetDatabase();
cache.KeyDelete("teamsList");
cache.KeyDelete("teamsSortedSet");
ViewBag.msg += "Team data removed from cache. ";
}

upvoted 3 times

aruni_mishra 1 year, 11 months ago

https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-cache-aside-leaderboard#:~:text=The%20ClearCachedTeams%20method%20removes%20any%20cached%20team%20statistics%20from%20the%20cache.

upvoted 1 times

kondapaturi 2 years, 4 months ago

IDatabase cache = Connection.GetDatabase(); - we need to use the IDatabase interface
cache.KeyDelete(“teams”); - To invalidate or remove the cache item, we can simply delete the key

upvoted 3 times

Question #5 *Topic 5*

DRAG DROP -

A company has multiple warehouses. Each warehouse contains IoT temperature devices which deliver temperature data to an Azure Service Bus queue.

You need to send email alerts to facility supervisors immediately if the temperature at a warehouse goes above or below specified threshold temperatures.

Which five 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:

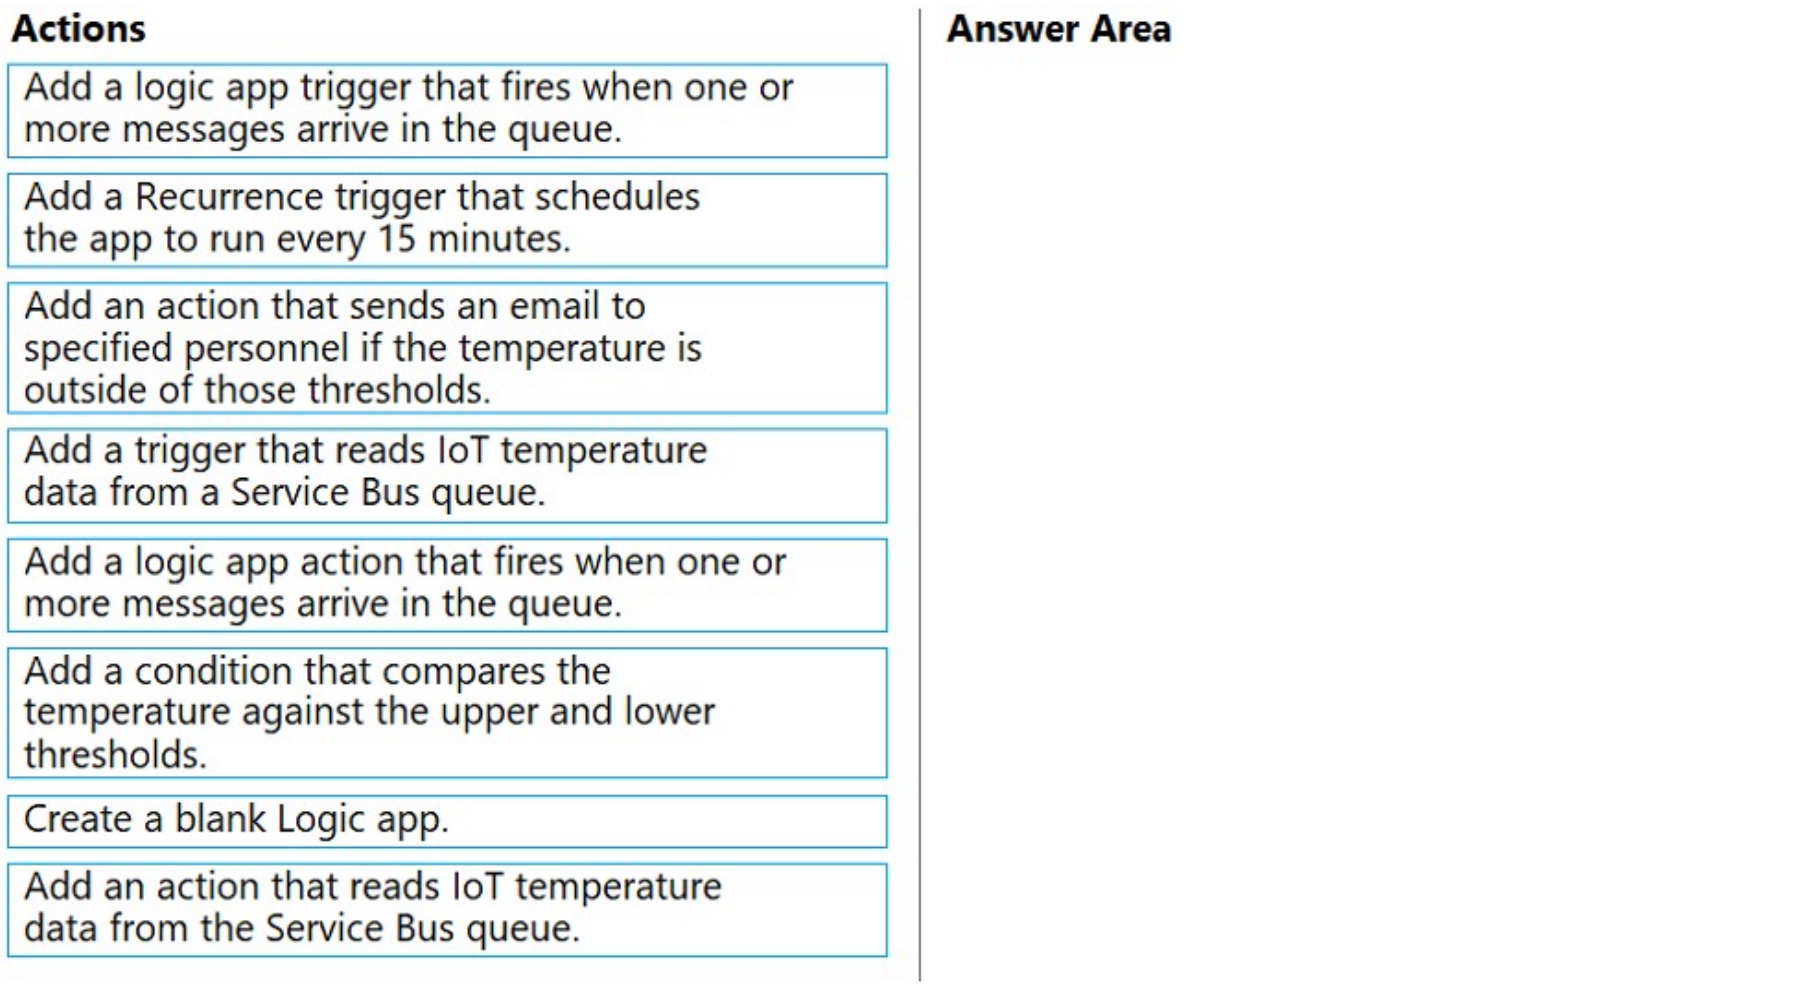

Correct Answer:

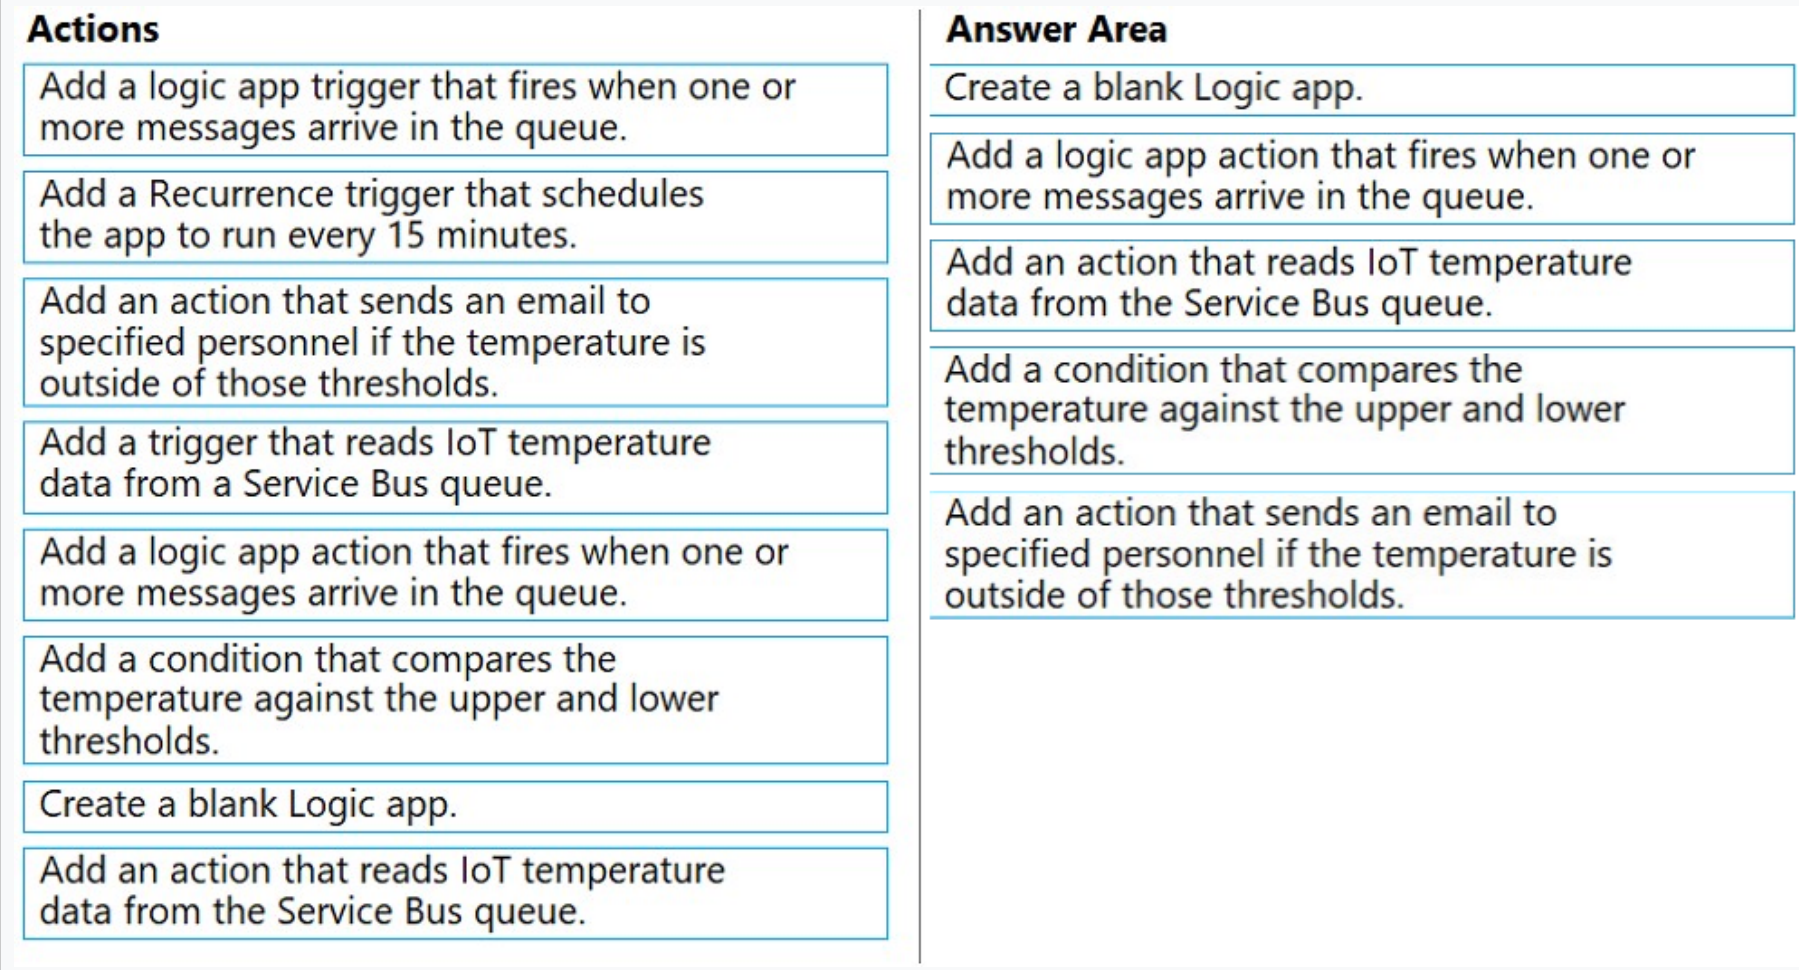

Step 1: Create a blank Logic app.

Create and configure a Logic App.

Step 2: Add a logical app trigger that fires when one or more messages arrive in the queue.

Configure the logic app trigger.

Under Triggers, select When one or more messages arrive in a queue (auto-complete).

Step 3: Add an action that reads IoT temperature data from the Service Bus queue

Step 4: Add a condition that compares the temperature against the upper and lower thresholds.

Step 5: Add an action that sends an email to specified personnel if the temperature is outside of those thresholds

Reference:

https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-monitoring-notifications-with-azure-logic-apps

agueda Highly Voted 2 years, 7 months ago

The 2nd step sholud be "logic app trigger". You need a trigger to fire the logic app

Ref: https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-monitoring-notifications-with-azure-logic-apps

upvoted 80 times

debanjan10 Highly Voted 2 years ago

Order:
1. Create blank logic app
2. Add logic app trigger when one or more messages arrive in the queue
3. Action to read IoT Temp data
4. Compare upper and lower temp
5. Send emails

upvoted 55 times

Knightie 1 year, 2 months ago

trigger instead of action, trigger is something kick starts asking someone to do the job..

upvoted 1 times

naivecoder786 Most Recent 7 months, 4 weeks ago

Logic Apps are out of Scope from AZ-204 now !

upvoted 8 times

warchoon 8 months ago

I think the trigger should bee after the action it uses. So 1 after 6.
Agree with the rest

upvoted 1 times

warchoon 8 months ago

Sorry
action should be created after trigger
https://learn.microsoft.com/en-us/azure/logic-apps/quickstart-create-first-logic-app-workflow
So the explanation is correct

upvoted 1 times

OPT_001122 11 months ago

this is a freaking question.

i try to remember it like a sentence
Black Logic Arrives in the Q Acts Compares Sends

Order:
1. Blank Logic
2. Arrive in the Queue
3. Action
4. Compare
5. Send

upvoted 5 times

Lucky_me 1 year, 4 months ago

The explanation is correct, the picture is incorrect.

upvoted 2 times

OPT_001122 11 months ago

Blank Logic Arrives in the Q Acts to Read Compares Sends

upvoted 1 times

liuliangzhou 1 year, 5 months ago

Order: create app -> trigger -> action -> condition -> do
1. Create blank logic app
2. Add logic app trigger when one or more messages arrive in the queue
3. Action to read IoT Temp data
4. Compare upper and lower temp
5. Send emails

upvoted 5 times

ivan0590 1 year, 6 months ago

Admin, please correct the given answer.
The second step is wrong, is a "logic app trigger", not an action.
Even the explanation says is a trigger!

upvoted 1 times

kondapaturi 2 years, 4 months ago

First , we need to create a blank logic app.
Then we need to add a trigger to the logic app which will fire when one or more messages arrives in the queue.
Then we need an action to read the IoT temperature data.
Then we need to have a condition that compares the temperature against the upper and lower thresholds.
And finally, we have to have an action that sends an email to the specified personnel.

upvoted 14 times

insanewriters 2 years, 4 months ago

Logic apps are supposed to be out of scope as of March 2021. Can anybody confirm they have gotten this question recently?

upvoted 14 times

glam 2 years, 5 months ago

Only Step 2 should be - Logic app triggers When one or more messages arrive in a queue .
Rest correct.

upvoted 4 times

Gajendran 2 years, 6 months ago

Steps and Answers are wrong. Steps given below the Answers are correct. Setp 2 should be - Logic app triggers When one or more messages arrive in a queue . Refer - https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-monitoring-notifications-with-azure-logic-apps

upvoted 2 times

Deputy_Cartman 2 years, 6 months ago

What is with this website and posting incorrect answers to questions and leaving them here?

upvoted 4 times

catsforthewin 2 years, 6 months ago

What is the correct answer for this question?

upvoted 2 times

jvyas 2 years, 6 months ago

Only step 2 needs to be replaced by " Add a trigger that fires when item arrives in queue" the rest of the steps are correct and in the same order as the provided answer.

upvoted 4 times

hobob 2 years, 7 months ago

Answer is incorrect for step 2. It should be to add a Service Bus Trigger when items arrive in a queue (ref: https://docs.microsoft.com/en-us/azure/connectors/connectors-create-api-servicebus#add-service-bus-trigger)

upvoted 10 times

pmsiva 2 years, 7 months ago

You are correct. Logic app must start with a Trigger.

upvoted 3 times

pmsiva 2 years, 7 months ago

Add the RSS trigger
Every logic app must start with a trigger, which fires when a specific event happens, or when a specific condition is met. Each time the trigger finds new items, it fires and the Logic Apps engine creates a logic app instance that starts and runs your workflow. If the trigger doesn't find new items, the trigger doesn't fire and doesn't create an instance or run the workflow at this check.

upvoted 1 times

Question #6 *Topic 5*

DRAG DROP -

You develop an ASP.NET Core MVC application. You configure the application to track webpages and custom events.

You need to identify trends in application usage.

Which Azure Application Insights Usage Analysis 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:

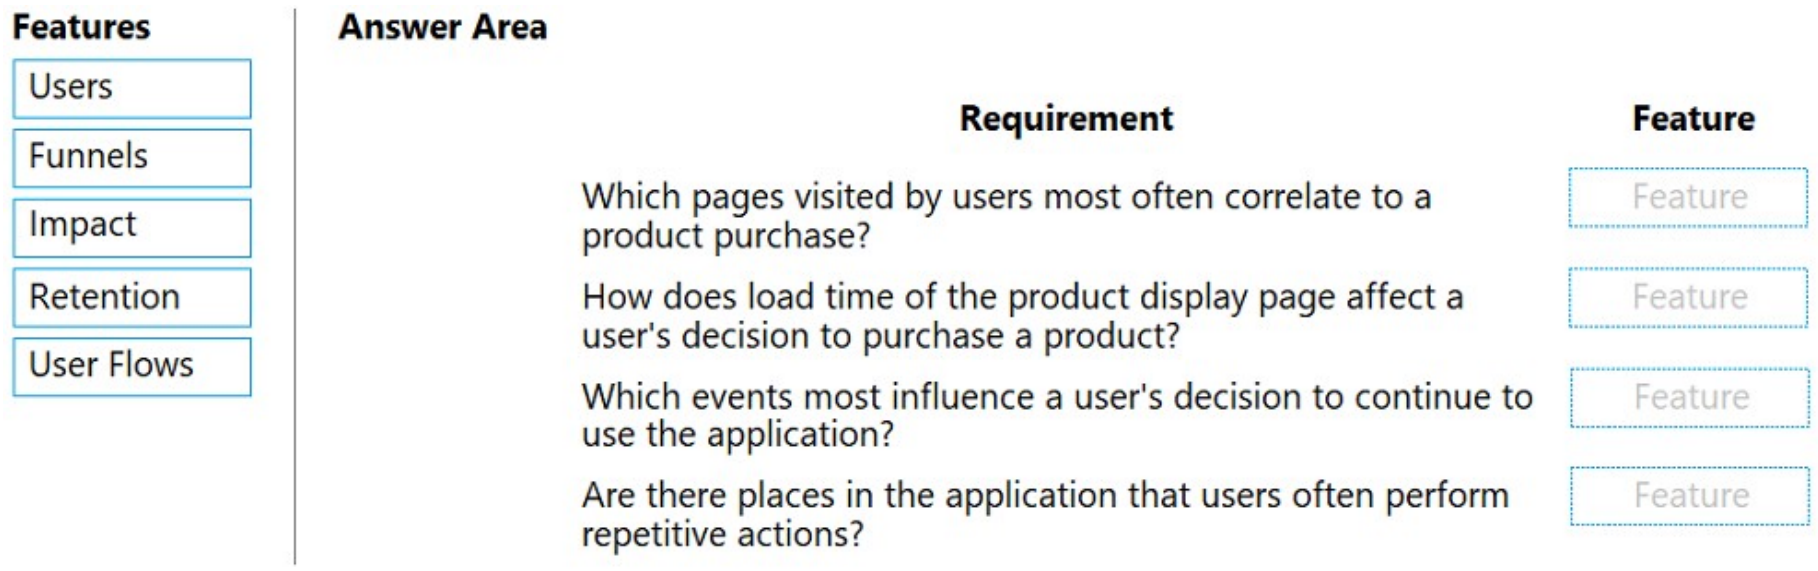

Correct Answer:

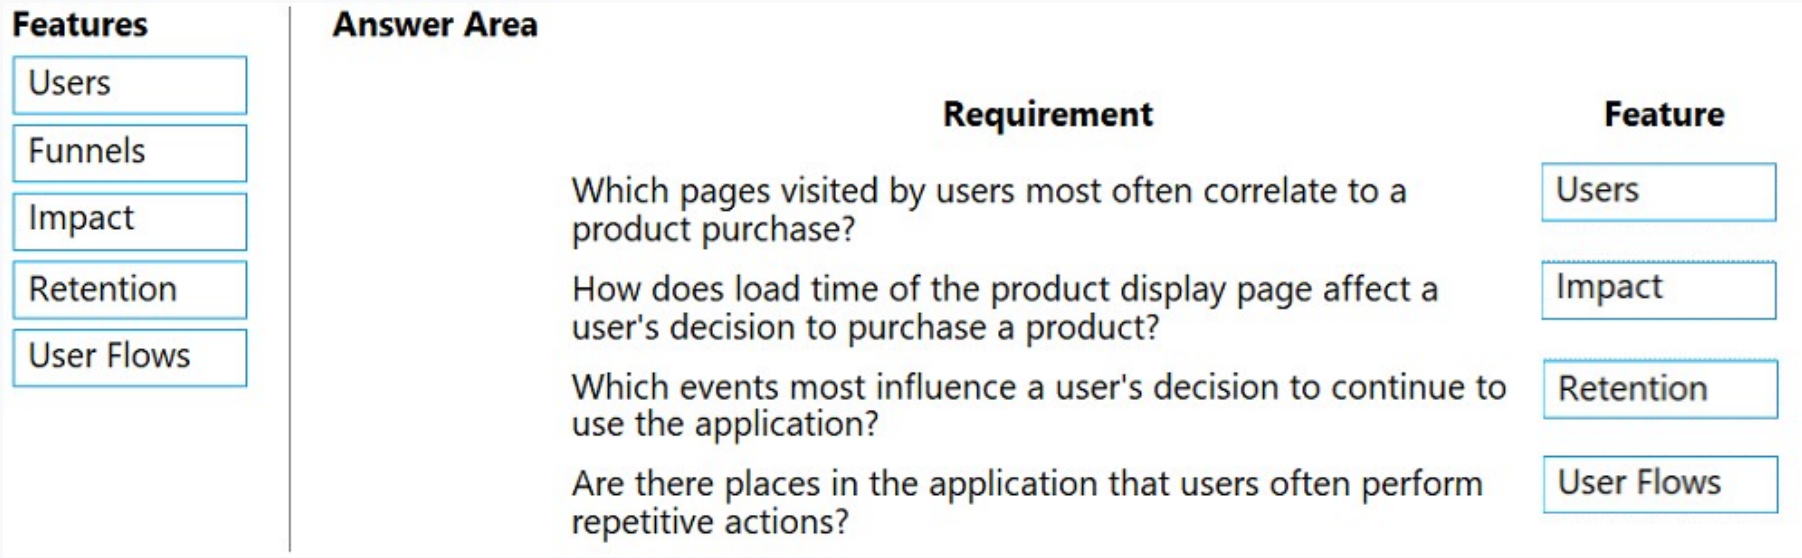

Box 1: Users -

Box 2: Impact -

One way to think of Impact is as the ultimate tool for settling arguments with someone on your team about how slowness in some aspect of your site is affecting whether users stick around. While users may tolerate a certain amount of slowness, Impact gives you insight into how best to balance optimization and performance to maximize user conversion.

Box 3: Retention -

The retention feature in Azure Application Insights helps you analyze how many users return to your app, and how often they perform particular tasks or achieve goals. For example, if you run a game site, you could compare the numbers of users who return to the site after losing a game with the number who return after winning. This knowledge can help you improve both your user experience and your business strategy.

Box 4: User flows -

The User Flows tool visualizes how users navigate between the pages and features of your site. It's great for answering questions like:

☞ How do users navigate away from a page on your site?

☞ What do users click on a page on your site?

☞ Where are the places that users churn most from your site?

☞ Are there places where users repeat the same action over and over?

Incorrect Answers:

Funnel: If your application involves multiple stages, you need to know if most customers are progressing through the entire process, or if they are ending the process at some point. The progression through a series of steps in a web application is known as a funnel. You can use Azure Application Insights Funnels to gain insights into your users, and monitor step-by-step conversion rates.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-impact

Dinima Highly Voted 2 years, 7 months ago

1st one Funnels

upvoted 55 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose:
Funnels
Impact
Retention
User Flows

upvoted 8 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000.

upvoted 4 times

clarionprogrammer 2 years, 6 months ago

The end-goal of a 'Funnel' is a product purchase. 'Users' is just about how users use your app.

upvoted 9 times

renzoku 1 year, 3 months ago

but you need to know which page is visited, with funnel you have to know it previously (what specific page) and get how many people are using this page, I think that probably could be "Users" because you have to verify N° users that visited your pages

upvoted 1 times

kondapaturi Highly Voted 2 years, 4 months ago

1.Funnels
2.Impact
3.Retention
4.User flow

upvoted 39 times

JH81 Most Recent 4 months ago

Got this on 6/28/2023 and passed with 850. Went with highly voted answer.

upvoted 4 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

Had cases:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 3 times

OPT_001122 12 months ago

1.Funnels
2.Impact
3.Retention
4.User flow

upvoted 5 times

gmishra88 1 year ago

The answers given are fully correct. Absolutely sure.

upvoted 1 times

gmishra88 1 year ago

Except the first one, which I'm not sure:
Users it says : "Find out when people use your web app, what pages they're most interested in"
Funnels it says: "You can use Application Insights Funnels to gain insights into your users, and monitor step-by-step conversion rates."

Pages most interested: Users
Conversion/Buying product: Funnel

After deep meditation I will bet on Funnel is what microsoft guy thought as the right answer. Or it might be mean to increase mental health of people so that they meditate

upvoted 4 times

serpevi 1 year, 1 month ago

Got it on 09/2022, went with 1.Funnels 2.Impact 3.Retention 4.User flow. Score 927

upvoted 2 times

Knightie 1 year, 2 months ago

the answer is correct.. they all match the keywords of the individual paragraph.
1. no of visit.
2. loading time.

3. user coming back.
4. repeated page flow. (Are there places where users repeat the same action over and over?)

upvoted 2 times

ivan0590 1 year, 6 months ago

After some investigation, I think the given answer is correct.
The question is "Which pages visited by users most often correlate to a product purchase?".
So, if I'm right, they want to know how often visiting a page leads to a purchase and in which of those pages this happens more often. Let's say if I have a web app with a catalogue of 500 products and every product has a detail page from which the user has the option to purchase the product. Now, I want to know the top 10 product detail pages that ended up leading to a purchase.
Maybe I'm wrong, but I don't think that Funnels would be appropriate here. When you create a funnel, you have to specify the steps of the flow. I think that the main goal of a funnel is to analyse the transitions between steps in order to check if users get stuck in a step. Or in other words, how many people completed the flow.
Funnels are to be used in known and specific flows, like when a user signup. That's not what the question is about. The question is about knowing what pages tend to lead more to a purchase. This is not a specific flow, you don't know the exact steps.

upvoted 6 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 2 times

AzureDJ 1 year, 7 months ago

Did you pass and did you go with?
1.Funnels (Funnels correlate to purchases, not general user events)
2.Impact
3.Retention
4.User flow

upvoted 2 times

Manivannan19 1 year, 8 months ago

1. Users
2. Impact
3. Retention
4. User Flows

upvoted 2 times

barretowski 1 year, 7 months ago

ta errado

upvoted 2 times

ytingyeu 1 year, 8 months ago

Doesn't what Funnels care about is "if users stop somewhere during the whole process", not the correlation between pages and products?

For example, creating a customer ticket takes 5 steps and Funnel shows lots of visitors give up creating one at the 4th step. Then we know the UX in the 4th step might need some enhancement.

upvoted 1 times

leonidn 1 year, 9 months ago

The Users feature does not allow to analyze flows. Funnels allows analyzing conversion rate through the specified flow. It does not allow us to find out what led to the final step. Users flow offers a means to identify user behavior. That is correlate the most to the first point.

1. UserFlow
2. Impact
3. Retention
4. UserFlow

upvoted 2 times

phvogel 1 year, 12 months ago

First one funnels: https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-funnels
To quote: "Before you create your funnel, decide on the question you want to answer. For example, you might want to know how many users are viewing the home page, viewing a customer profile, and creating a ticket." Just replace "creating a ticket" with "buying a product"

upvoted 7 times

nonoss 2 years ago

- 1st one can't be "users", because users always returns a number of users verifying a certain condition, see the screenshot here https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-segmentation

upvoted 1 times

HimanshuNankani 2 years ago

I guess it should be

UserFlow
Impact
Retention
UserFlow

The first one asks "Which pages..", how do we get which pages in Users?
And Funnel also signifies on how much customer interaction is there, and for that too it should be following a particular route specified while creating the funnel.

So, I think User Flow should be the right choice. Please correct me if I am wrong.

upvoted 4 times

myuv1131 2 years, 1 month ago

1. User
Users tool: How many people used your app and its features. Users are counted by using anonymous IDs stored in browser cookies. A single person using different browsers or machines will be counted as more than one user.
https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-segmentation

upvoted 1 times

Question #7 *Topic 5*

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 ג€"Location $Location -Organization $Org -AdminEmail $AdminEmail

D. New-AzureRmApiManagementBackendProxy -Url $ApiUrl

Correct 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.
Reference:
https://docs.microsoft.com/en-us/powershell/module/azurerm.apimanagement/new-azurermapimanagementbackendproxy?view=azurermps-6.13.0

Community vote distribution
A (68%) C (27%) 5%

Vano6k Highly Voted 2 years, 8 months ago
I think the correct answer is A because it is the one tuhat takes advantage of the swagger definition of the API?
upvoted 81 times

MiraA 2 years, 1 month ago
There is an example of importing an API from an Open Api Link using Import-AzApiManagementApi cmdlet:
https://docs.microsoft.com/en-us/powershell/module/az.apimanagement/import-azapimanagementapi#example-4--import-an-api-from-a-open-api-link

And as mentioned here:
https://docs.microsoft.com/en-us/powershell/module/azurerm.apimanagement/import-azurermapimanagementapi
the Import-AzureRmApiManagementApi cmdlet is obsoleted by Import-AzApiManagementApi cmdlet.

So I would select answer A.
upvoted 5 times

zero0 2 years, 4 months ago
Totally agree.
find the 2 line code here.
https://docs.microsoft.com/en-us/powershell/module/azurerm.apimanagement/import-azurermapimanagementapi?view=azurermps-6.13.0
upvoted 4 times

pac1311 2 years, 8 months ago
Think so too, the accepted answer has to be wrong.

upvoted 1 times

Eduarv2015 2 years, 8 months ago

It does not says that the API has swagger, D is correct.

upvoted 3 times

fadikh 2 years, 8 months ago

Swagger and OpenAPI specs are the same

upvoted 14 times

trance13 2 years, 7 months ago

They are not the same.

upvoted 2 times

dhishkiyaau 2 years, 2 months ago

OpenAPI = The specification itself, formerly known as Swagger specification. Swagger = Tools used in the implementation of OpenAPI.

upvoted 12 times

10x Highly Voted 2 years, 8 months ago

D should be the answer - it is about accessing API - not creating it (eg based on swagger)

upvoted 15 times

XYZ2 2 years, 6 months ago

but you should have one first, that's why you need to create it before able to use

upvoted 1 times

borfavor 2 years, 6 months ago

"The news API back end is implemented as a RESTful service and uses an OpenAPI specification." This clearly tells us the API already exists

upvoted 4 times

asdasdasg2 1 year, 9 months ago

It exists in reality, but not on API management. API management API != backend API

upvoted 4 times

katrang Most Recent 1 month, 2 weeks ago

A
When importing certain APIs, API Management configures the API backend automatically. For example, API Management configures the backend web service when importing:

An OpenAPI specification.
A SOAP API.
Azure resources, such as an HTTP-triggered Azure Function App or Logic App.

upvoted 1 times

uffuchsi 8 months, 1 week ago

Selected Answer: A

Since we already have an API back end, we need to import to Azure API Management.

https://docs.microsoft.com/en-us/azure/api-management/scripts/powershell-import-api-and-add-to-product?toc=/powershell/module/toc.json

upvoted 2 times

Sriniv 8 months, 2 weeks ago

Answer is B reason

To configure an Azure API Management instance for a RESTful service that uses an OpenAPI specification, you can use the New-AzApiManagementBackend cmdlet to define the backend service and the Set-AzApiManagementApi cmdlet to add an API to the API Management instance. Here's an example Azure PowerShell command:

New-AzApiManagementBackend -ResourceGroupName "myResourceGroup" -Name "myBackend" -Url "https://news-api.example.com" -Protocol "https" -Title "News API"

Set-AzApiManagementApi -Context "myApiManagementInstance" -SpecificationFormat "OpenApi" -SpecificationUrl "https://news-api.example.com/openapi.json" -Path "/news" -DisplayName "News API" -BackendUrl "https://news-api.example.com" -BackendProtocol "https" -BackendId "myBackend"

upvoted 1 times

Sriniv 8 months, 2 weeks ago

I think answer is b
New-AzApiManagementBackend -ResourceGroupName "myResourceGroup" -Name "myBackend" -Url "https://news-api.example.com" -Protocol "https" -Title "News API"
Set-AzApiManagementApi -Context "myApiManagementInstance" -SpecificationFormat "OpenApi" -SpecificationUrl "https://news-api.example.com/openapi.json" -Path "/news" -DisplayName "News API" -BackendUrl "https://news-api.example.com" -BackendProtocol "https" -BackendId "myBackend"

This command creates a new backend service named "myBackend" that points to the URL of the news API, and then creates a new API in the

specified Azure API Management instance that references the OpenAPI specification for the news API. The API is mapped to the "/news" path and is associated with the backend service. The BackendId parameter specifies the ID of the backend service that was created earlier.

upvoted 1 times

OPT_001122 11 months ago

Selected Answer: A

A. Import-AzureRmApiManagementApi -Context $ApiMgmtContext -SpecificationFormat "Swagger" -SpecificationPath $SwaggerPath -Path $Path

upvoted 1 times

POOOJAAAAAAAAAA 10 months, 1 week ago

not sure A or C ..
New-AzApiManagement -Name "myapim" -ResourceGroupName "myResourceGroup" `
-Location "West US" -Organization "Contoso" -AdminEmail "admin@contoso.com"
Reference: https://learn.microsoft.com/en-us/azure/api-management/powershell-create-service-instance

upvoted 1 times

elequiel 11 months, 3 weeks ago

Selected Answer: A

Letter A because you need to create using OpenAPI

upvoted 1 times

OPT_001122 12 months ago

Selected Answer: A

correct answer is A

upvoted 1 times

coffecold 1 year ago

-AzureRmApiManagement can be *-AzApiManagement in future exams
AzureRM will be replaced by AZ Powershell.

upvoted 2 times

gmishra88 1 year ago

Seems like Microsoft has no review mechanism. Look at how the vote is so divided. Microsoft should make this an open book exam. They can be assured people still will not get good scores. Isn't that their intention?

upvoted 2 times

macobuzi 2 months ago

More people will have to retake the exam, more money for Microsoft!

upvoted 1 times

Satish_Babu 1 year, 2 months ago

Selected Answer: D

D is correct

upvoted 1 times

HumbleYolo 1 year, 3 months ago

Selected Answer: C

whizlab has c

upvoted 4 times

vavra 1 year, 6 months ago

Selected Answer: C

It is a creation, and docs https://docs.microsoft.com/en-us/powershell/module/azurerm.apimanagement/new-azurermapimanagement?view=azurermps-6.13.0 says there are obligatory params

-ResourceGroupName <String>
-Name <String>
-Location <String>
-Organization <String>
-AdminEmail <String>

It must be C.

upvoted 2 times

SivajiTheBoss 1 year, 7 months ago

Selected Answer: A

A seems more promising

upvoted 1 times

ReniRechner 1 year, 8 months ago

Selected Answer: A

A: Import also creates. Importing also creates an already configured instance (opposed to C)
B: We don't need another backend, we need a frontend for an exisiting backend
C: creates a new APIM, but you would still need to manually configure it.
D: is meant to proxy a public API (https://petri.com/how-to-use-the-azure-api-management-to-proxy-a-public-api)

=> A looks most promising

upvoted 6 times

fearoffree 1 year, 8 months ago

It seems the trick here is to have the knowledge of API Proxy that could be used for verifying the API availability without a need the need to call the backend APIs. But the question is very vague!

upvoted 1 times

Question #8 *Topic 5*

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 SessionID property to the ReplyToSessionId property.

B. Assign the value of the hazard message MessageId property to the DevileryCount property.

C. Assign the value of the hazard message SessionID 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 SequenceNumber property to the DeliveryCount property.

F. Assign the value of the hazard message MessageId property to the SequenceNumber property.

Correct Answer: *AD*
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.
A: 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:

B, E: 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.
Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads

Community vote distribution
AD (100%)

Kitkit Highly Voted 2 years, 8 months ago
Answer is correct.
ReplyToSessionId (reply-to-group-id) This value augments the ReplyTo information and specifies which SessionId should be set for the reply when sent to the reply entity.
CorrelationId (correlation-id) 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.
upvoted 41 times

rustycables Highly Voted 2 years, 2 months ago
I have no clue what this is about, but I get the feeling we are supposed to arrive at the correct answer by elimination.
upvoted 25 times

coffecold 1 year ago
You're right : DeliveryCount, SequenceNumber are read only, two remaining.
upvoted 2 times

uffuchsi Most Recent 8 months, 1 week ago
I believe A + D is correct
upvoted 1 times

Sriniv 8 months, 2 weeks ago

Assign the value of the hazard message MessageId property to the CorrelationId property: When an alarm event is published to the Service Bus, set the MessageId property to a unique value that identifies the message. Then, when the alarm controller receives the message, set the CorrelationId property to the value of the MessageId property. This enables the controller to associate the message with the corresponding transaction record for auditing purposes.

upvoted 2 times

Sriniv 8 months, 2 weeks ago

Assign the value of the hazard message SessionID property to the ReplyToSessionId property: When an alarm event is published to the Service Bus, set the SessionID property to a unique value that identifies the session. Then, when the alarm controller receives the message, set the ReplyToSessionId property to the value of the SessionID property. This enables the controller to reply to the message and maintain the context of the session.

upvoted 1 times

Sriniv 8 months, 2 weeks ago

Therefore, the two actions that you should perform are A and D:

A. Assign the value of the hazard message SessionID property to the ReplyToSessionId property.
D. Assign the value of the hazard message MessageId property to the CorrelationId property.

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: AD

Answer is correct.

upvoted 2 times

OPT_001122 12 months ago

Selected Answer: AD

Answer is correct.

upvoted 3 times

Praks13 1 year, 1 month ago

AD
https://docs.microsoft.com/en-us/training/modules/discover-azure-message-queue/5-messages-payloads-serialization

upvoted 1 times

meoukg 1 year, 7 months ago

https://www.examtopics.com/discussions/microsoft/view/44385-exam-az-204-topic-4-question-11-discussion/
Got it on 03/2022, I chose as below:
A. Assign the value of the hazard message SessionID property to the ReplyToSessionId property.
D. Assign the value of the hazard message MessageId property to the CorrelationId property.

upvoted 2 times

ScubaDiver123456 1 year, 9 months ago

Selected Answer: AD

I arrived at A+D using the explanations given here

https://docs.microsoft.com/en-us/learn/modules/discover-azure-message-queue/5-messages-payloads-serialization

upvoted 3 times

ehurfheiz 1 year, 9 months ago

Selected Answer: AD

Answer seems correct : AD

upvoted 2 times

glam 2 years, 5 months ago

correct.

upvoted 3 times

SnakePlissken 2 years, 5 months ago

Answer is correct. Nice example to show the purpose of CorrelationId and ReplyToSessionId.

upvoted 4 times

kwaazaar 2 years, 6 months ago

what reply entity? this question makes no sense.
a simple picture of the setup would greatly clarify things

upvoted 6 times

MrZoom 2 years, 7 months ago

Also, all the other answers are really kind of obviously wrong...

upvoted 2 times

Question #9 *Topic 5*

You are developing an Azure function that connects to an Azure SQL Database instance. The function is triggered by an Azure Storage queue.
You receive reports of numerous System.InvalidOperationExceptions with the following message:
`Timeout expired. The timeout period elapsed prior to obtaining a connection from the pool. This may have occurred because all pooled connections were in use and max pool size was reached.`
You need to prevent the exception.
What should you do?

A. In the host.json file, decrease the value of the batchSize option

B. Convert the trigger to Azure Event Hub

C. Convert the Azure Function to the Premium plan

D. In the function.json file, change the value of the type option to queueScaling

Correct Answer: *C*
With the Premium plan the max outbound connections per instance is unbounded compared to the 600 active (1200 total) in a Consumption plan.
Note: The number of available connections is limited partly because a function app runs in a sandbox environment. One of the restrictions that the sandbox imposes on your code is a limit on the number of outbound connections, which is currently 600 active (1,200 total) connections per instance. When you reach this limit, the functions runtime writes the following message to the logs: Host thresholds exceeded: Connections.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/manage-connections https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#service-limits

Community vote distribution
A (90%) 10%

VK7Az204 Highly Voted 2 years, 7 months ago
A is the correct answer
upvoted 69 times

djffcnhhyiyaigyvuy Highly Voted 2 years, 7 months ago
The error shown is a SQL Server error not an Azure functions error.
https://forums.asp.net/t/2004198.aspx?Timeout+expired+The+timeout+period+elapsed+prior+to+obtaining+a+connection+from+the+pool+This+may+have+occurred+because+all+pooled+connections+were+in+use+and+max+pool+size+was+reached+
The actual issue appears to be that you have too many simultaneous functions running, the solution is to limit the batch size (# functions that can run in parallel)
https://social.msdn.microsoft.com/Forums/azure/en-US/a2955297-1c14-45f2-b799-6346b340519a/how-does-batchsize-works-in-hostjson?forum=AzureFunctions
upvoted 29 times

rdemontis 2 years, 7 months ago
Correct! We have to consider we are speaking about Azure Function with Azure Storage Queue Trigger. This is a case where execution could happens in parallel based on the queue batchSize property
(see https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue).

I think this is the batchSize property referred in the question (and not that in the functions aggregator. That's another thing).

Besides, we have to consider the origin of the error message: Azure Sql Database!! So upgrading the plan the problem could only get worse thing.
upvoted 8 times

coffecold 1 year ago
This article explains well and easy : https://medium.com/@hammadarif/throttling-the-scalability-of-azure-functions-v2-ddb8625eeedc
It gives also the answer why upgrading to premium won't help.
upvoted 1 times

coffecold 1 year ago
B.t.w usage in host.json :
{
"queues": {
"batchSize": 1,
"newBatchThreshold": 7
}

}
It has nothing to do with the aggrator batchsize
upvoted 2 times

NPE_ Most Recent 5 days ago
The answer should be A.
The error message shows that there is not enough connections, which means that the concurrency is too high. Too many instances are running parallel. So we have to reduce the concurrency of the app.

https://github.com/Azure/azure-functions-host/wiki/host.json
If you search for "batchSize" on the above official microsoft github wiki page for host.json, you will find the explanation for the batchSize:
// The number of queue messages to retrieve and process in
// parallel (per job function). The default is 16 and the maximum is 32.
"batchSize": 16,
upvoted 1 times

Lola2023 1 month, 2 weeks ago
Selected Answer: C
Explanation:
With the Premium plan the max outbound connections per instance is unbounded compared to the 600 active (1200 total) in a Consumption plan.
Note: The number of available connections is limited partly because a function app runs in a sandbox environment. One of the restrictions that the sandbox imposes on your code is a limit on the number of outbound connections, which is currently 600 active (1,200 total) connections per instance. When you reach this limit, the functions runtime writes the following message to the logs: Host thresholds exceeded: Connections.
upvoted 1 times

BaoNguyen2411 3 months ago
got this question on 29/06/2023
upvoted 2 times

macobuzi 2 months ago
What option did you choose?
upvoted 2 times

JOSEEVILLASMIL 7 months, 1 week ago
Answer A, the error is at SQL Instance no a function error, and the questions does not say the plan you have... so you cannot asume that you are using consumtion plan.
upvoted 1 times

Sriniv 8 months, 2 weeks ago
None of the options provided are correct I think
upvoted 1 times

Sriniv 8 months, 2 weeks ago
the most appropriate solution in this case is to optimize the function code and increase the Max Pool Size property of the connection string used by the function.
upvoted 1 times

alexein74 9 months, 1 week ago
Selected Answer: A
A is the correct answer
upvoted 1 times

OPT_001122 12 months ago
Selected Answer: A
A is the correct answer
upvoted 3 times

gmishra88 1 year ago
This is a question for a machine. I have to remember all options in host.json. How many of these random questions can be asked by Microsoft.
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads
https://learn.microsoft.com/en-us/azure/azure-functions/functions-host-json#aggregator
According to the above link batchSize is about application insights. Another marvel from Microsoft to make it without application insights anywhere and cherry on top to ask that mistake in an exam. Especially when batching will make people think it has something to do with batching the calls. Microsoft wins again
https://learn.microsoft.com/en-us/azure/azure-functions/functions-host-json#aggregator
upvoted 2 times

gmishra88 1 year ago
This is the way to limit scale out, but guess that is not there in the options
https://learn.microsoft.com/en-us/azure/azure-functions/event-driven-scaling#limit-scale-out
upvoted 1 times

gmishra88 1 year ago

And yes, there is another batchSize that could be for the storage queue. But then we have to imagine that could be the unwritten scenario.
upvoted 1 times

Knightie 1 year, 1 month ago

it can not be C when Pool connection is running out, it's able the DB connections pool is out, changing the function plan only allow more callers to access the DB but it does not increase the connections pool size, so it will still be out of connections in DB.
upvoted 1 times

Knightie 1 year, 2 months ago

Selected Answer: A

It's the DB connection Pool..... you scale up the app also no use.. the connectors from DB Pool not release fast enough for the next app to operate.
upvoted 1 times

Dani_ac7 1 year, 3 months ago

Selected Answer: A

A is the correct answer
upvoted 2 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose A. In the host.json file, decrease the value of the batchSize option
upvoted 4 times

ReniRechner 1 year, 8 months ago

Selected Answer: C

A: seems not correct for me.
see: https://docs.microsoft.com/en-us/azure/azure-functions/functions-host-json#aggregator
it just seems to affect metrics. Besides an increase would help => more elements per function => less functions => less connections

B: might help (if scaling can be limited), but for itself is not a solution

C: does not help since source of problem (to many incoming connections to SQL Server) unchanged. There is no problem with number of outgoing connections in function

D: searching for "queueScaling" only resulted in AZ-204 question results ;-)

So what should I choose?

While doing research for point B I git to this site
https://medium.com/microsoftazure/azure-functions-limiting-throughput-and-scalability-of-a-serverless-app-5b1c381491e3

I found the hint that a premium plan can limit instances, which is verified here:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-premium-plan?tabs=portal#maximum-function-app-instances

So I would go with C but for the option to limit scaling and not for the explanation in the result.
upvoted 1 times

iamstudying 1 year, 7 months ago

I agree with your reason for not A.. batching is supposed to reduce number of IO requests and thus connections to SQL DB

The Medium post you referenced is outdated. You can actually limit the scaling on both Consumption and Premium plans:
https://docs.microsoft.com/en-us/azure/azure-functions/event-driven-scaling#limit-scale-out

I don't think there is a right answer here... pick A and prAy, buddies
upvoted 2 times

massnonn 1 year, 8 months ago

Selected Answer: A

here: https://stackoverflow.com/questions/56489553/host-json-meaning-of-batchsize
upvoted 2 times

Question #10 *Topic 5*

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

Correct Answer: *A*
The session state provider for Azure Cache for Redis enables you to share session information between different instances of an ASP.NET web application.
The same connection can be used by multiple concurrent threads.
Redis supports both read and write operations.
The output cache provider for Azure Cache for Redis enables you to save the HTTP responses generated by an ASP.NET web application.
Note: Using the Azure portal, you can also configure the eviction policy of the cache, and control access to the cache by adding users to the roles provided. These roles, which define the operations that members can perform, include Owner, Contributor, and Reader. For example, members of the Owner role have complete control over the cache (including security) and its contents, members of the Contributor role can read and write information in the cache, and members of the
Reader role can only retrieve data from the cache.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching

Community vote distribution
A (100%)

Prakash4691 Highly Voted 2 years, 7 months ago
Correct answer
upvoted 33 times

THAMORAY Highly Voted 2 years, 5 months ago
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.
This is the unique solution
A: Yes.
upvoted 8 times

OPT_001122 Most Recent 12 months ago
Selected Answer: A
Correct answer
upvoted 3 times

Eltooth 1 year, 4 months ago
Selected Answer: A
A is correct answer.
upvoted 3 times

mattvasc 1 year, 8 months ago
Selected Answer: A
Redis will do the trick.
upvoted 1 times

glam 2 years, 5 months ago
A. Yes

upvoted 5 times

roshansir 2 years, 5 months ago

A is the correct answer

upvoted 3 times

Question #11 *Topic 5*

HOTSPOT -

You are debugging an application that is running on Azure Kubernetes cluster named cluster1. The cluster uses Azure Monitor for containers to monitor the cluster.

The application has sticky sessions enabled on the ingress controller.

Some customers report a large number of errors in the application over the last 24 hours.

You need to determine on which virtual machines (VMs) the errors are occurring.

How should you complete the Azure Monitor query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
let startTimestamp = [                    ▼] ;
                     ago(1d)
                     since(1d)
                     totimespan(1d)
                     date(now() - 1d)
let ContainerIDs = KubePodInventory
| where ClusterName == "Cluster1"
| [                                        ▼] ;
                     top ContainerID
                     union ContainerID
                     sample ContainerID
                     distinct ContainerID
ContainerLog
| [                                        ▼]
                     fork containerIDs
                     where ContainerID in (ContainerIDs)
                     restrict ContainerID in (ContainerIDs)
                     join ContainerID = = ContainerIDs.ContainerID
| where TimeGenerated > startTimestamp
| where LogEntrySource = = "stderr"
| [                                        ▼]
                     project by Computer
                     summarize by Computer
                     partition count() by Computer
                     summarize count() by Computer
```

Correct Answer:

Answer Area

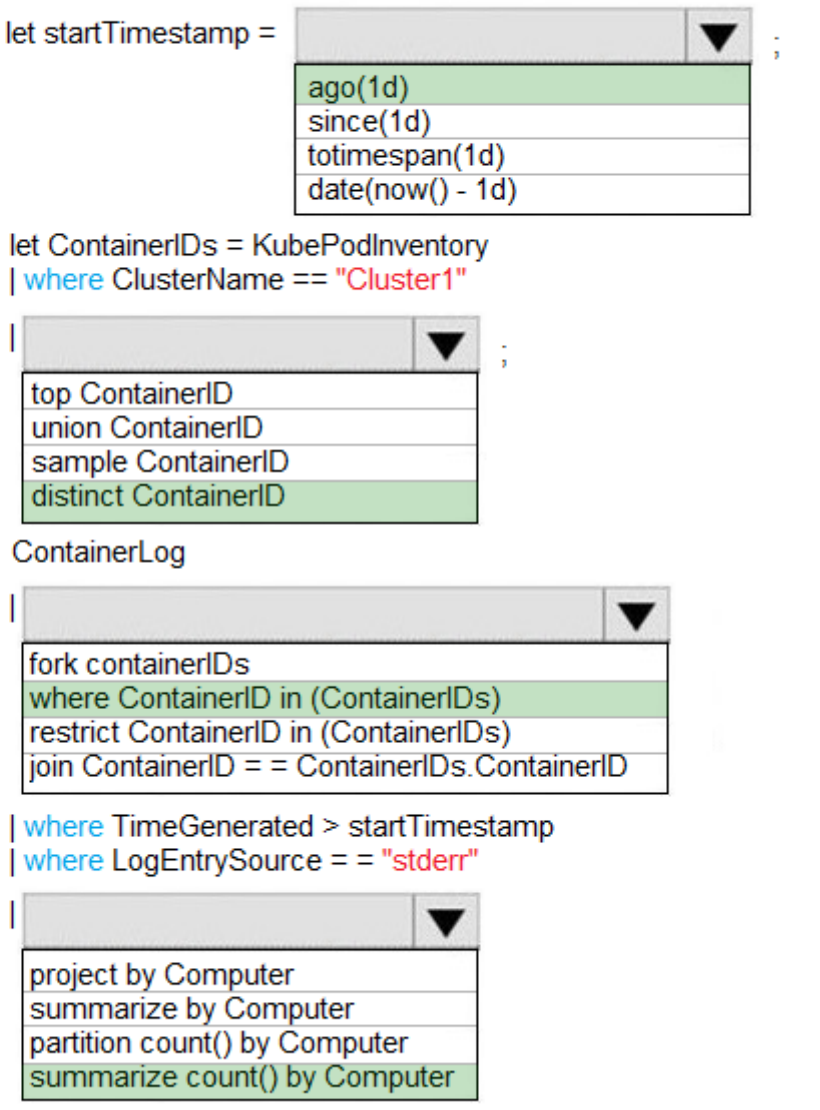

Box 1: ago(1d)

Box 2: distinct containerID -

Box 3: where ContainerID in (ContainerIDs)

Box 4: summarize Count by Computer

Summarize: aggregate groups of rows

Use summarize to identify groups of records, according to one or more columns, and apply aggregations to them. The most common use of summarize is count, which returns the number of results in each group.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/get-started-queries https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/query-optimization

Marusyk Highly Voted 2 years, 8 months ago

the answer is correct

upvoted 39 times

ReyPirata 2 months, 1 week ago

On my exam 2023-08-20. Scored 925
ago(1)
distinct ContainerID
where ContainerID in (ContainerIDs)
summarize count() by Computer

upvoted 4 times

pac1311 Highly Voted 2 years, 9 months ago

Isn't kubernetes out of scope for AZ-204?

upvoted 13 times

hicloud 2 months, 1 week ago

I guess if kubernetes is not the only product involved in the question then the question still appears

upvoted 1 times

gmishra88 1 year ago

There are many questions from kubernetes

upvoted 2 times

Chilred 2 years, 9 months ago

No, i think there are still some question from az203

upvoted 3 times

phuria 2 years, 8 months ago
You'll notice it really isn't about kubernetes per say, but how to query your logs from Log Analytics
upvoted 15 times

surprise0011 6 months, 2 weeks ago
yep, it is more about monitoring
upvoted 1 times

WH16 Most Recent 2 months, 3 weeks ago
Kubernetes out if scope as of August 2023
upvoted 3 times

macobuzi 2 months ago
Better safe than sorry!
upvoted 1 times

rockstar89 3 months ago
Got it exam 31/07/2023
upvoted 4 times

juanckar 3 months, 3 weeks ago
This was on the exam (July 2023). Went with proposed. Scored 917
upvoted 6 times

uffuchsi 8 months, 1 week ago
The question asks for errors in last 24 hours, so startTimestap should be 'ago (1d)'.

We need to distinct the logs for each container from set of containers. So, distinct based on ContainerID and where condition on container Ids. We need to determine the VM. So, summarize count() by computer.

https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/get-started-queries
https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/query-optimization
upvoted 1 times

Aterniss 10 months, 2 weeks ago
Got this 15/12/2022
upvoted 5 times

surprise0011 6 months, 2 weeks ago
thanks for mentioning. Maybe Kubernetes is out of scope but this question is more about monitoring
upvoted 1 times

elequiel 1 year ago
Got it in exam 20/10/2022
upvoted 5 times

Knightie 1 year, 2 months ago
the VMs here mean Containers.... sigh....
upvoted 1 times

petitbilly 1 year, 7 months ago
Got it in exam 03/22
upvoted 3 times

oescm 1 year, 8 months ago
Got this one 02/2022, so this question is NOT out of scope. I went with the provided answer
upvoted 7 times

lugospod 1 year, 9 months ago
Got this one 01/2022. Went with most voted (to avoid writing answers again)
upvoted 9 times

MasterQuestMaster 1 year, 9 months ago
I got this as well.
upvoted 3 times

glam 2 years, 5 months ago
correct
upvoted 4 times

diligent176 2 years, 9 months ago
ago(1d) is correct.
date(now()-1d) is not correct syntax (query won't run or compile at all).
upvoted 6 times

HimanshuNankani 2 years ago

Should have told the correct syntax as well.

upvoted 1 times

m_siri 2 years, 11 months ago

I think the answer to the 1st drop should be;
let start Timestamp = date(now() - 1d)

because we need to take the logs occured within the last 24 hours, not for the previous day

upvoted 4 times

adilkhan 9 months ago

no ago is correct..!

upvoted 1 times

mitcheld 2 years, 10 months ago

I think it should be ago.
Search for "Summarize: aggregate groups of rows" at https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/get-started-queries

upvoted 1 times

ALittleBunny 2 years, 10 months ago

Wrong. It should be ago(1d) as ago(a_timespan) subtracts the given timespan from the current UTC clock time. https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/examples#search-application-level-events-described-as-cryptographic

upvoted 6 times

stylebc 2 years, 10 months ago

You are not right, correct is ago(1d), as from ago function description:
https://docs.microsoft.com/en-us/azure/data-explorer/kusto/query/agofunction
a_timespan - Interval to subtract from the current UTC clock time (now()).

upvoted 19 times

Question #12 *Topic 5*

HOTSPOT -

You plan to deploy a web app to App Service on Linux. You create an App Service plan. You create and push a custom Docker image that contains the web app to Azure Container Registry.

You need to access the console logs generated from inside the container in real-time.

How should you complete the Azure CLI command? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

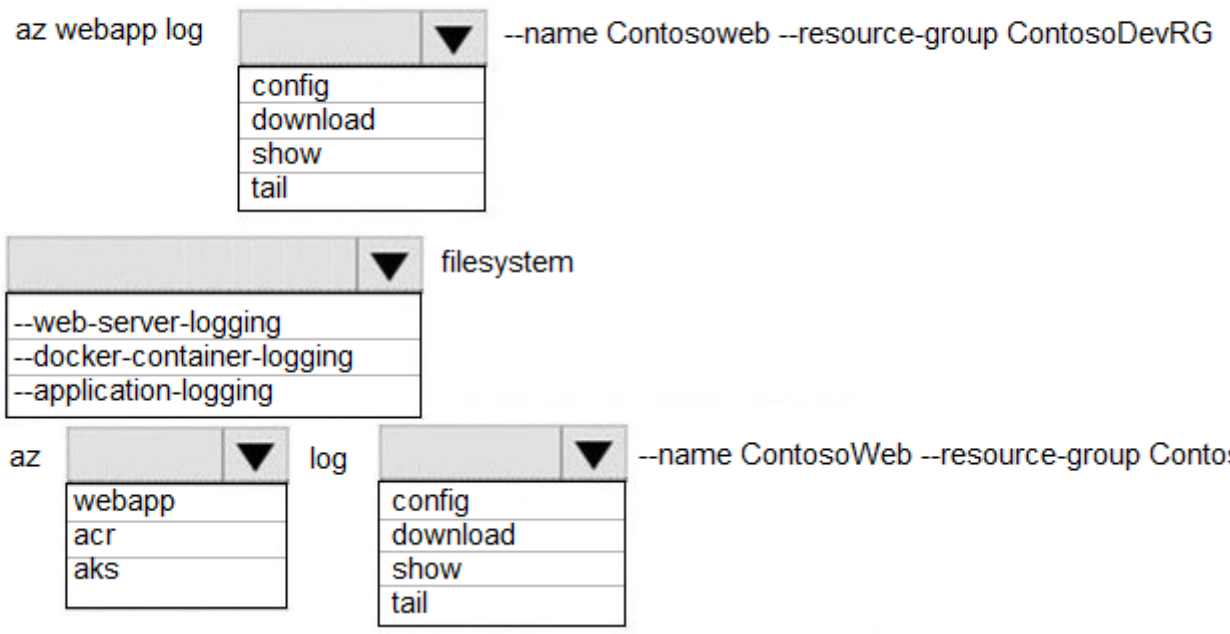

Correct Answer:

Answer Area

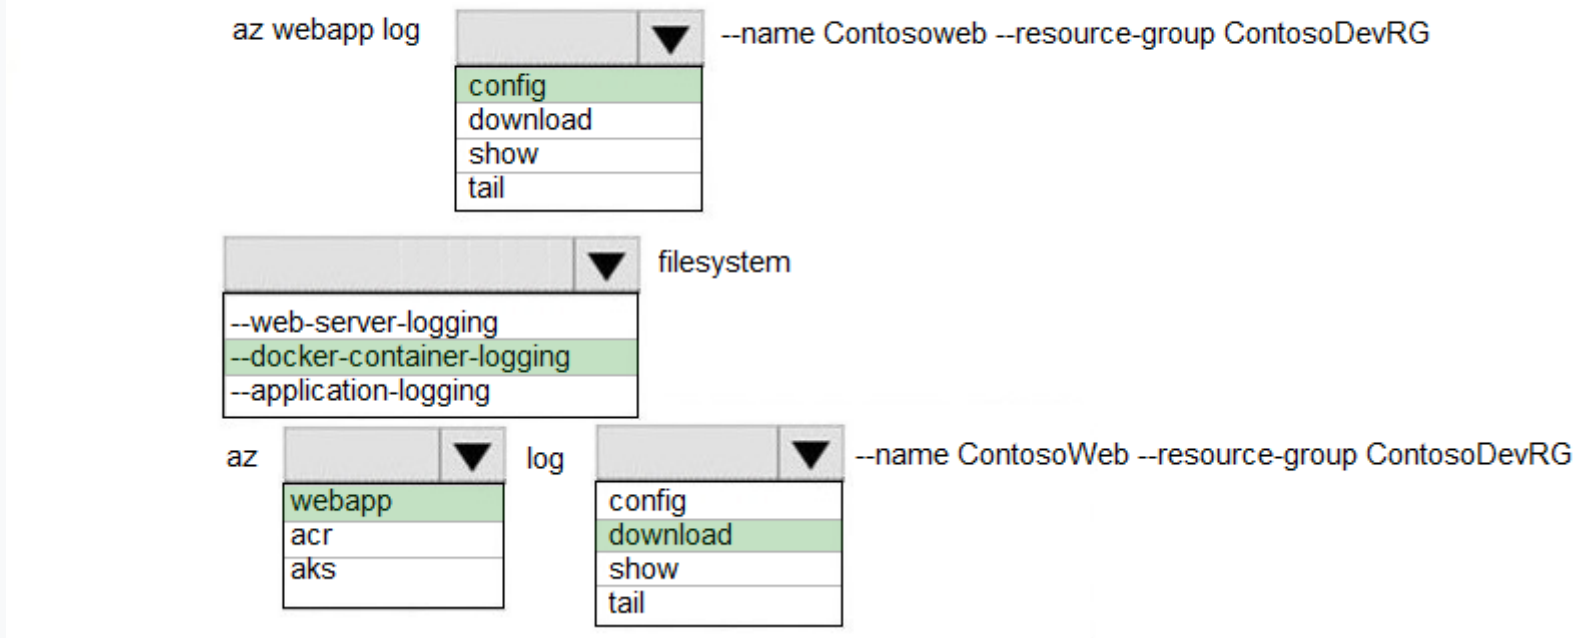

Box 1: config -

To Configure logging for a web app use the command:

az webapp log config

Box 2: --docker-container-logging

Syntax include:

az webapp log config [--docker-container-logging {filesystem, off}]

Box 3: webapp -

To download a web app's log history as a zip file use the command: az webapp log download

Box 4: download -

Reference:

https://docs.microsoft.com/en-us/cli/azure/webapp/log

granitoula Highly Voted 2 years, 12 months ago

It never said dowload, it says "access real time"
Will go with TAIL:

https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az_webapp_log_tail
upvoted 143 times

OPT_001122 12 months ago
Box 1: config
Box 2: --docker-container-logging
Box 3: webapp

above 3 are correct and
Box 4 : Tail
upvoted 12 times

Kellyb85 2 years, 11 months ago
I think download
in the question--> You plan to deploy a web app to App Service on Linux.
az-webapp-log-tail --> https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az_webapp_log_tail --> This command may not work with web apps running on Linux.
upvoted 6 times

TakumaK 2 years, 5 months ago
See the MSDN doc again. It obviously says the download is also "This command may not work with web apps running on Linux.". And the question clearly mentions "in real-time". Ok. let's think do you want to download the log zip file in real time? How?
upvoted 1 times

Ave 2 years, 11 months ago
It says the same for download.
upvoted 9 times

iiiihhhh 2 years, 9 months ago
There is a request for clarification on github:
WebApp:az webapp log tail and insufficient information about linux https://github.com/Azure/azure-cli/issues/16203
But there is no answer.
upvoted 1 times

Adol 2 years, 9 months ago
You're right and it's tail according to a course in Udemy for the exact question
upvoted 9 times

AfroYeti 2 years, 9 months ago
No see, that interesting becasue there is a notice on the port itself and it reads "tail might not be available for linux containers"

Now the question is, is that answer the "safe" answer based on the possibility that it might not be available?
upvoted 1 times

DSK 2 years, 8 months ago
The suggested download command will not work because download needs a local file system location.
upvoted 1 times

monniq 2 years, 8 months ago
It says "This command may not work with web apps running on Linux." on https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az_webapp_log_tail
So, it can't be the answer.
upvoted 8 times

rdemontis 2 years, 7 months ago
I think TAIL is correct because it is the only way to access console log in real time. The sentence "This command may not work with web apps running on Linux" is present for every option, even for download, see https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az_webapp_log_download
upvoted 3 times

jay158 2 years, 4 months ago
"This command may not work with web apps running on Linux." is removed https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az_webapp_log_tail.
Warning is only for 'az webapp log download'
upvoted 6 times

Robert12345Robert 2 years, 3 months ago
az webapp log download
Download a web app's log history as a zip file.
This command may not work with web apps running on Linux.

It is download that may not work, not tail.
upvoted 8 times

RaviKS Highly Voted 2 years, 9 months ago

Correct Answers are
config
docker-container-logging
webapp
tail
upvoted 70 times

juanckar Most Recent 3 months, 3 weeks ago
This was on the exam (July 2023). Went with highly voted. Scored 917
upvoted 4 times

JH81 4 months ago
Got this on 6/28/2023 and passed with 850. Went with highly voted answer.
upvoted 1 times

ExamDev 4 months, 1 week ago
Ok we got Linux and documentation says:
- az webapp log download
Download a web app's log history as a zip file.
!!! This command may not work with web apps running on Linux. !!!

https://learn.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az-webapp-log-download
upvoted 1 times

Vmwarevirtual 5 months, 1 week ago
Box 1: config
Box 2: --docker-container-logging
Box 3: webapp
Box 4 : Tail
Question appeared in the exam I took at 27/May 2023
upvoted 2 times

AzBeginner 5 months, 1 week ago
On Exam : 22 May 2023
upvoted 1 times

aragones 5 months, 3 weeks ago
Got this 2023-05-12.

Had cases:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada
upvoted 2 times

uffuchsi 8 months, 1 week ago
To access the console logs generated from inside the container, first, turn on container logging by running the following command:

az webapp log config --name <app-name> --resource-group <resource-group-name> --docker-container-logging filesystem

Once container logging is turned on, run the following command to see the log stream:

az webapp log tail --name <app-name> --resource-group <resource-group-name>

If you don't see console logs immediately, check again in 30 seconds.

To stop log streaming at any time, type Ctrl+C.

https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-linux#access-diagnostic-logs-1
upvoted 3 times

elequiel 1 year ago
Got it in exam 20/10/2022
upvoted 4 times

Knightie 1 year, 2 months ago
No Download, but Tail
https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az-webapp-log-tail
upvoted 1 times

mahidharKollipara 1 year, 2 months ago
it shold be tail. Download does not work in linux.
upvoted 1 times

Eltooth 1 year, 4 months ago
config
docker-container-logging

webapp
tail

upvoted 2 times

Eltooth 1 year, 4 months ago

https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az-webapp-log-config

upvoted 1 times

Sandeep12093 1 year, 4 months ago

We don't want to download.also it said real time so definitely answer is "tail"

upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 2 times

KhaiHoang 1 year, 7 months ago

the last answer should be "tail"

Start live log tracing for a web app.
https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az-webapp-log-tail

upvoted 1 times

iamdamzy 1 year, 10 months ago

Definitely got be *tail* - "access in real-time is the keyword"
https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest#az_webapp_log_tail

upvoted 1 times

Question #13 *Topic 5*

You develop and deploy an ASP.NET web app to Azure App Service. You use Application Insights telemetry to monitor the app.
You must test the app to ensure that the app is available and responsive from various points around the world and at regular intervals. If the app is not responding, you must send an alert to support staff.
You need to configure a test for the web app.
Which two test types can you use? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. integration

B. multi-step web

C. URL ping

D. unit

E. load

Correct Answer: *BC*
There are three types of availability tests:
➪ URL ping test: a simple test that you can create in the Azure portal.
➪ Multi-step web test: A recording of a sequence of web requests, which can be played back to test more complex scenarios. Multi-step web tests are created in
Visual Studio Enterprise and uploaded to the portal for execution.
➪ Custom Track Availability Tests: If you decide to create a custom application to run availability tests, the TrackAvailability() method can be used to send the results to Application Insights.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability

Community vote distribution
BC (100%)

Ash111 Highly Voted 2 years, 11 months ago
Given answer is correct
upvoted 42 times

Secure01 2 years, 11 months ago
True:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability
There are three types of availability tests:
URL ping test
Multi-step web test
Custom Track Availability Tests
upvoted 11 times

pac1311 Highly Voted 2 years, 8 months ago
Correctemundo!
upvoted 9 times

BaoNguyen2411 Most Recent 3 months ago
got this question on 29/07/2023
upvoted 3 times

uffuchsi 8 months, 1 week ago
B and C
upvoted 1 times

Sriniv 8 months, 2 weeks ago
Note that multi-step web tests, load tests, and unit tests are not suitable for this scenario as they do not test the availability and responsiveness of the app from various points around the world at regular intervals. Multi-step web tests can simulate user scenarios, but they do not test the app from various points around the world. Load tests can test the performance of the app, but they do not test the availability and responsiveness of the app. Unit tests are used to test individual code components and are not suitable for testing the app as a whole.
upvoted 1 times

Sriniv 8 months, 2 weeks ago
Answer is A and C reason is below
Integration tests: These tests simulate the user experience by using a browser to interact with the web app, and they can check for specific page

content, errors, and response times. You can configure integration tests using the Azure Portal or Visual Studio.

URL ping tests: These tests send a GET request to a specific URL of the web app to check for a valid HTTP response and response time. URL ping tests are the simplest and most common type of web app test. You can configure URL ping tests using the Azure Portal or the Application Insights API.

upvoted 1 times

elequiel 1 year ago

Got it in exam 20/10/2022

upvoted 6 times

Eltooth 1 year, 4 months ago

Selected Answer: BC

B and C are correct answers.

upvoted 2 times

Eltooth 1 year, 4 months ago

Multi-step web tests are deprecated:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-multistep

upvoted 1 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with the given answers

upvoted 4 times

lugospod 1 year, 9 months ago

Got similar one 01/2022. But tests needed to also test for SSL and custom headers.. went with multistep and Standard test (not ping or custom track).

upvoted 7 times

Henksc 1 year, 9 months ago

Correct, however Multi-step web test is deprecated.
Source:
Multi-step web tests have been deprecated. We recommend using TrackAvailability() to submit custom availability tests instead of multi-step web tests. With TrackAvailability() and custom availability tests, you can run tests on any compute you want and use C# to easily author new tests.

https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-multistep

upvoted 3 times

ehurfheiz 1 year, 9 months ago

Selected Answer: BC

Answer seems to be correct : BC

upvoted 1 times

glam 2 years, 5 months ago

correct

upvoted 4 times

roshansir 2 years, 5 months ago

correct answer

upvoted 3 times

paru123456789 2 years, 8 months ago

Answer:
Multi-step
Url ping

upvoted 5 times

10x 2 years, 8 months ago

b. and c.

upvoted 3 times

DRAG DROP -

A web service provides customer summary information for e-commerce partners. The web service is implemented as an Azure Function app with an HTTP trigger.

Access to the API is provided by an Azure API Management instance. The API Management instance is configured in consumption plan mode. All API calls are authenticated by using OAuth.

API calls must be cached. Customers must not be able to view cached data for other customers.

You need to configure API Management policies for caching.

How should you complete the policy statement?

Select and Place:

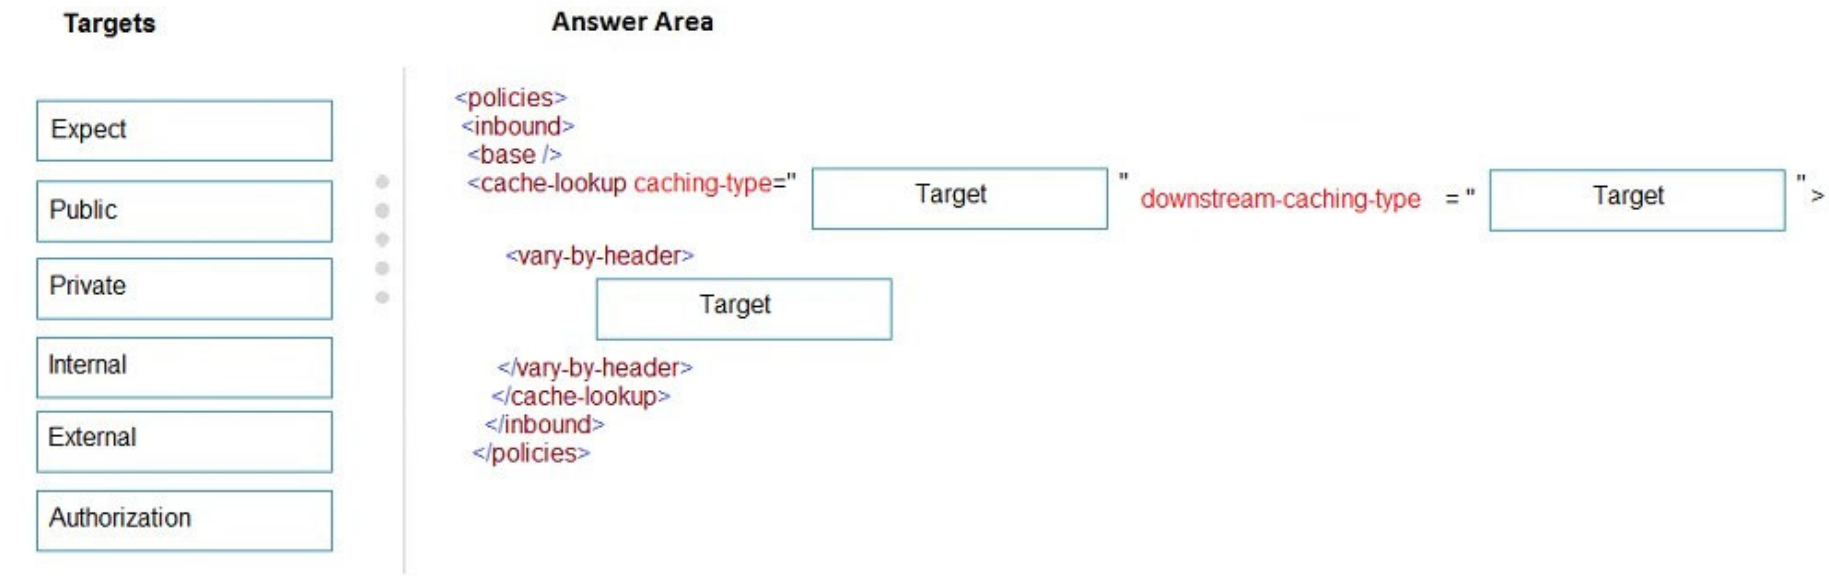

Correct Answer:

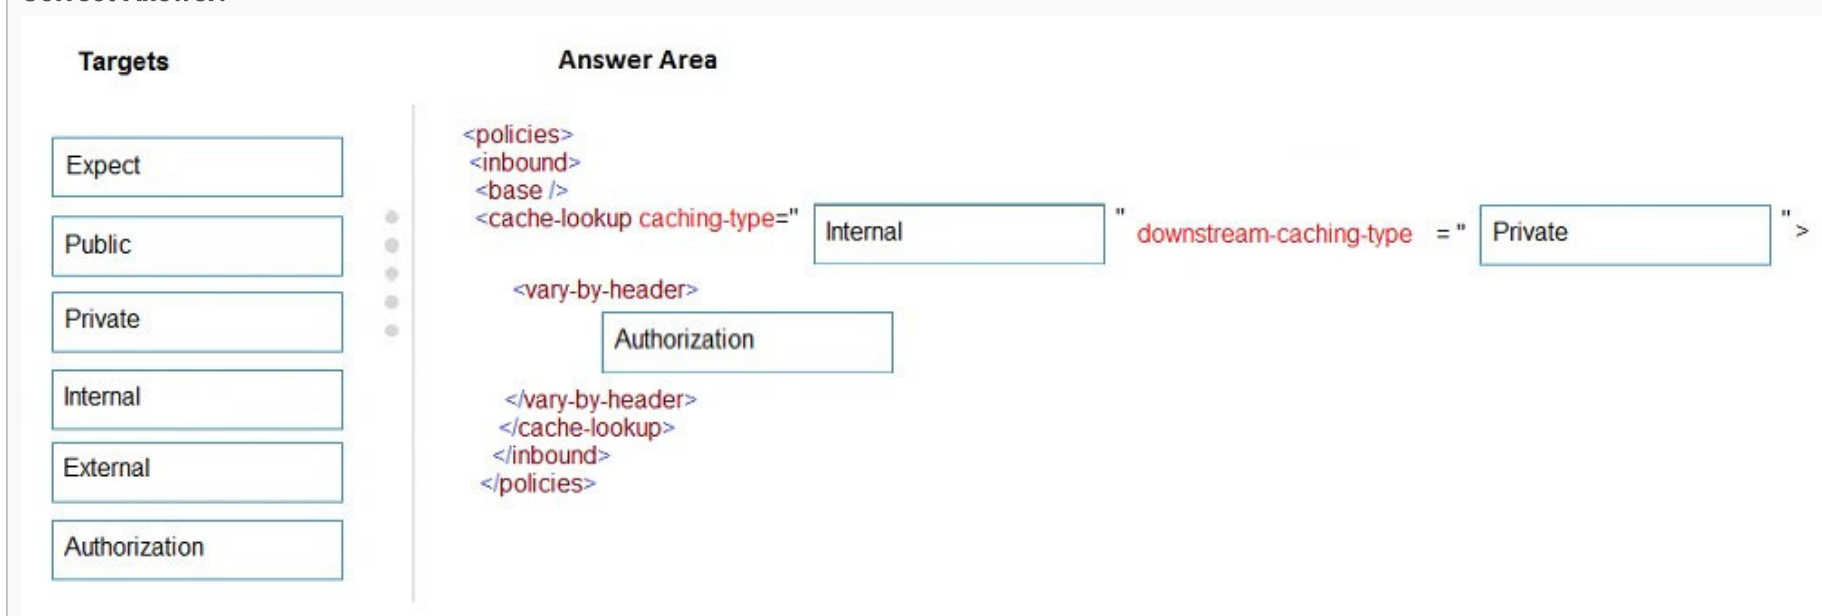

Box 1: internal -

caching-type

Choose between the following values of the attribute:

- internal to use the built-in API Management cache,
- external to use the external cache as Azure Cache for Redis prefer-external to use external cache if configured or internal cache otherwise.

Box 2: private -

downstream-caching-type

This attribute must be set to one of the following values.

- none - downstream caching is not allowed.
- private - downstream private caching is allowed.
- public - private and shared downstream caching is allowed.

Box 3: Authorization -

<vary-by-header>Authorization</vary-by-header>

<!-- should be present when allow-private-response-caching is "true"-->

Note: Start caching responses per value of specified header, such as Accept, Accept-Charset, Accept-Encoding, Accept-Language, Authorization, Expect, From,

Host, If-Match -

Reference:

https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies

cgreene Highly Voted 2 years, 11 months ago

Internal Cache isn't available for Consumption plan. So shouldn't it be External (using Redis)?

upvoted 88 times

Egger1992 2 years, 11 months ago

Yes, you are right. In the feature-based comparison of the tiers, it says that built-in cache is not available. https://docs.microsoft.com/en-us/azure/api-management/api-management-features
And here it says that "internal" would be using the built-in cache. https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies

upvoted 18 times

Ananth2020 Highly Voted 2 years, 11 months ago

Consumption Pricing Tier of API Management does not support Built-in Cache. Therefore caching-type should be set to "external".

Reference: https://docs.microsoft.com/en-us/azure/api-management/api-management-features

upvoted 35 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/07/2023

upvoted 4 times

uffuchsi 8 months, 1 week ago

Caching type must be external since consumption plan does not support internal caching.

Downstream caching type must be private as it is not shared.

Vary by header must be authorization as it is cached per customer.

https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies

upvoted 5 times

adilkhan 9 months ago

https://learn.microsoft.com/en-us/azure/api-management/cache-lookup-policy

upvoted 1 times

OPT_001122 12 months ago

External
Private
Authorization

upvoted 8 times

Eltooth 1 year, 4 months ago

External
Private
Authorization

https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies#attributes

https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#architecture

upvoted 10 times

Sandeep12093 1 year, 4 months ago

Yes, I agree consumption plan doesn't support internal caching..
Check microsoft pricing page for api management

upvoted 3 times

adone 1 year, 7 months ago

I don't find any example with private downstream and a vary by header. For me the vary by header only make sense for public downstream. I opt for:
External
Public
Authorization

In normal situation, I would have put private downstream without a vary by header.

upvoted 2 times

oescm 1 year, 8 months ago

Got this one 02/2022. Went with:

External

Private
Authorization

upvoted 16 times

leonidn 1 year, 9 months ago

Private is because cached data is isolated between customers
External is because internal is not a feature of the Consumption plan
Authorization to enable caching responses per value of specified header

upvoted 9 times

Gautam47 2 years, 1 month ago

Cache should be external

upvoted 3 times

chandru1dev 2 years, 3 months ago

internal
caching-type Choose between the following values of the attribute:
- internal to use the built-in API Management cache,
- external to use the external cache as described in Use an external Azure Cache for Redis in Azure API Management,
- prefer-external to use external cache if configured or internal cache otherwise.
https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies

upvoted 2 times

kondapaturi 2 years, 4 months ago

Extermal - When you look at the pricing plans for the Azure API management, when it comes to the Consumption plan, it only supports external cache.
Private - we need to ensure that the data is not available to other customers, we need to use the cache type of Private caching.
Authorization - the calls are being authenticated by OAuth, the caching can vary by the Authorization header.

upvoted 13 times

glam 2 years, 5 months ago

Answers are:
External
Private
Authorization

upvoted 17 times

anandhprakash 2 years, 5 months ago

Internal cache is not supported in consumption plan, so the cache type should be External.
Answers are:
External
Private
Authorization

upvoted 6 times

roshansir 2 years, 5 months ago

External
Private
Authorization

upvoted 1 times

Question #15 *Topic 5*

You are developing applications for a company. You plan to host the applications on Azure App Services.

The company has the following requirements:

☞ Every five minutes verify that the websites are responsive.

☞ Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly.

☞ Generate alerts if a website is experiencing issues.

☞ If a website fails to load, the system must attempt to reload the site three more times.

You need to implement this process with the least amount of effort.

What should you do?

A. Create a Selenium web test and configure it to run from your workstation as a scheduled task.

B. Set up a URL ping test to query the home page.

C. Create an Azure function to query the home page.

D. Create a multi-step web test to query the home page.

E. Create a Custom Track Availability Test to query the home page.

Correct Answer: *D*

You can monitor a recorded sequence of URLs and interactions with a website via multi-step web tests.

Incorrect Answers:

A: Selenium is an umbrella project for a range of tools and libraries that enable and support the automation of web browsers.

It provides extensions to emulate user interaction with browsers, a distribution server for scaling browser allocation, and the infrastructure for implementations of the W3C WebDriver specification that lets you write interchangeable code for all major web browsers.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-multistep

Community vote distribution

B (74%) E (15%) 6%

kishe Highly Voted 2 years, 5 months ago

IMO, B is correct, a URL ping test is enough.

upvoted 66 times

mattvasc 1 year, 8 months ago

And how a URL ping will meet the requirement:
"☞ Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly." ?

Definitly not B.

upvoted 8 times

uffuchsi 8 months, 1 week ago

URL ping test validate whether an endpoint is responding. It also measures the performance associated with that response. You can also set custom success criteria, coupled with more advanced features like parsing dependent requests and allowing for retries.

https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability

upvoted 3 times

ezelans 1 year, 7 months ago

please, read this article and you will check it. There is a setting called "Test timeout"

https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability

upvoted 10 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000.

upvoted 5 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose：B

upvoted 3 times

AymanAkk 10 months, 1 week ago

Thanks for mentioning the date. were there a lot of questions from exam topics ?
upvoted 3 times

jay158 Highly Voted 2 years, 4 months ago
The Given Answer D is correct.
Ping URL, does not satisfies : Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly
upvoted 25 times

jay158 2 years, 4 months ago
Second thought
Correct answer is B. ie URL Ping Test
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability
upvoted 20 times

anirbanzeus 2 years, 4 months ago
Not really, all the objectives can be achieved through Ping-test. Multi-step test in a complex solution requiring more effort and dives into much deeper details that requested in the question.
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability#create-a-test : All the options listed here map1-1 with the question requirements.
upvoted 11 times

katrang Most Recent 1 month, 2 weeks ago
Selected Answer: B
From the documentation:

The name URL ping test is a bit of a misnomer. These tests don't use the Internet Control Message Protocol (ICMP) to check your site's availability. Instead, they use more advanced HTTP request functionality to validate whether an endpoint is responding. They measure the performance associated with that response. They also add the ability to set custom success criteria, coupled with more advanced features like parsing dependent requests and allowing for retries.
upvoted 1 times

Paolo12 2 months ago
Selected Answer: B
In URL-Ping test it's possible to set dependent requests, retries, etc...
upvoted 1 times

aragones 5 months, 2 weeks ago
Got this 2023-05-12.

FYI:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada
upvoted 3 times

motekim 7 months, 1 week ago
It looks like its B though: https://learn.microsoft.com/en-us/previous-versions/azure/azure-monitor/app/monitor-web-app-availability
upvoted 1 times

motekim 7 months, 1 week ago
if the answer is D. Create a multi-step web test to query the home page., then it's actually E. from Msoft site: We recommend using TrackAvailability to submit custom availability tests instead of multistep web tests. This option is the long-term supported solution for multi-request or authentication test scenarios. With TrackAvailability() and custom availability tests, you can run tests on any compute you want and use C# to easily author new tests. hmm
upvoted 1 times

NombreFalso 8 months, 1 week ago
Selected Answer: D
Ping test seems to be the least effort solution
upvoted 1 times

Sriniv 8 months, 2 weeks ago
Answer B as it satisfies all requirements.Below is the reason for URL ping test.
The Azure App Service has built-in URL ping testing that can be configured to check whether a website is responsive every five minutes, and generate alerts if there are issues. The URL ping test can also check that the website responds within a specified time threshold and that dependent requests such as images and JavaScript files load properly. Additionally, the URL ping test will automatically retry up to three times if the website fails to load.
upvoted 2 times

adilkhan 9 months ago
B is correct because:
1. The test requests images, scripts, style files, and other files that are part of the webpage under test. The recorded response time includes the time taken to get these files. The test fails if any of these resources can't be successfully downloaded within the timeout for the whole test. If the option isn't enabled, the test only requests the file at the URL that you specified. Enabling this option results in a stricter check. The test might fail for cases that aren't noticeable from manually browsing through the site.
2. When the test fails, it's retried after a short interval. A failure is reported only if three successive attempts fail. Subsequent tests are then

performed at the usual test frequency. Retry is temporarily suspended until the next success. This rule is applied independently at each test location. We recommend this option. On average, about 80 percent of failures disappear on retry.
https://learn.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability
upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: B

B. Set up a URL ping test to query the home page.
upvoted 2 times

Dhanishetty 11 months, 2 weeks ago

Selected Answer: B

https://learn.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability#:~:text=Setting,to%2016%20locations.
upvoted 1 times

OPT_001122 12 months ago

Selected Answer: B

https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability

Test timeout - Decrease this value to be alerted about slow responses. The test is counted as a failure if the responses from your site haven't been received within this period. If you selected Parse dependent requests, all the images, style files, scripts, and other dependent resources must have been received within this period.
upvoted 1 times

nvtienanh 1 year ago

Selected Answer: E

Multi-step web tests have been deprecated. We recommend using TrackAvailability() to submit custom availability tests instead of multi-step web tests.
I think E is correct.
upvoted 2 times

azurepaul 1 year, 1 month ago

I vote e based on the dependent files bit - and as called out in other comments:

Multi-step web tests have been deprecated. We recommend using TrackAvailability() to submit custom availability tests instead of multi-step web tests. With TrackAvailability() and custom availability tests, you can run tests on any compute you want and use C# to easily author new tests.
upvoted 3 times

Knightie 1 year, 2 months ago

Selected Answer: A

Ping Test is not enough to ensure the content loaded time, ping will get the 200 response and the image and other stuff will carry on until completion.

Selenium can check if the loaded content is missing or not but it needs much efforts choose this answer only if we have no other better choice.

MultiStep content match only support String (Plain Text), not able to check if images are loaded properly.

So in the end, Selenium.
upvoted 2 times

renzoku 1 year, 4 months ago

Selected Answer: D

It mention websites (pluras) then how can you use single URL ping for a amount of endpoints, maybe it could be more complex to do than using multi-step web test or custom track availability
upvoted 1 times

Question #16 *Topic 5*

You develop and add several functions to an Azure Function app that uses the latest runtime host. The functions contain several REST API endpoints secured by using SSL. The Azure Function app runs in a Consumption plan.
You must send an alert when any of the function endpoints are unavailable or responding too slowly.
You need to monitor the availability and responsiveness of the functions.
What should you do?

A. Create a URL ping test.

B. Create a timer triggered function that calls TrackAvailability() and send the results to Application Insights.

C. Create a timer triggered function that calls GetMetric("Request Size") and send the results to Application Insights.

D. Add a new diagnostic setting to the Azure Function app. Enable the FunctionAppLogs and Send to Log Analytics options.

Correct Answer: *B*
You can create an Azure Function with TrackAvailability() that will run periodically according to the configuration given in TimerTrigger function with your own business logic. The results of this test will be sent to your Application Insights resource, where you will be able to query for and alert on the availability results data.
This allows you to create customized tests similar to what you can do via Availability Monitoring in the portal. Customized tests will allow you to write more complex availability tests than is possible using the portal UI, monitor an app inside of your Azure VNET, change the endpoint address, or create an availability test even if this feature is not available in your region.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-azure-functions

Community vote distribution
B (77%) A (23%)

mlantonis Highly Voted 2 years, 5 months ago
Based on the doc, I believe B is correct.

https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-azure-functions
upvoted 26 times

qzrkfrcglkzzvfrcro Highly Voted 2 years, 3 months ago
I think the answer should be URL Ping Test.
It can send an alert when the website is unavailable or unresponsive. See test timeouts and alerts in the link.
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability
upvoted 20 times

ks321 3 months, 1 week ago
Note that the API is protected by SSL. Does URL ping test support auth?
upvoted 1 times

Knightie 1 year, 2 months ago
Ping Test only check for response code 200, it will not give you any info after the initial 200 ok code, contents and information will keep loading after the 200 code recieved, so if the 200 there but the page not responsing or the background not processing or not completing, then the ping test will not be able to know...it will still be 200 but the response might be incomplete.
upvoted 2 times

Knightie 1 year, 2 months ago
"responsiveness of the functions"<=== this one will need more than the ping test
upvoted 3 times

troy89 1 year, 8 months ago
But you need a URL test per endpoint, so one URL test wouldn't be enough.
upvoted 2 times

uffuchsi Most Recent 8 months, 1 week ago
It has to be C - The URL ping test uses advanced HTTP request functionality to validate whether an endpoint is responding. They measure the performance associated with that response. They also add the ability to set custom success criteria, coupled with more advanced features like parsing dependent requests and allowing for retries.

https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability#actions-if-you-see-failures
upvoted 2 times

uffuchsi 8 months, 1 week ago

Apologies - I meant A, not C
upvoted 2 times

Sriniv 8 months, 2 weeks ago

Adding a new diagnostic setting (Option D) can enable logging for the Azure Function app, but it will not provide the ability to monitor the availability and responsiveness of the REST API endpoints or generate an alert when they are unavailable or responding too slowly.
upvoted 1 times

Sriniv 8 months, 2 weeks ago

Creating a timer triggered function (Option B or C) can also be used to call the REST API endpoints and report the results to Application Insights. However, this approach requires additional development effort and maintenance overhead, which is not necessary for this scenario. Moreover, it will not provide an immediate alert when the function endpoints are unavailable or responding too slowly.
upvoted 1 times

Sriniv 8 months, 2 weeks ago

the best option is to create a URL ping test to monitor the availability and responsiveness of the REST API endpoints. This can be easily configured in Azure Application Insights with minimal effort and can provide immediate alerts when the endpoints are unavailable or responding too slowly.
upvoted 1 times

Dhanishetty 11 months, 2 weeks ago

Selected Answer: B

its B.....
upvoted 1 times

OPT_001122 12 months ago

Selected Answer: B

B is correct
upvoted 1 times

AZAdam22 1 year, 3 months ago

Selected Answer: A

A - Because the URL ping test automatically measures availability and response times.
upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.
https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-azure-functions
upvoted 3 times

yohannvf 1 year, 5 months ago

Selected Answer: B

"You need to monitor the availability and responsiveness of the functions." You want it to be on application insights
upvoted 2 times

yohannvf 1 year, 5 months ago

"You need to monitor the availability and responsiveness of the functions." You want it to be on application insights
upvoted 1 times

iamstudying 1 year, 7 months ago

Selected Answer: A

I would go with A - URL Ping satisfy all the requirements. It does not check SSL validity BUT there is no requirement to do so, buddies

https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-standard-tests
upvoted 2 times

adone 1 year, 8 months ago

This scenario does not make any sense.
Consumption plans are not intended for high availability because the app may scale to zero when idle which may will leads to cold start of some minute(s) so it makes no sense to monitor it for availability. With URL Ping, you will pay for the compute resources generated by your monitoring request to your URL Ping which makes this option ridiculous with consumption plan. But monitoring your application with an URL Ping might keep your function app up and running if this is what you expect :D

The option B is the less worse but keep in mind that consumption plan are not intended for high availability (does not need to be monitored). They are intended for asynchronous executions, even high workloads and parallel executions, which does not require to be highly available. For high availability, choose Premium/Dedicated plans.
upvoted 9 times

leonidn 1 year, 9 months ago

Selected Answer: B

I guess an SSL validity check is required. URL ping does not provide this feature. There is no Standard test option. So then, option B is the only one that is viable.
upvoted 3 times

ScubaDiver123456 1 year, 9 months ago

I'm confused. I dont see a requirement to test SSL validity. It only says the API is protected by SSL.

upvoted 2 times

HiteshRamnani 1 year, 10 months ago

B is correct

upvoted 1 times

avenger34 1 year, 11 months ago

"Standard tests are a single request test that is similar to the URL ping test but more advanced. In addition to validating whether an endpoint is responding and measuring the performance, Standard tests also includes SSL certificate validity, proactive lifetime check, HTTP request verb (for example GET,HEAD,POST, etc.), custom headers, and custom data associated with your HTTP request."

URL Ping doesn't support checking SSL certificate

https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability

upvoted 5 times

avenger34 1 year, 11 months ago

The ref url:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-standard-tests

upvoted 2 times

Question #17 *Topic 5*

DRAG DROP -

You are developing an application to retrieve user profile information. The application will use the Microsoft Graph SDK.

The app must retrieve user profile information by using a Microsoft Graph API call.

You need to call the Microsoft Graph API from the application.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

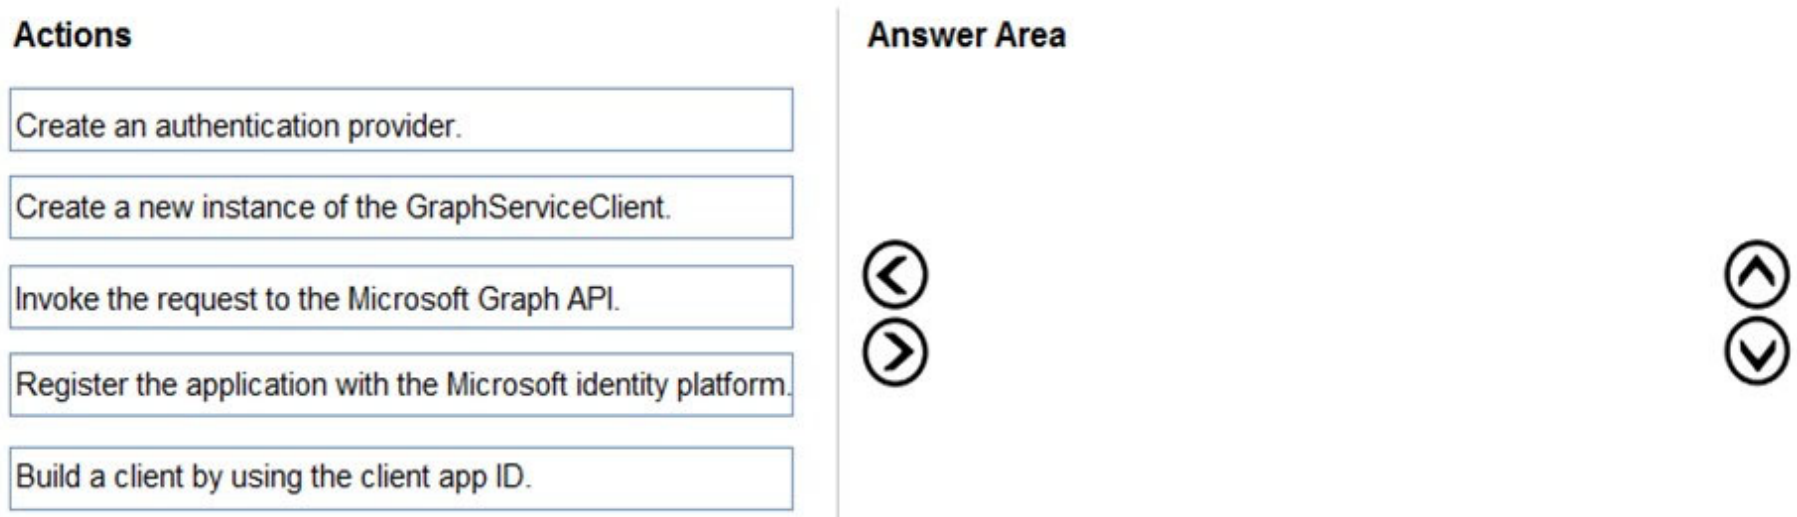

Correct Answer:

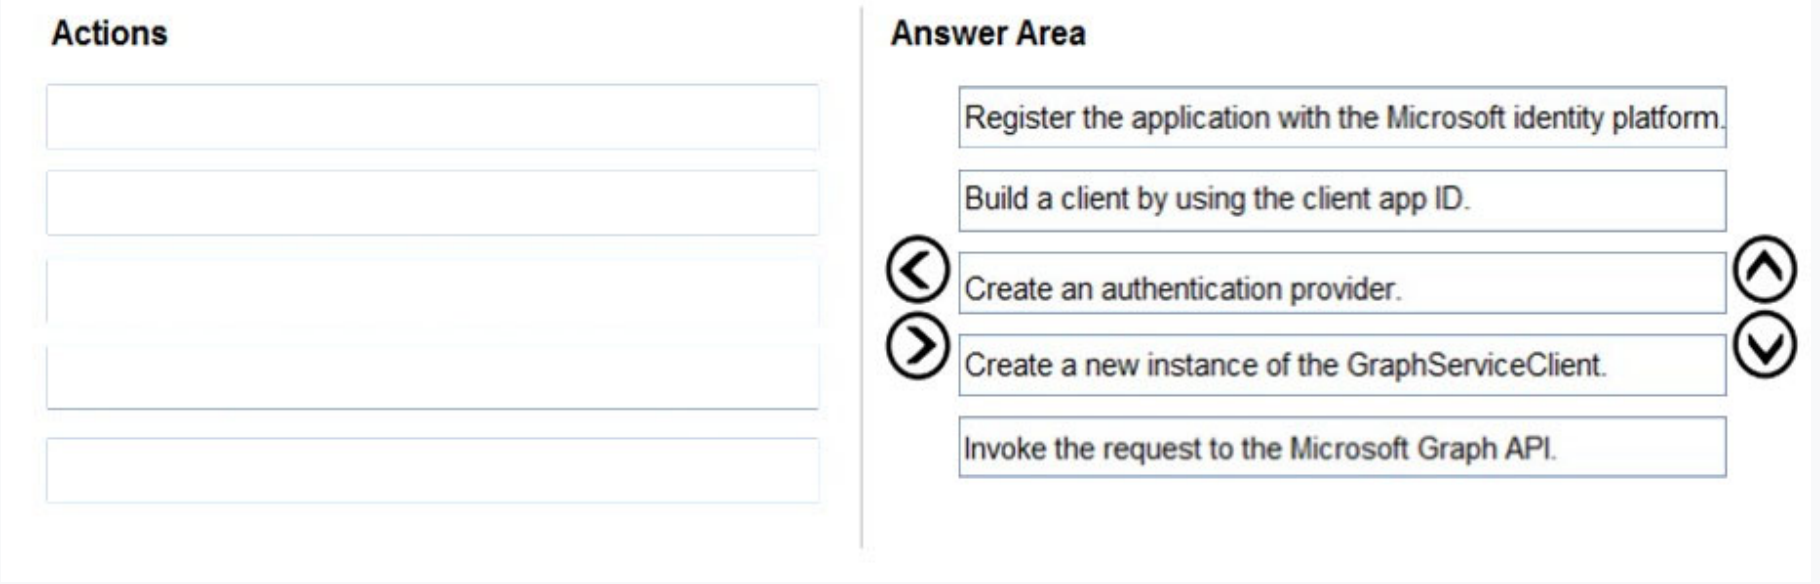

Step 1: Register the application with the Microsoft identity platform.

To authenticate with the Microsoft identity platform endpoint, you must first register your app at the Azure app registration portal

Step 2: Build a client by using the client app ID

Step 3: Create an authentication provider

Create an authentication provider by passing in a client application and graph scopes.

Code example:

```
DeviceCodeProvider authProvider = new DeviceCodeProvider(publicClientApplication, graphScopes);
// Create a new instance of GraphServiceClient with the authentication provider.
GraphServiceClient graphClient = new GraphServiceClient(authProvider);
```

Step 4: Create a new instance of the GraphServiceClient

Step 5: Invoke the request to the Microsoft Graph API

Reference:

https://docs.microsoft.com/en-us/graph/auth-v2-service

https://docs.microsoft.com/en-us/graph/sdks/create-client

kishe Highly Voted 2 years, 5 months ago

Answer is correct

upvoted 43 times

warchoon 7 months, 4 weeks ago

Answer is correct, but the question is not.
3,4 are the parts of 2 not actually steps

upvoted 2 times

BishopeL Highly Voted 2 years, 2 months ago

I got this in my exam this morning. The provided answer is correct.

upvoted 32 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 4 times

warchoon 7 months, 4 weeks ago

Check out the JavaScript example https://learn.microsoft.com/en-us/graph/sdks/create-client?tabs=Javascript
The C# sample does not clearly describe the authentication provider.

upvoted 1 times

OPT_001122 10 months, 3 weeks ago

The answer is correct - RBCCI

upvoted 9 times

NombreFalso 8 months, 1 week ago

Hahaha, that's so classic, creating acronyms for memorizing stuff like this.

upvoted 4 times

AAcademy 10 months, 2 weeks ago

I lov that, thanks

upvoted 3 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 7 times

OPT_001122 11 months, 1 week ago

Thanks for mentioning the date

upvoted 1 times

Question #18 *Topic 5*

DRAG DROP -

You develop and deploy an Azure Logic App that calls an Azure Function app. The Azure Function App includes an OpenAPI (Swagger) definition and uses an

Azure Blob storage account. All resources are secured by using Azure Active Directory (Azure AD).

The Logic App must use Azure Monitor logs to record and store information about runtime data and events. The logs must be stored in the Azure Blob storage account.

You need to set up Azure Monitor logs and collect diagnostics data for the Azure Logic App.

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:

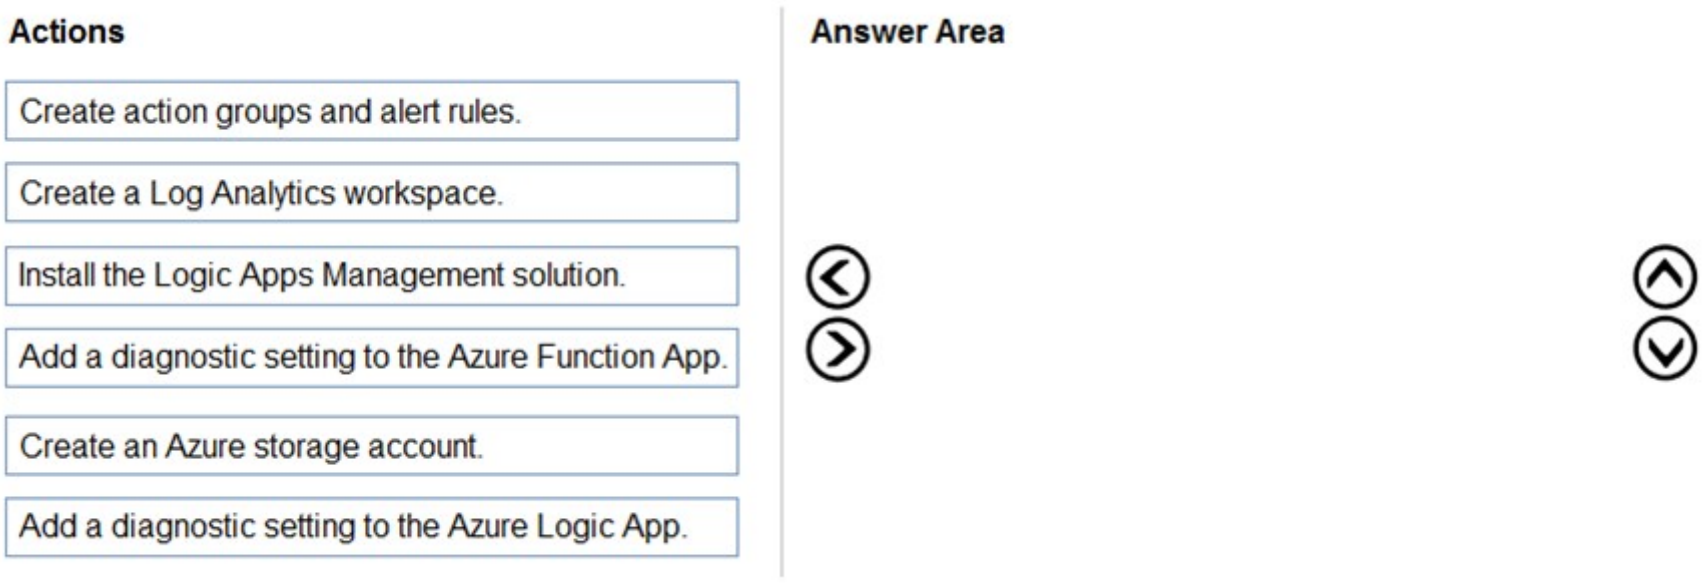

Correct Answer:

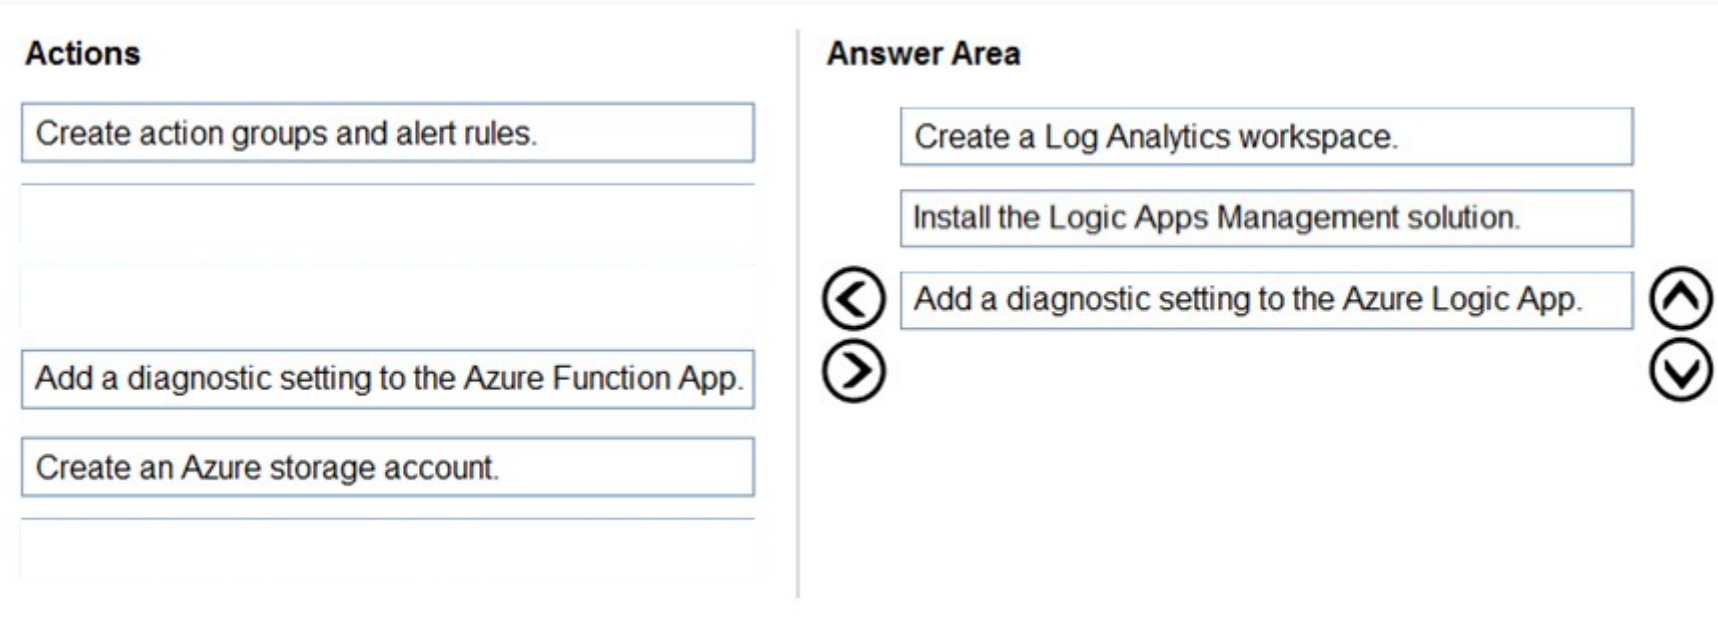

Step 1: Create a Log Analytics workspace

Before you start, you need a Log Analytics workspace.

Step 2: Install the Logic Apps Management solution

To set up logging for your logic app, you can enable Log Analytics when you create your logic app, or you can install the Logic Apps Management solution in your

Log Analytics workspace for existing logic apps.

Step 3: Add a diagnostic setting to the Azure Logic App

Set up Azure Monitor logs -

1. In the Azure portal, find and select your logic app.
2. On your logic app menu, under Monitoring, select Diagnostic settings > Add diagnostic setting.

Reference:

https://docs.microsoft.com/en-us/azure/logic-apps/monitor-logic-apps-log-analytics

jadda Highly Voted 2 years, 4 months ago

After the March 26 changes, Logic apps are out of scope for the AZ-204 exam

upvoted 31 times

hawgie 1 month, 1 week ago

Got this one today. So much for being out of scope. 2023-09-19

upvoted 4 times

dhishkiyaau 2 years, 2 months ago

Can you refer the source?

upvoted 2 times

mandusya 1 year, 10 months ago

confirmed
https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B

upvoted 3 times

HiteshRamnani 1 year, 10 months ago

thank you

upvoted 1 times

jay158 Highly Voted 2 years, 3 months ago

Answer is correct
https://docs.microsoft.com/en-us/azure/logic-apps/monitor-logic-apps-log-analytics

upvoted 12 times

JH81 Most Recent 4 months ago

I got this on 6/29/2023 and passed with 850. Answer is correct.

upvoted 8 times

OPT_001122 10 months, 3 weeks ago

The given answer is correct.

CIA

upvoted 4 times

Knightie 1 year, 2 months ago

eliminate action and alert.
eliminate storage account creaton because the question already has an account.
eliminate diagnotics on function app, you need to put it on the logic app instead.

Then arrange the rest to get the answer.

upvoted 4 times

learner06 2 years, 1 month ago

dont we need to create Azure storage account in order to store logs in blob?

upvoted 2 times

xahah22222 2 years ago

Such an account is already there

upvoted 2 times

Jurgen1234 2 years, 4 months ago

First one should Azure storage account as you need to store them in a blob storage

upvoted 1 times

Jurgen1234 2 years, 4 months ago

Nevermind you already have a storage account

upvoted 1 times

markra 2 years, 5 months ago

The given answer is correct. See the reference provided in a solution.

upvoted 4 times

Question #19 *Topic 5*

DRAG DROP -

You develop an application. You plan to host the application on a set of virtual machines (VMs) in Azure.

You need to configure Azure Monitor to collect logs from the application.

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

Select and Place:

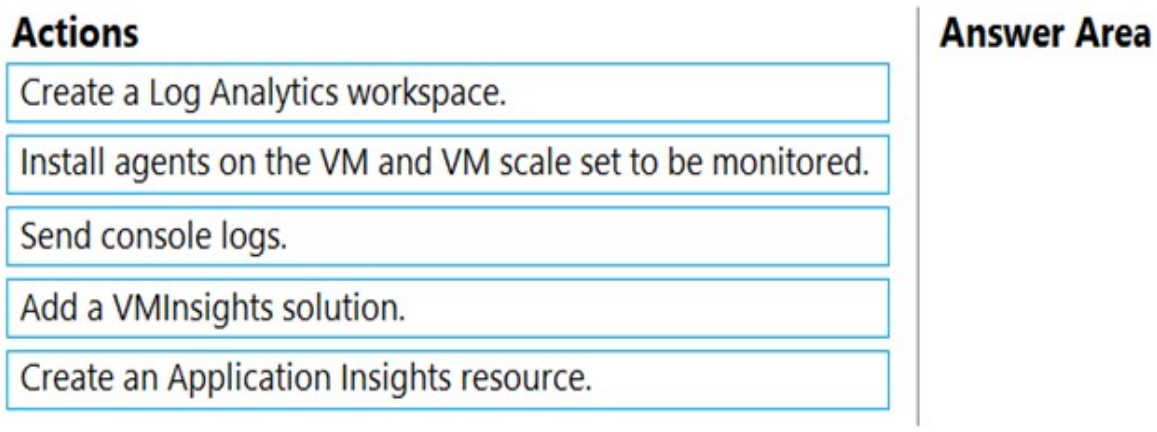

Correct Answer:

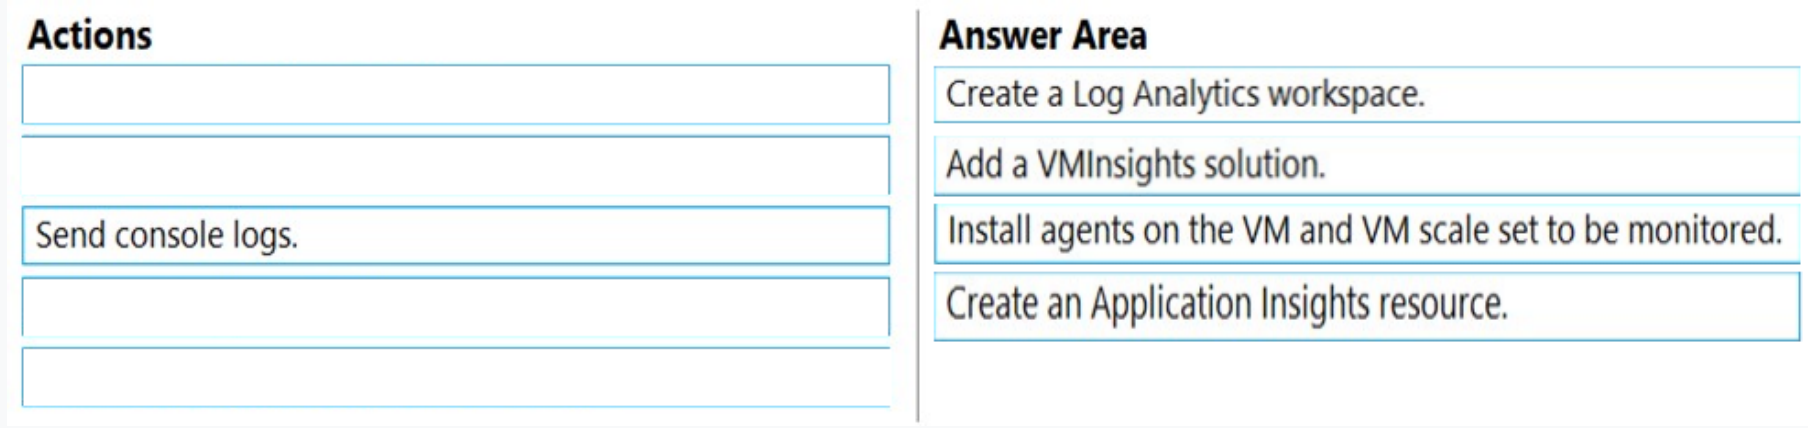

Step 1: Create a Log Analytics workspace.

First create the workspace.

Step 2: Add a VMInsights solution.

Before a Log Analytics workspace can be used with VM insights, it must have the VMInsights solution installed.

Step 3: Install agents on the VM and VM scale set to be monitored.

Prior to onboarding agents, you must create and configure a workspace. Install or update the Application Insights Agent as an extension for Azure virtual machines and VM scalet sets.

Step 4: Create an Application Insights resource

Sign in to the Azure portal, and create an Application Insights resource.

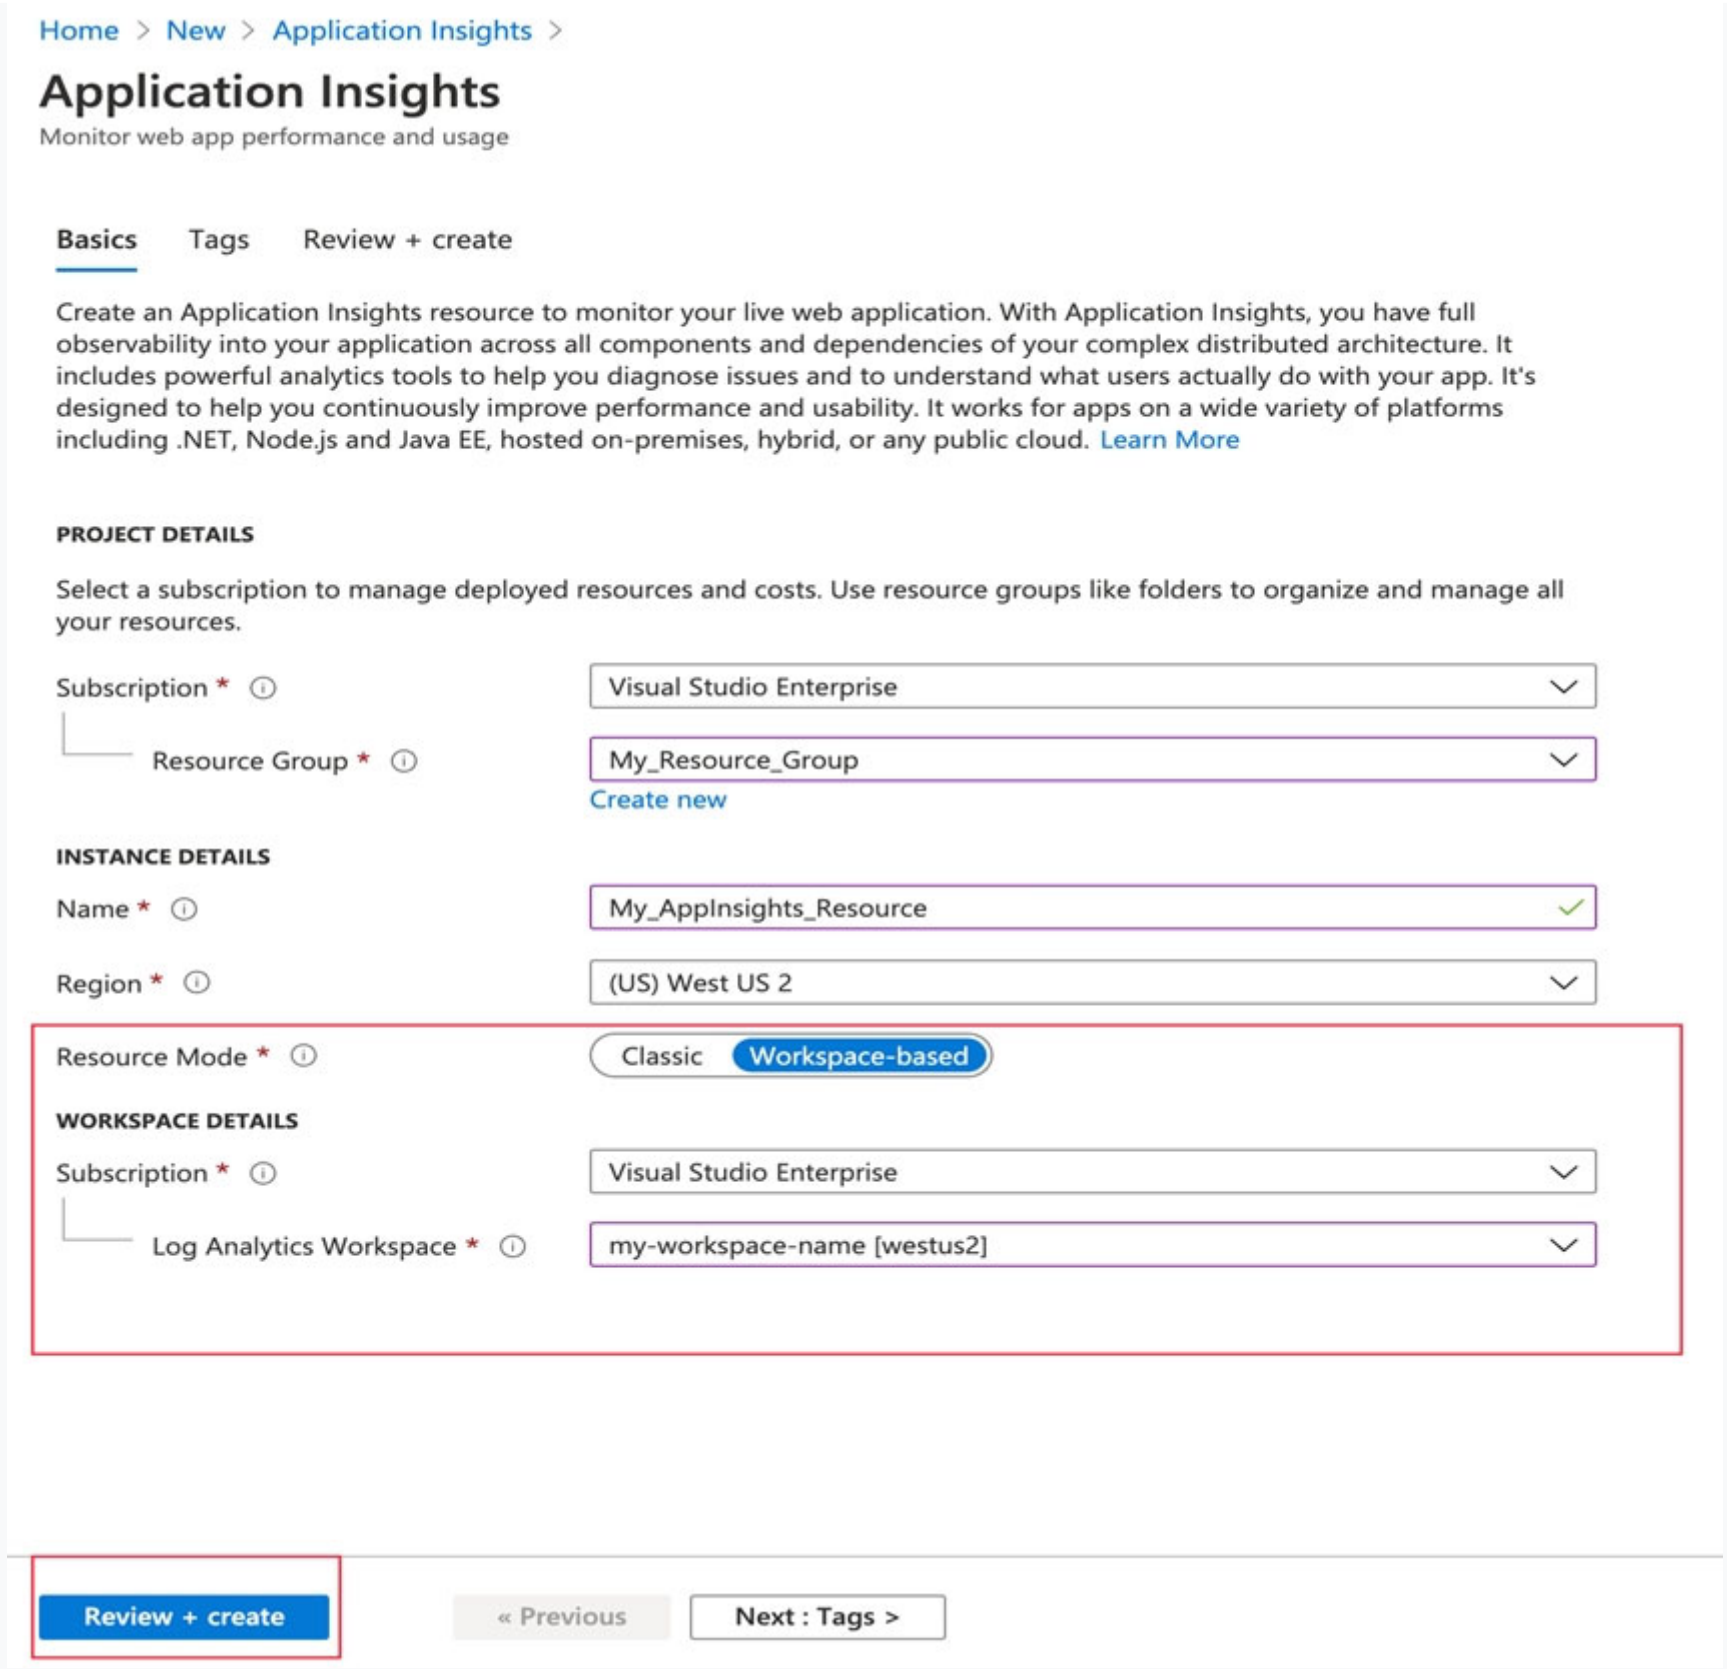

Once a workspace-based Application Insights resource has been created, configuring monitoring is relatively straightforward.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/vm/vminsights-configure-workspace https://docs.microsoft.com/en-us/azure/azure-monitor/app/create-workspace-resource

aradice Highly Voted 2 years, 4 months ago
correct https://docs.microsoft.com/en-us/azure/azure-monitor/vm/vminsights-configure-workspace?tabs=CLI
https://docs.microsoft.com/en-us/azure/azure-monitor/app/create-workspace-resource
upvoted 23 times

vilainchien 1 year, 8 months ago
correct https://docs.microsoft.com/en-us/azure/azure-monitor/vm/vminsights-overview#configuring-vm-insights
upvoted 2 times

rolling_potato_ Highly Voted 10 months ago
Was on my exam today (3rd of Jan 2023)
upvoted 7 times

adilkhan 9 months ago
what was your answer and your score please?
upvoted 1 times

WH16 Most Recent 2 months, 3 weeks ago
This was on the exam (9th August 2023). Went with proposed. Scored 910.
upvoted 2 times

coffecold 1 year ago
You must make the workspace in the portal
In Monitor you must add Application insights with the workspace
In Monitor you must activate VMInsights and add your VM's and assign a workspace to it
Then install and run the agents in each VM

1. Create Log Analytics workspace (in Azure Portal)
2. Create Application Insights resource (in Monitor, Application Insights with workspace attached)
3.Add VMInsights solution (== activate VMInsights in Monitor, choose VM's, attach workspace).
4.Install agents on VM

upvoted 5 times

wsguruprasath 1 year, 7 months ago

Create Log Analytics workspace.
Add VMInsights solution to workspace.
Install agents on virtual machine and virtual machine scale set to be monitored.

https://docs.microsoft.com/en-us/azure/azure-monitor/vm/vminsights-overview

upvoted 1 times

RajMasilamani 2 years, 1 month ago

Configuring VM insights
The steps to configure VM insights are as follows. Follow each link for detailed guidance on each step:

1.Create Log Analytics workspace.
2.Add VMInsights solution to workspace.
3.Install agents on virtual machine and virtual machine scale set to be monitored.

upvoted 5 times

ning 2 years, 2 months ago

Why application insights???

upvoted 1 times

nyanko 1 year, 8 months ago

Is VMInsights required to monitor application log ?
https://docs.microsoft.com/en-us/azure/azure-monitor/app/azure-vm-vmss-apps

upvoted 1 times

ning 2 years, 2 months ago

Add more stuff ...

If you want to monitor VM, CPU, memory, disk, you should us VM insights, if you want to monitor an app running inside VM, you need application insights ...

Not sure what is the question asking ...

upvoted 3 times

jvyas 2 years, 1 month ago

It does say "You need to configure Azure Monitor to collect logs from the application." Hope this helps to discern.

upvoted 1 times

ning 2 years, 2 months ago

I think this is the answer
1. Create log name space
2. Create app insights
3. Install app insights agents on VM
4. Send console log for the app

upvoted 4 times

Maharaja 2 years, 1 month ago

Thank you. Can you provide some ms doc link supporting your answer?

upvoted 2 times

adilkhan 9 months ago

why 2 is application insights?

upvoted 1 times

ning 2 years, 1 month ago

After lots of research ...

1. Log analytics name space
2. Application insights
3. VM solutions
4. Install agents

upvoted 5 times

adilkhan 9 months ago

why 2 is application insights?

upvoted 1 times

finnishr 2 years, 2 months ago

100% correct https://docs.microsoft.com/en-us/azure/azure-monitor/vm/vminsights-overview

upvoted 7 times

Question #20 *Topic 5*

You develop and deploy an Azure App Service web app. The app is deployed to multiple regions and uses Azure Traffic Manager. Application Insights is enabled for the app.
You need to analyze app uptime for each month.
Which two solutions will achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Azure Monitor logs

B. Application Insights alerts

C. Azure Monitor metrics

D. Application Insights web tests

Correct Answer: *BD*
Reference:
https://azure.microsoft.com/en-us/blog/creating-a-web-test-alert-programmatically-with-application-insights/

Community vote distribution
AC (55%) BC (30%)

coffecold Highly Voted 1 year ago
A and C. Metrics will give you the uptime. Logs will give you the causes of the downtime.
wrong B: Alerts are not required
wrong D: Web tests has nothing to do with uptime.
upvoted 17 times

qwerty112233 Highly Voted 1 year, 1 month ago
"You need to analyze app uptime for each month." how web tests and alerts will help you to analyze previous data?
upvoted 9 times

edmond 10 months, 2 weeks ago
Certainly web tests don't help, but for alerts, you can get use of the availability alerts, refer to https://learn.microsoft.com/en-us/azure/azure-monitor/app/availability-alerts.
I don't think the Metrics logs work here, hence BC makes more sense.
upvoted 2 times

edmond 10 months, 2 weeks ago
Sorry I meant the Monitor logs don't work in this case, the Monitor metrics and Application Insight alerts make more sense here.
upvoted 2 times

micro9000 11 months, 3 weeks ago
We are analyzing the app uptime not the resources or errors that our azure resources or the application generates, if you go to your App Insight instance -> Availability then click the SLA Report you will see the uptime report of your application there. That is the SLA report of your application.
upvoted 2 times

EliteAllen Most Recent 3 months ago
Selected Answer: AC
A. Azure Monitor logs: By querying logs, you can analyze detailed information about app uptime and other metrics over a specified time period, such as a month.

C. Azure Monitor metrics: Metrics provide numerical data about your system's operation, including uptime. You can use Azure Monitor metrics to track and analyze the availability and responsiveness of your app over time, including on a monthly basis.
upvoted 3 times

motekim 7 months, 1 week ago
I like A, D why?
A: https://www.azure.cn/en-us/support/sla/application-insights/index.html
d: https://learn.microsoft.com/en-us/azure/azure-monitor/app/availability-overview
upvoted 1 times

Sriniv 8 months, 2 weeks ago
Definitely not A &B due to following reason
Option A, Azure Monitor logs, is not the best solution for analyzing app uptime. While Azure Monitor logs can provide information on app usage and performance, they may not provide the necessary information to calculate the app uptime percentage for each month.

Option B, Application Insights alerts, can be useful for notifying you of any issues with the app, but they may not provide the necessary information to calculate the app uptime percentage for each month.

upvoted 1 times

Osip21 9 months ago

Selected Answer: CD

ChatGPT:
C. Azure Monitor metrics
D. Application Insights web tests

Azure Monitor metrics can be used to track the availability of the app by analyzing the uptime percentage. Application Insights web tests can be configured to periodically check the availability of the app and provide uptime data.

upvoted 3 times

alexein74 9 months, 1 week ago

Selected Answer: AC

A. Azure Monitor logs
C. Azure Monitor metrics

Azure Monitor logs: Azure Monitor logs provides detailed diagnostic and operational information from Azure resources and allows you to store and analyze the logs. By querying Azure Monitor logs, you can analyze app uptime for each month by searching for uptime and downtime events.

Azure Monitor metrics: Azure Monitor metrics allow you to monitor the performance and health of your Azure resources. With Azure Monitor metrics, you can analyze app uptime by looking at the percentage of availability for the app, and you can set up alerts if uptime falls below a certain threshold.

upvoted 2 times

rotimislaw 10 months ago

Selected Answer: BC

B. Azure Monitor metrics - can be set to count uptime
C. Application Insights alerts - flag when servers go down

upvoted 3 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose：A、C

upvoted 4 times

bertwin 11 months ago

Selected Answer: AC

A & C i would say
https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/metrics-supported

upvoted 2 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 3 times

OPT_001122 11 months, 1 week ago

correct ans?

upvoted 1 times

somename20221106 11 months, 2 weeks ago

What is the correct answer ?

upvoted 3 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: AC

A and C.

upvoted 3 times

Yumi21 11 months, 3 weeks ago

Selected Answer: AC

Shouldn't it be Monitor logs instead of alerts?

upvoted 1 times

g2000 9 months, 2 weeks ago

https://learn.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability#alerts

upvoted 1 times

sof123 1 year ago

BC - Azure monitor metrics and Application insights alerts

upvoted 1 times

Chief2110 1 year, 1 month ago

Selected Answer: BC

I think this should be it. BC

upvoted 3 times

erreyesarroyo 1 year, 1 month ago

????? The reference did not even mention it...

upvoted 1 times

Question #21 *Topic 5*

DRAG DROP -

You develop and deploy an Azure App Service web app. The web app accesses data in an Azure SQL database.

You must update the web app to store frequently used data in a new Azure Cache for Redis Premium instance.

You need to implement the Azure Cache for Redis features.

Which feature should you implement? To answer, drag the appropriate feature 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:

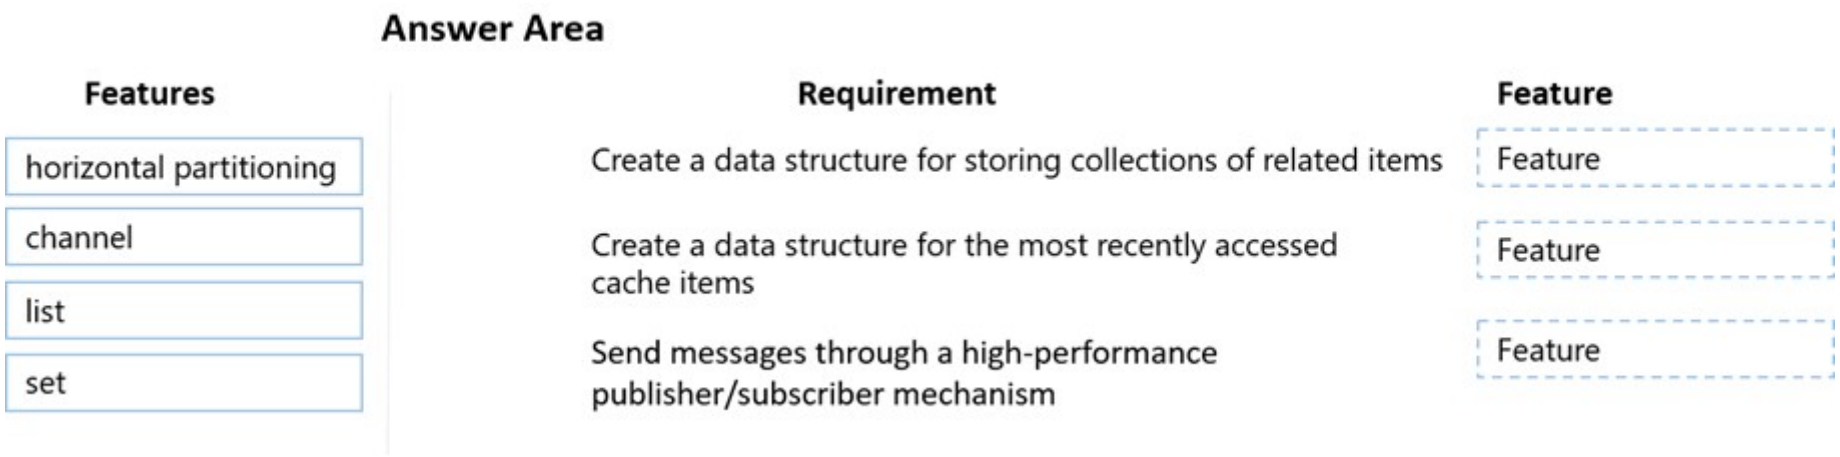

Correct Answer:

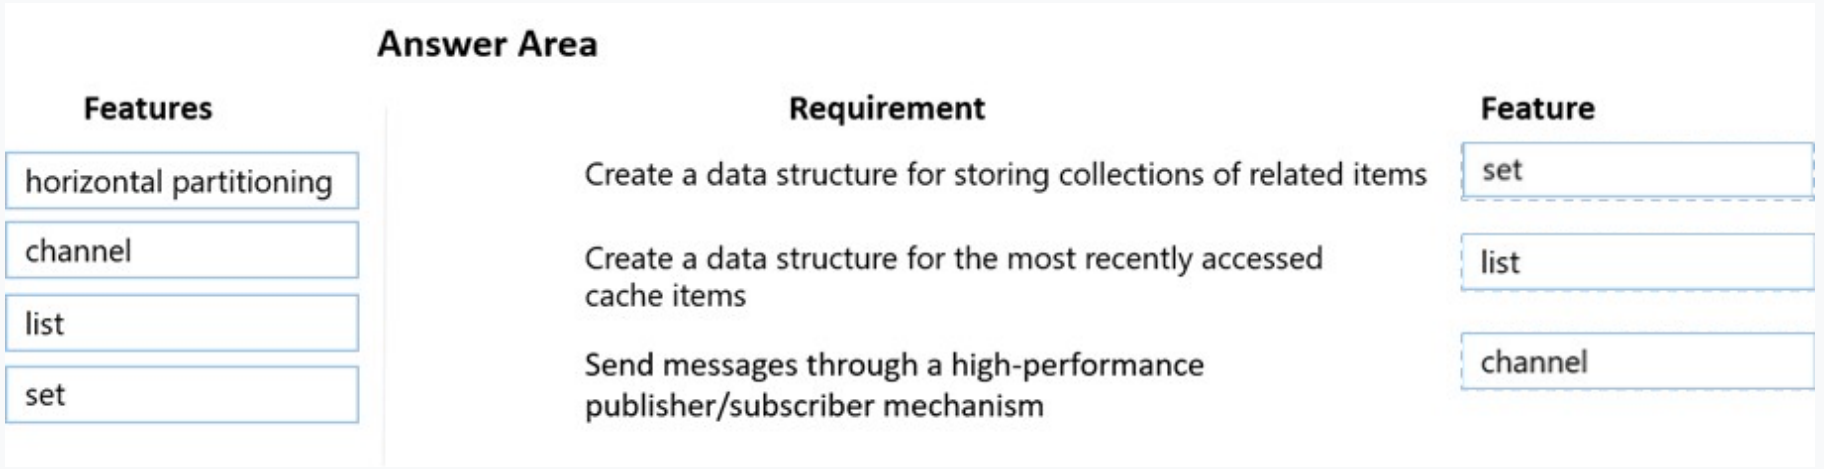

Reference:

https://www.red-gate.com/simple-talk/development/dotnet-development/overview-of-azure-cache-for-redis/ https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching

baroo1 Highly Voted 8 months, 1 week ago

This is correct!
https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching

If you look for the right words (related > set | high-performance > channel | recent > list) in documentation, it helps you validate this answer. Ignore the other comments as it only leads to confusion.

upvoted 9 times

elequiel Highly Voted 1 year ago

Got it in exam 20/10/2022

upvoted 8 times

OPT_001122 11 months, 3 weeks ago

thnx for the date

upvoted 1 times

juanckar Most Recent 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 5 times

Sriniv 8 months, 2 weeks ago

represents a collection of unique items, where each item can only appear once in the set. While a set can be used to store a collection of related items, it may not be the most suitable data structure for this requirement. This is because a set does not maintain the order of the items and does not provide an efficient way to search for a specific item.

upvoted 1 times

Sriniv 8 months, 2 weeks ago

Here are the features that match the given requirements:

Feature: horizontal partitioning
Requirement: Create a data structure for storing collections of related items

Feature: list
Requirement: Create a data structure for the most recently accessed cache items

Feature: channel
Requirement: Send messages through a high-performance publisher/subscriber mechanism

upvoted 2 times

aruni_mishra 11 months ago

Set: for related items
https://learn.microsoft.com/en-us/azure/architecture/best-practices/caching#use-tags-to-cross-correlate-cached-items

List: to find the most recently accessed items
https://learn.microsoft.com/en-us/azure/architecture/best-practices/caching#find-recently-accessed-items

upvoted 5 times

Knightie 1 year, 1 month ago

https://redis.io/docs/data-types/sets/

upvoted 1 times

qwerty112233 1 year, 1 month ago

Correct
https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching

upvoted 3 times

sghaha 1 year, 6 months ago

caching (in Korean)
https://docs.microsoft.com/ko-kr/azure/architecture/best-practices/caching

upvoted 2 times

Question #22 *Topic 5*

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 use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Telemetry.Context.Cloud.RoleInstance

B. Telemetry.Id

C. Telemetry.Name

D. Telemetry.Context.Operation.Id

E. Telemetry.Context.Session.Id

Correct Answer: *BD*
Example:
public async Task Enqueue(string payload)
{
// StartOperation is a helper method that initializes the telemetry item
// and allows correlation of this operation with its parent and children. var operation = telemetryClient.StartOperation<DependencyTelemetry>
("enqueue " + queueName);

operation.Telemetry.Type = "Azure Service Bus";
operation.Telemetry.Data = "Enqueue " + queueName;
var message = new BrokeredMessage(payload);
// Service Bus queue allows the property bag to pass along with the message.
// We will use them to pass our correlation identifiers (and other context)
// to the consumer.
message.Properties.Add("ParentId", operation.Telemetry.Id);
message.Properties.Add("RootId", operation.Telemetry.Context.Operation.Id);
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking

Community vote distribution
BD (100%)

cyberbull Highly Voted 3 years, 1 month ago
the Given answer is correct IMHO
upvoted 43 times

10x 2 years, 8 months ago
Yes. Makes sense to me too.
upvoted 8 times

mlantonis Highly Voted 2 years, 5 months ago
Correct Answer: B and D
message.Properties.Add("ParentId", operation.Telemetry.Id);
message.Properties.Add("RootId", operation.Telemetry.Context.Operation.Id);

Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking#enqueue
upvoted 18 times

alexein74 Most Recent 9 months, 1 week ago
Selected Answer: BD
Answer is correct
upvoted 1 times

OPT_001122 11 months, 3 weeks ago
Selected Answer: BD
correct ans

upvoted 1 times

AzureDJ 1 year, 7 months ago

Answer is correct

upvoted 2 times

Ravindu 2 years ago

The given answer is correct.

upvoted 1 times

glam 2 years, 5 months ago

Correct

upvoted 2 times

Immy_08 2 years, 5 months ago

The given answer is correct, also available on udemy course

upvoted 3 times

vb3d 2 years, 7 months ago

Options form the answer mentioned in the code on https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking#enqueue

upvoted 2 times

paru123456789 2 years, 8 months ago

Answer:
Telemetry.Id
Telermetry.Context.Operation.Id

upvoted 3 times

Question #23 *Topic 5*

HOTSPOT -

You are using Azure Front Door Service.

You are expecting inbound files to be compressed by using Brotli compression. You discover that inbound XML files are not compressed. The files are 9 megabytes (MB) in size.

You need to determine the root cause for the issue.

To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

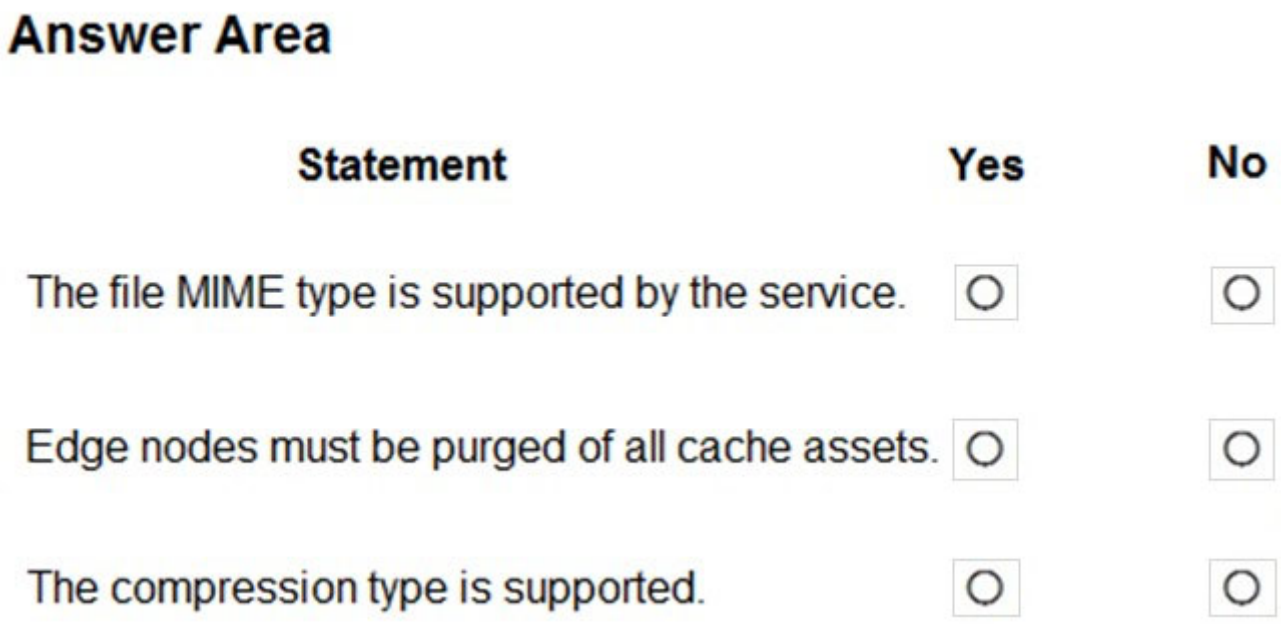

Correct Answer:

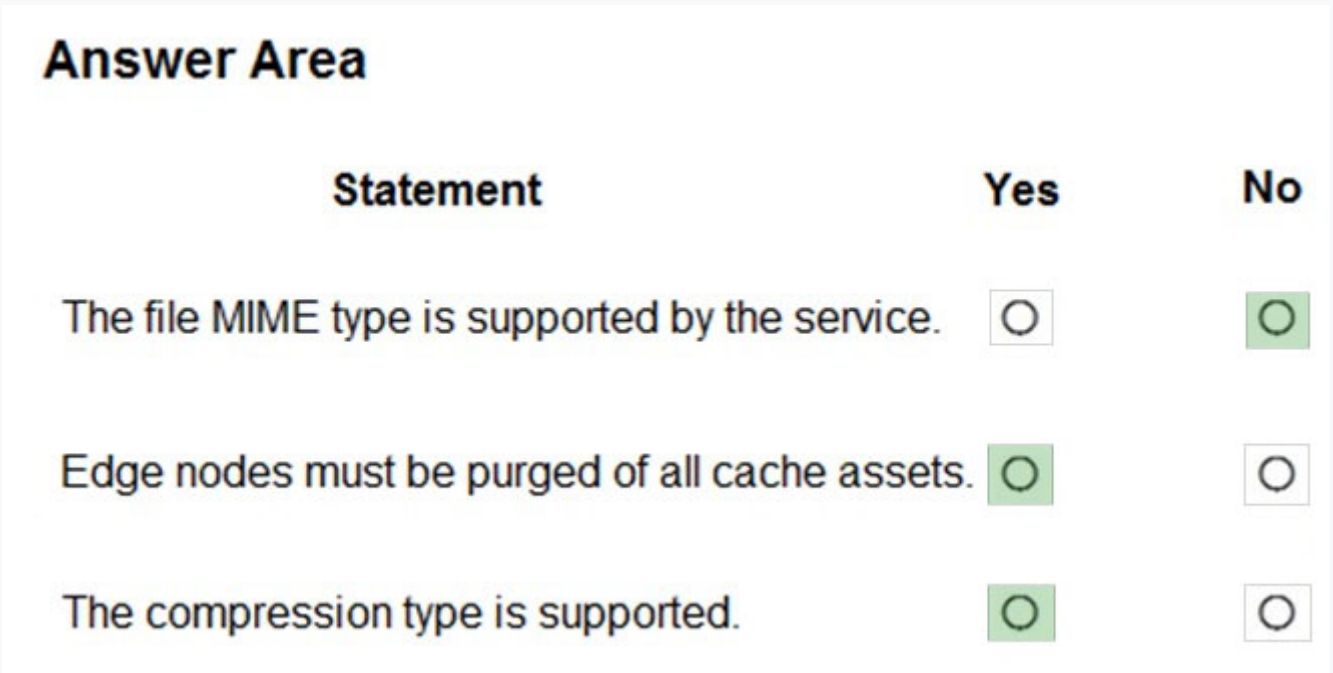

Box 1: No -

Front Door can dynamically compress content on the edge, resulting in a smaller and faster response to your clients. All files are eligible for compression.

However, a file must be of a MIME type that is eligible for compression list.

Box 2: No -

Sometimes you may wish to purge cached content from all edge nodes and force them all to retrieve new updated assets. This might be due to updates to your web application, or to quickly update assets that contain incorrect information.

Box 3: Yes -

These profiles support the following compression encodings: Gzip (GNU zip), Brotli

Reference:

https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching

bensagar10 Highly Voted 2 years, 11 months ago

1.Yes
2.No Refer - https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching ,You don't need to purge the cache assets. Here the issue is that the file size needs to be less than 8MB
3.Yes

upvoted 179 times

uffuchsi 8 months, 1 week ago

No. The file MIME type is supported by the service. So, this option does not help you to determine the root cause.

No. Edge nodes are not required to be purged. So, this option does not help you to determine the root cause.

Yes. The compression type is not supported due to size limits. This options provides root cause for not compression.

https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching#file-compression

upvoted 2 times

warchoon 7 months, 4 weeks ago

The size is not part of the compression type. Brotli supported. So no on 3 too

upvoted 2 times

nekkilodeon 6 months, 1 week ago

Correct
Xml is supported
Compression types suported brotli and gzip
File size limit of 8mb for compression
https://learn.microsoft.com/en-us/azure/frontdoor/standard-premium/how-to-compression

upvoted 2 times

Tealon 2 years, 10 months ago

I agree with this answer.

upvoted 1 times

profesorklaus 2 years, 9 months ago

I agree also.

upvoted 1 times

AakashNeedsEmAll Highly Voted 2 years, 11 months ago

application/xml is a MIME type and therefore, is supported. Answer 1 should be yes.

upvoted 26 times

WH16 Most Recent 2 months, 3 weeks ago

This was on the exam (9th August 2023). Went with proposed. Scored 910.

upvoted 1 times

sundeepch1 4 months ago

1. No
2. Yes
3. Yes

upvoted 1 times

wrak 8 months, 1 week ago

IMHO, Yes/No answers should be based on the details provided in the question, hence:
1. Yes - application/xml is supported
2. No - if we will purge cache it will still not compress the file
3. Yes - statement is about compression type and Brotli is supported

upvoted 2 times

rolling_potato_ 10 months ago

Was on my exam today (3rd of Jan 2023), went with given answer

upvoted 2 times

adilkhan 9 months ago

you mean No, Yes and Yes? your score please

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Why are the admins not correcting the answers having gone thru the highly voted discussion comments!!!

upvoted 5 times

OPT_001122 11 months, 3 weeks ago

Yes
No
Yes

upvoted 2 times

Eltooth 1 year, 4 months ago

Yes
No
Yes

upvoted 5 times

minaritochuck 1 year, 6 months ago

https://docs.microsoft.com/en-us/azure/frontdoor/standard-premium/how-to-compression

In Azure Front Door, only eligible files are compressed. To be eligible for compression, a file must:

Be of a MIME type
Be larger than 1 KB
Be smaller than 8 MB
These profiles support the following compression encodings:

gzip (GNU zip)
brotli
If the request supports more than one compression type, brotli compression takes precedence.

I guess the root cause is because xml is 9MB in size.

1. YES : xml = MIME type
2. NO : I don't think purging cache is necessary..
3. YES - brotli is supported

upvoted 4 times

Azprep 1 year, 6 months ago

1. Yes
2. No
3. Yes

upvoted 2 times

Prasu69 1 year, 7 months ago

1. NO - the file must be of a MIME type to be eligible for compression. application/xml or text/xml are not supported.
2. Yes - size should be in 1- 8MB chunks. (9 MB divides into 8MB chunk and 1 MB chunk)
3. Yes - Brotli Compression method is supported.

upvoted 1 times

Prasu69 1 year, 7 months ago

1. NO - the file must be of a MIME type to be eligible for compression. application/xml or text/xml are not supported.
2. No - size should be in 1- 8MB chunks.
3. Yes - Brotli Compression method is supported.

upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 3 times

SivajiTheBoss 1 year, 7 months ago

Question: What is the root cause for the issue, So
Answer:
1. No
2. No
3. Yes (because the file size greater than 8 mb)

upvoted 1 times

kozchris 1 year, 8 months ago

no - the files are getting stored, there is no blocking of the XML files
yes - maybe files were cached prior to some setting getting configured
yes - you need to make sure the compression is supported.

upvoted 1 times

exnaniantwort 1 year, 9 months ago

I think Azure Front Door is out of syllabus

upvoted 4 times

Wengatz 1 year, 1 month ago

I agree, but it was on the test September 22. I complained about it on the feedback form.

upvoted 1 times

Question #24 *Topic 5*

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.mp4?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:

Answer Area

Setting	Action
Caching behavior	Bypass cache Override Set if missing
Cache expiration duration	1 second 1 minute 1 hour 1 day
Query string caching behavior	Ignore query strings Bypass caching for query strings Cache every unique URL

Correct Answer:

Answer Area

Setting	Action
Caching behavior	Bypass cache **Override** Set if missing
Cache expiration duration	1 second 1 minute **1 hour** 1 day
Query string caching behavior	Ignore query strings Bypass caching for query strings **Cache every unique URL**

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.
Reference:
https://docs.microsoft.com/en-us/azure/cdn/cdn-query-string

RaviKS Highly Voted 2 years, 9 months ago
Given Answers are correct
upvoted 67 times

mlantonis Highly Voted 2 years, 5 months ago
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.
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.
upvoted 27 times

mlantonis 2 years, 5 months ago
Reference:

https://docs.microsoft.com/en-us/azure/cdn/cdn-query-string

https://docs.microsoft.com/en-us/azure/cdn/cdn-caching-rules
upvoted 9 times

mlantonis 2 years, 5 months ago
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.

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.
upvoted 19 times

jaf19f Most Recent 2 months, 2 weeks ago
I got this question (12-Aug-2023) and I chose the given answer - 932 passed
upvoted 2 times

rolling_potato_ 10 months ago
Got it in my exam today (3rd of Jan)
upvoted 2 times

adilkhan 9 months ago
are these questions enough to pass the exam?
upvoted 1 times

petitbilly 1 year, 7 months ago

Got it in exam 03/22

upvoted 3 times

AzureDJ 1 year, 7 months ago

Hope your exam went well. This answer looks right. Did you choose it?

upvoted 1 times

HimanshuNankani 2 years ago

1. Override - because we are explicitly defining the duration to be 1 hour .
We DO NEED Cache (hence it cannot be bypass cache)
We need EXACT 1 HOUR (so can't use Set If Missing which could change the specific value)

2. 1 Hour as required

3. We do need to 'Cache every unique url' because we need to provide cache for different qualities that is specified in the query parameter.

upvoted 10 times

MK22 2 years, 2 months ago

Provided answer is correct

upvoted 3 times

glam 2 years, 5 months ago

correct.

upvoted 2 times

rakelcoelho 2 years, 7 months ago

Didn't understand. On the first box it can be «override» or «set if missing»?

upvoted 4 times

titombo 2 years, 6 months ago

Me too, but according MS page the concepts are:

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.

What I don't understand is:

origin-provided cache duration: is the cache provided in Azure, right?
So in other words it would something like:
Override cache means that it will always use the cache provided in Azure.
Set if missing, this means it can use the client cache request in the request header property.

Am I wrong?

upvoted 3 times

Cricketer 2 years, 6 months ago

Not really, origin means your backend servers

upvoted 1 times

profesorklaus 2 years, 9 months ago

What hasn't been mentioned here is when you try to do on azure portal you should use cdn from Akamai or Verizon but not from Microsoft. Microsoft has default caching rules which can not be changed using portal

upvoted 1 times

yerisof 2 years, 10 months ago

why is this not ignore query string

upvoted 1 times

fredrodlima 2 years, 10 months ago

I guess is because each video with different qualities should be cached for upcoming requests to it from users

upvoted 19 times

Tealon 2 years, 10 months ago

That is correct.

upvoted 3 times

TakumaK 2 years, 5 months ago

As the item in the question. the query has the quality field. this is intended to maintain the cache based on the quality.

upvoted 2 times

Nyaku 2 years, 11 months ago

For more reference to the answer .. https://docs.microsoft.com/en-us/azure/cdn/cdn-caching-rules

upvoted 11 times

Question #25 *Topic 5*

HOTSPOT -

You are developing an ASP.NET Core time sheet application that runs as an Azure Web App. Users of the application enter their time sheet information on the first day of every month.

The application uses a third-party web service to validate data.

The application encounters periodic server errors due to errors that result from calling a third-party web server. Each request to the third-party server has the same chance of failure.

You need to configure an Azure Monitor alert to detect server errors unrelated to the third-party service. You must minimize false-positive alerts.

How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

```
"type": "Microsoft.Insights/metricAlerts",
"properties": {
"criteria": {
   "odata.type": ". . .",
   "allOf": [
     {
       "criterionType": "[                              ]",
                          DynamicThresholdCriterion
                          SingleResourceMultipleMetricCriteria

       "metricName": "[          ]",
                      Http4xx
                      Http5xx

       "alertSensitivity": "[          ]"
                            Low
                            High
}]}}
```

Correct Answer:

```
"type": "Microsoft.Insights/metricAlerts",
"properties": {
"criteria": {
   "odata.type": ". . .",
   "allOf": [
     {
       "criterionType": "[DynamicThresholdCriterion                ]",
                          DynamicThresholdCriterion
                          SingleResourceMultipleMetricCriteria

       "metricName": "[          ]",
                      Http4xx
                      Http5xx

       "alertSensitivity": "[          ]"
                            Low
                            High
}]}}
```

Box 1: DynamicThresholdCriterion

Box 2: Http5xx -

Server errors are in the 5xx range.

Client errors are in the 4xx range

Box 3: Low -

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-dynamic-thresholds

OPT_001122 Highly Voted 10 months, 3 weeks ago

DynamicThresholdCriterion
Http5XX
Low
upvoted 9 times

coffecold Highly Voted 1 year ago
DynamicThresholdCriterion: Dynamic thresholds in metric alerts use advanced machine learning to learn metric
Http5XX, just a name , we are monitoring server errors, so 500 range error reponses
AlertSensitvity : would set is to high. "You must minimize false-positive alerts" The thresholds will be tight and close to the metric series pattern. An alert rule will be triggered on the smallest deviation, resulting in more alerts.
upvoted 6 times

Yumi21 11 months, 3 weeks ago
I think sensitivity should be low due to "false positive" requirements, if we set to high we report on every smallest deviation, and it means there is a highly chance of false positive
upvoted 11 times

b1ade 11 months, 1 week ago
Agree. Definition according to MS:

High: The thresholds will be tight and close to the metric series pattern. An alert rule will be triggered on the smallest deviation, resulting in more alerts.
Medium: The thresholds will be less tight and more balanced. There will be fewer alerts than with high sensitivity (default).
Low: The thresholds will be loose with more distance from metric series pattern. An alert rule will only trigger on large deviations, resulting in fewer alerts.
upvoted 6 times

warchoon Most Recent 7 months, 4 weeks ago
"SingleResourceMultipleMetricCriteria" is used for "odata.type", not "cryterionType".
upvoted 2 times

ozzyy123 9 months ago
It says "errors unrelated to the third-party service". So, I think it should be Http4xx. Not 5xx.
upvoted 1 times

surprise0011 6 months, 3 weeks ago
but you need to detect server errors:
"You need to configure an Azure Monitor alert to detect server errors..."
upvoted 1 times

sghaha 1 year, 6 months ago
about alerts-dynamic-thresholds (in Korean)
https://docs.microsoft.com/ko-kr/azure/azure-monitor/alerts/alerts-dynamic-thresholds
upvoted 2 times

ArturKon 1 year, 1 month ago
You can provide us links in other language. Just change ko-kr in URL to, (for ex.) en-us
upvoted 4 times

kerpse 1 year ago
Or in preferred language by just removing the language code, e.g.
https://docs.microsoft.com/azure/azure-monitor/alerts/alerts-dynamic-thresholds
upvoted 3 times

Question #26 *Topic 5*

You are developing a web application that uses Azure Cache for Redis. You anticipate that the cache will frequently fill and that you will need to evict keys.
You must configure Azure Cache for Redis based on the following predicted usage pattern: A small subset of elements will be accessed much more often than the rest.
You need to configure the Azure Cache for Redis to optimize performance for the predicted usage pattern.
Which two eviction policies will achieve the goal?
NOTE: Each correct selection is worth one point.

A. noeviction

B. allkeys-lru

C. volatile-lru

D. allkeys-random

E. volatile-ttl

F. volatile-random

Correct Answer: *BC*
B: The allkeys-lru policy evict keys by trying to remove the less recently used (LRU) keys first, in order to make space for the new data added.
Use the allkeys-lru policy when you expect a power-law distribution in the popularity of your requests, that is, you expect that a subset of elements will be accessed far more often than the rest.
C: volatile-lru: evict keys by trying to remove the less recently used (LRU) keys first, but only among keys that have an expire set, in order to make space for the new data added.
Note: The allkeys-lru policy is more memory efficient since there is no need to set an expire for the key to be evicted under memory pressure.
Reference:
https://redis.io/topics/lru-cache

Community vote distribution
BC (100%)

sghaha Highly Voted 1 year, 6 months ago
my choices are BC
upvoted 9 times

qwerty112233 1 year, 1 month ago
agree!
upvoted 3 times

alexein74 Highly Voted 9 months, 1 week ago
Selected Answer: BC
B. allkeys-lru
C. volatile-lru

allkeys-lru : This eviction policy evicts the least recently used (LRU) key among all keys, regardless of the key's time-to-live (TTL) value. This policy is suitable when you have a small subset of elements that will be accessed much more often than the rest, as the most frequently used elements will be kept in cache while the least frequently used elements will be evicted.

volatile-lru: This eviction policy evicts the least recently used (LRU) key among keys with an expiring TTL. This policy is suitable when you have a small subset of elements that will be accessed much more often than the rest and also when you want to expire keys that are not used frequently.
upvoted 7 times

alinuxguru70 Most Recent 1 month, 1 week ago
allkeys-lru would include volatile-lru, it is redundant. Therefore B allkeys-lru and E volatile-ttl.
upvoted 1 times

Ciupaz 1 month, 2 weeks ago
"least recently used", not "less recently used". Take care.
upvoted 1 times

AymanAkk 10 months, 1 week ago
Selected Answer: BC
correct answers.

upvoted 1 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose：B、C

upvoted 2 times

OPT_001122 11 months ago

Selected Answer: BC

B. allkeys-lru
C. volatile-lru

upvoted 2 times

ExamTopicsAnkur 11 months ago

Selected Answer: BC

https://redis.io/docs/reference/eviction/

upvoted 2 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This question was on it.

upvoted 4 times

Question #27 *Topic 5*

DRAG DROP -

An organization has web apps hosted in Azure.

The organization wants to track events and telemetry data in the web apps by using Application Insights.

You need to configure the web apps for Application Insights.

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:

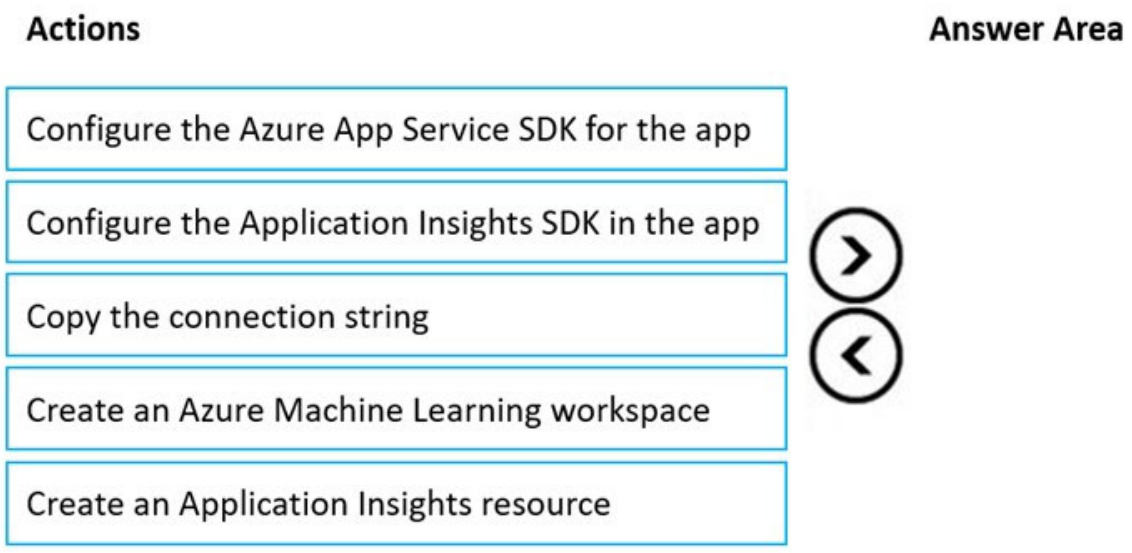

Correct Answer:

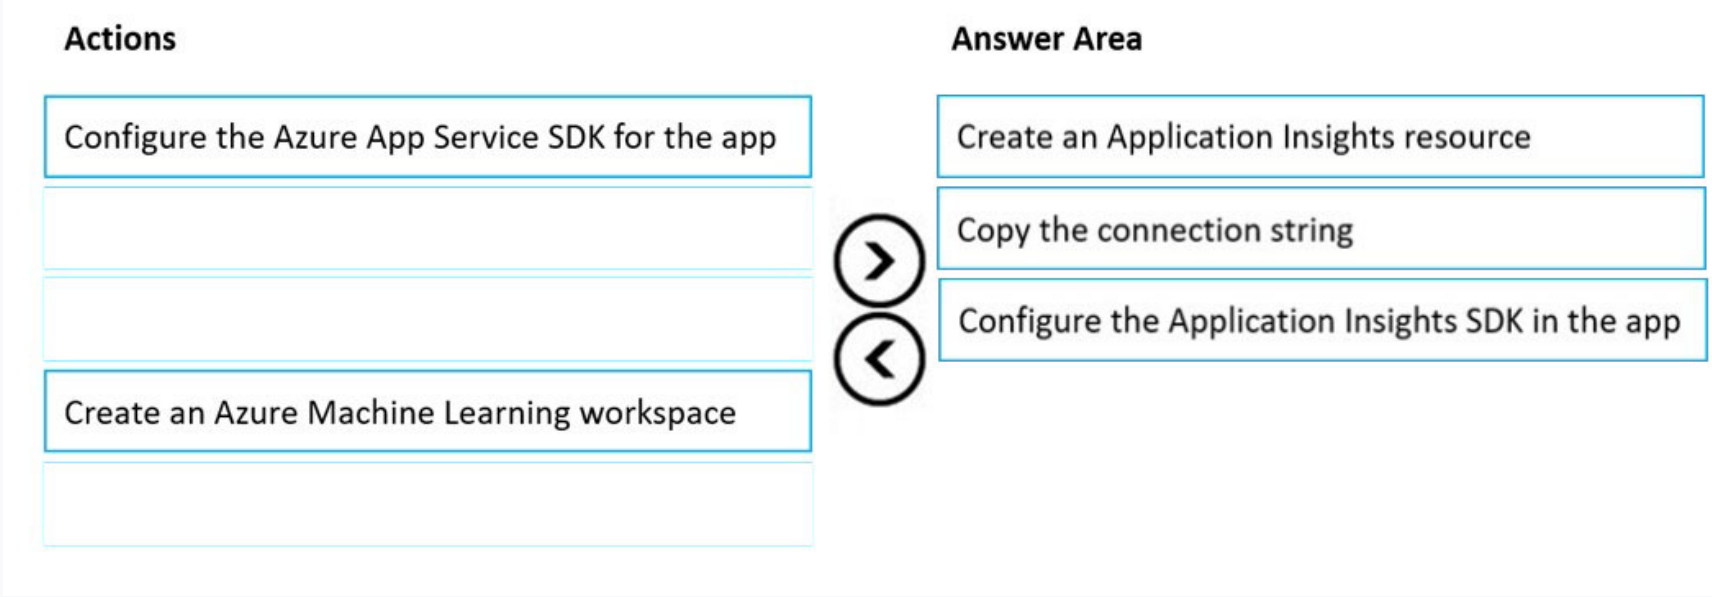

Step 1: Create an Application Insights resource

Creating an Application Insights workspace-based resource us a prerequisite.

Step 2: Copy the connection string

A connection string identifies the resource that you want to associate with your telemetry data. It also allows you to modify the endpoints that your resource will use as a destination for your telemetry. You'll need to copy the connection string and add it to your application's code or to an environment variable.

Step 3: Configure the Application Insights SDK in the app

The Application Insights SDK for ASP.NET Core can monitor your applications no matter where or how they run.

Install the Application Insights SDK NuGet package for ASP.NET Core.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/asp-net-core

osagiestar Highly Voted 1 year ago

The given answer is correct.
1. Create an Application Insights resource
2. Copy the instrumentation key
3. Install the SDK in your app
https://learn.microsoft.com/en-us/azure/azure-monitor/app/create-new-resource

upvoted 13 times

NombreFalso Highly Voted 8 months, 1 week ago

Why would you make a Machine Learning Workspace LOL

upvoted 9 times

threeslice91 Most Recent 4 months ago

The order of 2nd and 3rd does it matter?

upvoted 2 times

Heighte 6 months, 3 weeks ago

bad question, connection string != instrumentation key

upvoted 3 times

elequiel 1 year ago

Got it in exam 20/10/2022

upvoted 6 times

willchenxa 1 year, 1 month ago

1. Create resource
2. install Application insight SDK
3. setup connection string

according to provided link: https://docs.microsoft.com/en-us/azure/azure-monitor/app/asp-net-core

upvoted 4 times

AymanAkk 10 months, 1 week ago

yes for a setup , but here we have the " copy " word

upvoted 2 times

finnishr 1 year, 1 month ago

Correctmundo.

upvoted 2 times

Question #28 *Topic 5*

An organization hosts web apps in Azure. The organization uses Azure Monitor.
You discover that configuration changes were made to some of the web apps.
You need to identify the configuration changes.
Which Azure Monitor log should you review?

A. AppServiceAppLogs

B. AppServiceEnvironmentPlatformlogs

C. AppServiceConsoleLogs

D. AppServiceAuditLogs

Correct Answer: *B*
The log type AppServiceEnvironmentPlatformLogs handles the App Service Environment: scaling, configuration changes, and status logs.
Incorrect:
AppServiceAppLogs contains logs generated through your application.
AppServiceAuditLogs logs generated when publishing users successfully log on via one of the App Service publishing protocols.
Reference:
https://docs.microsoft.com/en-us/azure/app-service/troubleshoot-diagnostic-logs

Community vote distribution
B (100%)

serpevi Highly Voted 1 year, 1 month ago

Got it 10/09/2022. Went with B. Score 927

upvoted 13 times

coffecold Highly Voted 1 year ago

AppServiceEnvironmentPlatformlogs only logs when app services are hosted in an App Service Environment witch is a deployment of Azure App Service into a subnet in your Azure Virtual Network (VNet). It is a Premium service option

The problem is, all other anwers are wrong :
AppServiceAppLogs : logs from inside your app
AppServiceConsoleLogs : logs from the console
AppServiceAuditLogs : logs who are logged in or out.

but...
AppServiceFileAuditLogs will do the job for premium tiers.
https://ruslany.net/2020/01/how-to-monitor-azure-app-service-content-changes-and-publishing-activity/

upvoted 5 times

Dilmurod 7 months ago

AppServiceEnvironmentPlatformLogs - App Service Environment: scaling, configuration changes, and status logs
AppServiceFileAuditLogs - File changes made to the site content; only available for Premium tier and above
So, the correct answer is AppServiceEnvironmentPlatformlogs

upvoted 1 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 2 times

OPT_001122 11 months, 1 week ago

Selected Answer: B

You discover that configuration changes were made to some of the web apps.

hence the ans is correct B

upvoted 2 times

seckinfo 11 months, 2 weeks ago

Ref: https://learn.microsoft.com/en-us/azure/app-service/troubleshoot-diagnostic-logs#supported-log-types

upvoted 3 times

qwerty112233 1 year, 1 month ago

Correct

upvoted 2 times

Question #29 *Topic 5*

You develop and deploy an Azure App Service web app to a production environment. You enable the Always On setting and the Application Insights site extensions.
You deploy a code update and receive multiple failed requests and exceptions in the web app.
You need to validate the performance and failure counts of the web app in near real time.
Which Application Insights tool should you use?

A. Profiler

B. Smart Detection

C. Live Metrics Stream

D. Application Map

E. Snapshot Debugger

Correct Answer: *C*
Live Metrics Stream -
Deploying the latest build can be an anxious experience. If there are any problems, you want to know about them right away, so that you can back out if necessary. Live Metrics Stream gives you key metrics with a latency of about one second.
With Live Metrics Stream, you can:
* Validate a fix while it's released, by watching performance and failure counts.
* Etc.

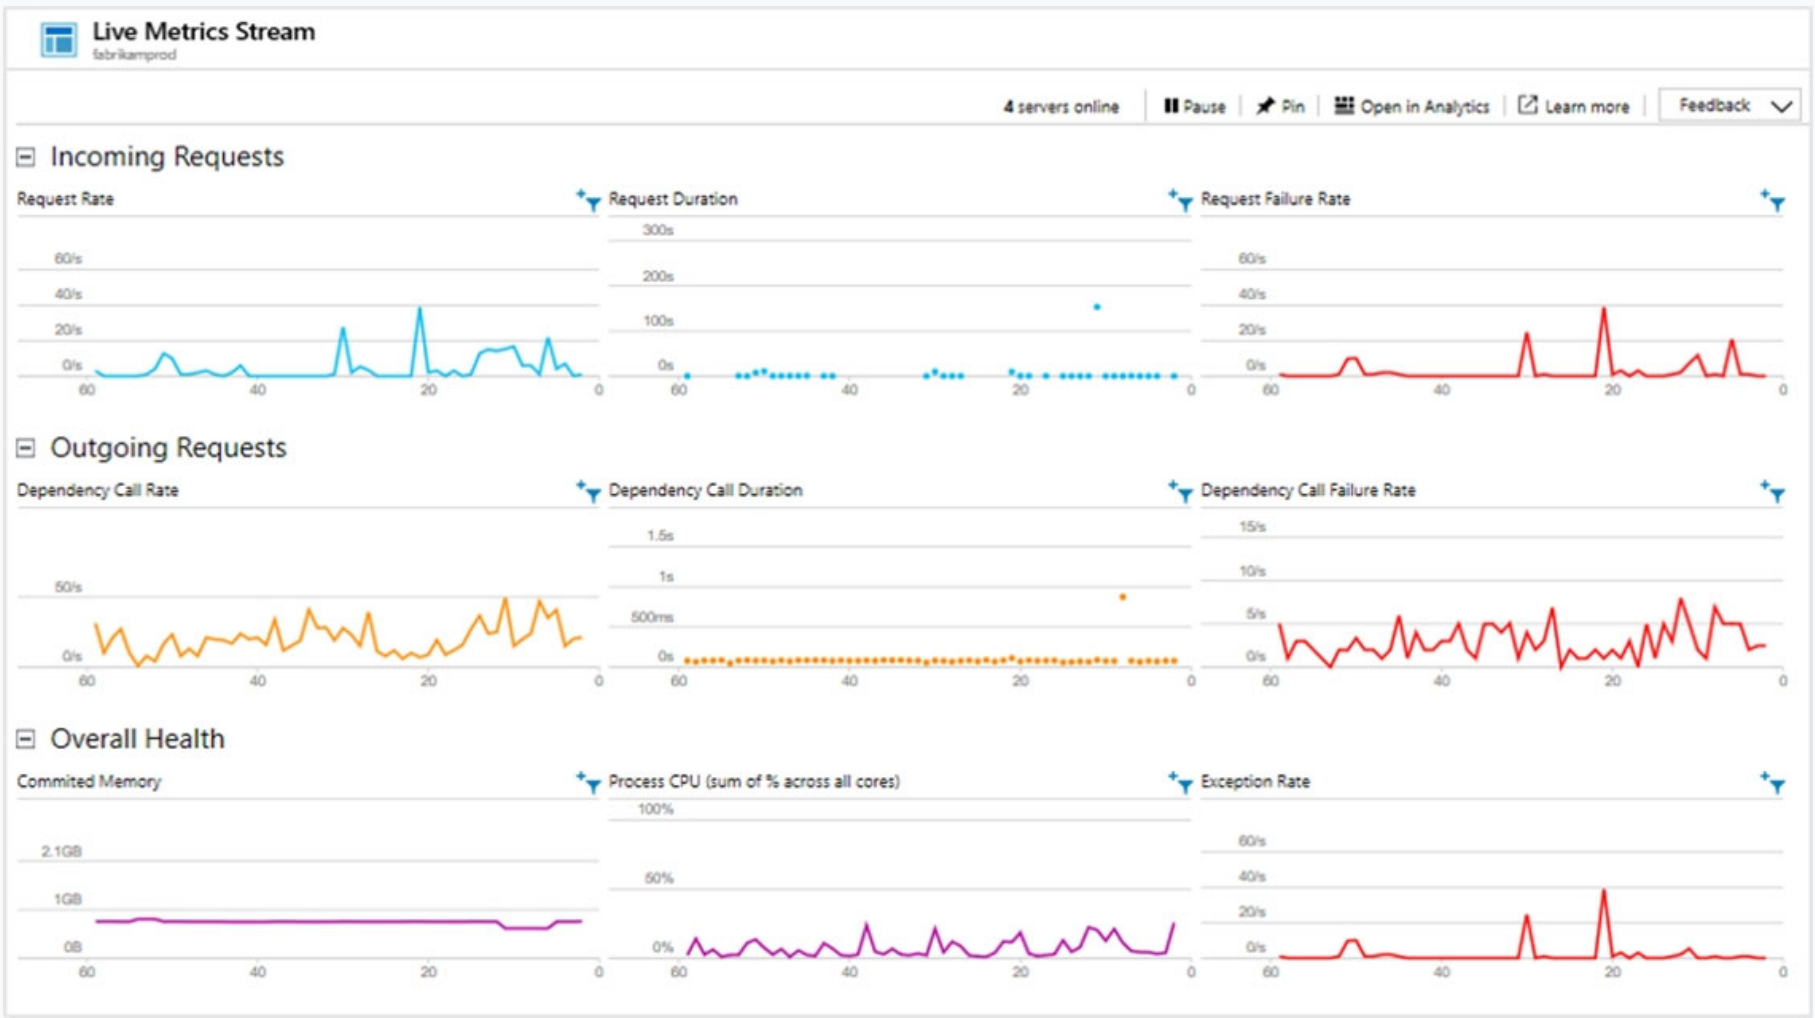

Incorrect:
* Profiler
Azure Application Insights Profiler provides performance traces for applications running in production in Azure. Profiler:
Captures the data automatically at scale without negatively affecting your users.
Helps you identify the ג€hotג€ code path spending the most time handling a particular web request.
* Snapshot debugger
When an exception occurs, you can automatically collect a debug snapshot from your live web application. The snapshot shows the state of source code and variables at the moment the exception was thrown. The Snapshot Debugger in Azure Application Insights monitors exception telemetry from your web app. It collects snapshots on your top-throwing exceptions so that you have the information you need to diagnose issues in production.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/live-stream

Community vote distribution
C (75%) B (25%)

adilkhan Highly Voted 9 months ago

C is correct

upvoted 5 times

g2000 Most Recent 9 months, 2 weeks ago

the key is near real time.
https://learn.microsoft.com/en-us/azure/azure-monitor/app/live-stream

upvoted 2 times

Tornik 9 months, 2 weeks ago

Selected Answer: B

Isnt it B ?
"Smart detection automatically warns you of potential performance problems and failure anomalies in your web application. It performs proactive analysis of the telemetry that your app sends to Application Insights. "

https://learn.microsoft.com/en-us/azure/azure-monitor/alerts/proactive-diagnostics

upvoted 2 times

Tornik 9 months, 2 weeks ago

This is the same question just Smart Detection as answer, so which is it ? I think it might be Live Metrics Stream after all as it says "Near Realtime" while Smart Detection send out notification which is not realtime ?
https://www.pass4success.com/microsoft/discussions/exam-az-204-topic-11-question-59-discussion

upvoted 1 times

OPT_001122 11 months, 1 week ago

Selected Answer: C

C. Live Metrics Stream

upvoted 3 times

Perplex 1 year ago

Selected Answer: C

correct

upvoted 3 times

gmishra88 1 year ago

Another beautiful question with a red-herring: "always on". People who know that is needed for setting Insights profiler will want to answer that. A common thing in many exams, but Microsoft is taking this personally and checking how to trick

upvoted 1 times

Mousavi 1 year, 1 month ago

why not A ?

upvoted 1 times

kampatra 1 year, 1 month ago

Ans Correct

upvoted 1 times

Question #30 *Topic 5*

HOTSPOT -

You deploy an ASP.NET web app to Azure App Service.

You must monitor the web app by using Application Insights.

You need to configure Application Insights to meet the requirements.

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

NOTE: Each correct selection is worth one point.

Hot Area:

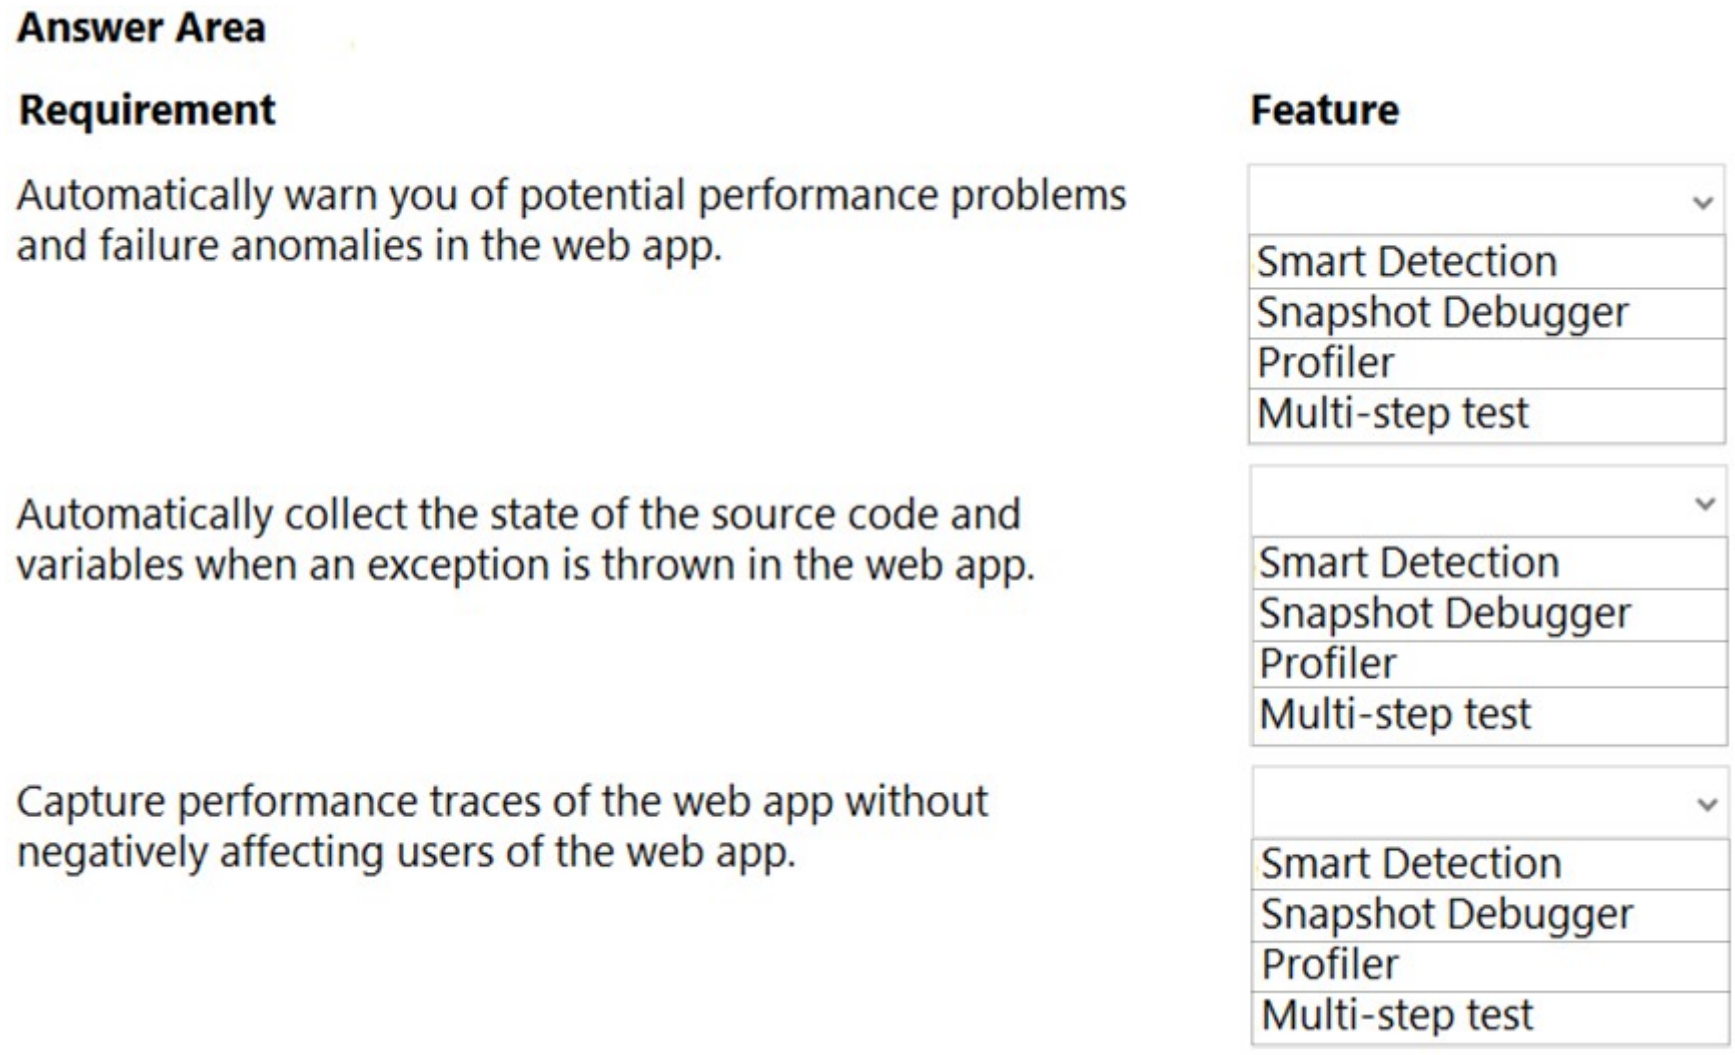

Correct Answer:

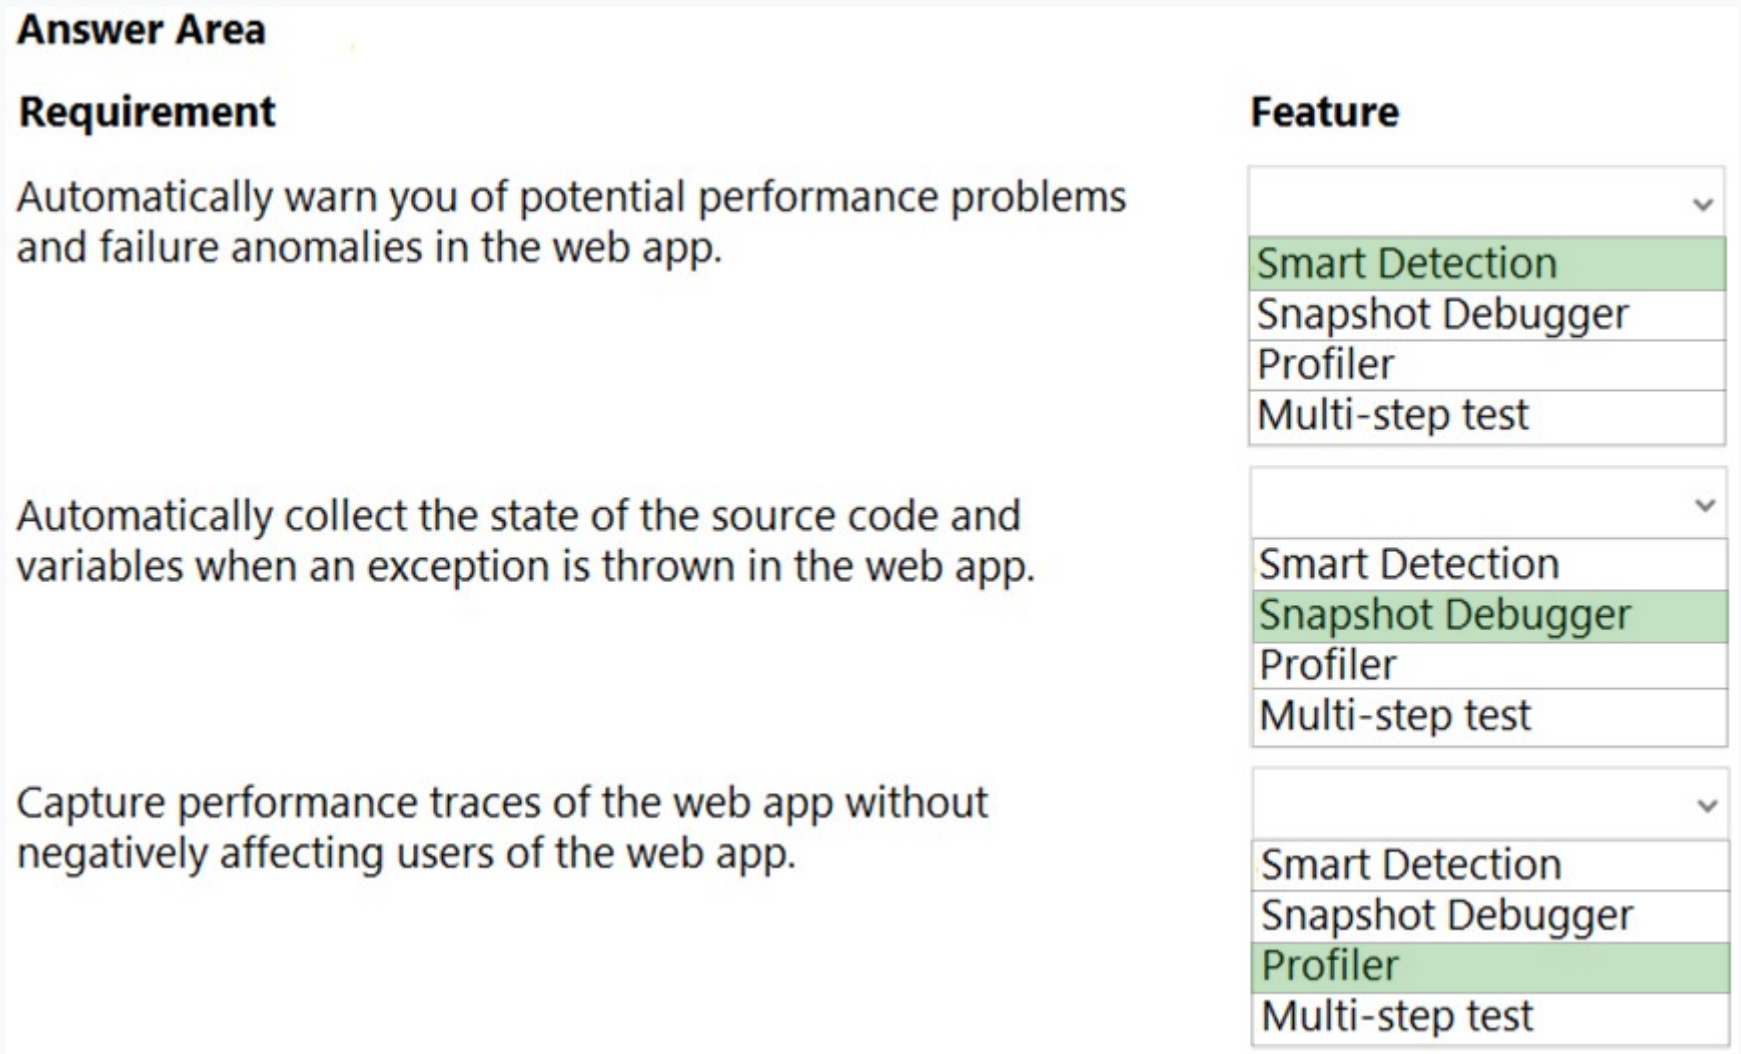

Box 1: Smart Detection -

Smart detection automatically warns you of potential performance problems and failure anomalies in your web application. It performs proactive analysis of the telemetry that your app sends to Application Insights. If there is a sudden rise in failure rates, or abnormal patterns in client or server performance, you get an alert. This feature needs no configuration. It operates if your application sends enough telemetry.

Box 2: Snapshot Debugger -
When an exception occurs, you can automatically collect a debug snapshot from your live web application. The snapshot shows the state of source code and variables at the moment the exception was thrown. The Snapshot Debugger in Azure Application Insights monitors exception telemetry from your web app. It collects snapshots on your top-throwing exceptions so that you have the information you need to diagnose issues in production.

Box 3: Profiler -
Azure Application Insights Profiler provides performance traces for applications running in production in Azure. Profiler:
Captures the data automatically at scale without negatively affecting your users.
Helps you identify the ג€hotג€ code path spending the most time handling a particular web request.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/proactive-diagnostics https://docs.microsoft.com/en-us/azure/azure-monitor/snapshot-debugger/snapshot-debugger https://docs.microsoft.com/en-us/azure/azure-monitor/profiler/profiler-overview

OPT_001122 Highly Voted 11 months, 1 week ago
Answer is correct
upvoted 9 times

KingChuang 10 months, 1 week ago
On my exam 2022-12-26.
Chose：Given Answer (1、2、3)
upvoted 5 times

ExamNow Highly Voted 1 year, 1 month ago
1.Smart detection
2.Snapshot
3.user profile
upvoted 7 times

TheExamMaster2020 Most Recent 11 months, 2 weeks ago
Did my exam on 15th November 2022. This question was on it.
upvoted 4 times

devsb 1 year, 1 month ago
Answer is correct
upvoted 4 times

Question #31 *Topic 5*

You are building a web application that performs image analysis on user photos and returns metadata containing objects identified. The image analysis is very costly in terms of time and compute resources. You are planning to use Azure Redis Cache so duplicate uploads do not need to be reprocessed.
In case of an Azure data center outage, metadata loss must be kept to a minimum.
You need to configure the Azure Redis cache instance.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Configure Azure Redis with AOF persistence.

B. Configure Azure Redis with RDB persistence.

C. Configure second storage account for persistence.

D. Set backup frequency to the minimum value.

Correct Answer: *BD*
RDB persistence - When you use RDB persistence, Azure Cache for Redis persists a snapshot of your cache in a binary format. The snapshot is saved in an
Azure Storage account. The configurable backup frequency determines how often to persist the snapshot. If a catastrophic event occurs that disables both the primary and replica cache, the cache is reconstructed using the most recent snapshot.
Note: Azure Cache for Redis supports zone redundant configurations in the Premium and Enterprise tiers. A zone redundant cache can place its nodes across different Azure Availability Zones in the same region. It eliminates data center or AZ outage as a single point of failure and increases the overall availability of your cache.
Incorrect:
Not A: Zone redundancy doesn't support AOF persistence or work with geo-replication currently.
Not C: No need for a second storage account.
Reference:
https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-premium-persistence

Community vote distribution
AC (70%) BD (30%)

finnishr Highly Voted 1 year, 1 month ago
The key here is "In case of an Azure data center outage, metadata loss must be kept to a minimum."

So the correct answer is AC.
upvoted 20 times

Bear_Polar 10 months, 4 weeks ago
I think this is correct.

"RDB is NOT good if you need to minimize the chance of data loss in case Redis stops working". source: https://redis.io/docs/management/persistence/

since the question doesn't mention ZRS for Redis cache, we can use AOF with a 2nd storage to store replica cache data. check the config with AOF: https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-premium-persistence#set-up-data-persistence
upvoted 4 times

serpevi Highly Voted 1 year, 1 month ago
got it 10/09/2022. went with AC. Score 927
upvoted 12 times

dzits 11 months, 3 weeks ago
Very confused why it's AC if it states on MS website that Zone redundancy doesn't support AOF persistence or work with geo-replication currently. https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-zone-redundancy (as mentioned by coffecold). I would go with provided answer B,D unless someone can prove why this wouldn't be the case otherwise?
upvoted 2 times

AndySmith Most Recent 5 days, 11 hours ago
Selected Answer: AC
Technically here could be just 2 valid answer:
AC and BD.
But, we need to select the option mostly matching the requirements.

Since, they ask to minimize data loss (rather than speed of recovery) then AC is the right answer, because:
1) With AOF persistence every write transaction is logged. So, in theory we can recover everything up to the moment of failure.
2) Additionally, we'll enable Secondary Storage for AOF, which will help to restore Redis image if primary zone was fully destroyed.

Why RDB is worse? Even if we minimize the period of snapshotting to minimum (15 mins), then we can lose up to 15 mins of data in case of data center outage. Though advantage is that we can restore the snapshot much faster than from the AOF logs.

You can find full info about AOF and RDB settings in Azure Redis Cache on this link:
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-premium-persistence?tabs=premium

upvoted 2 times

Dianahu 3 months ago

https://redis.io/docs/management/persistence/ Ok, so what should I use? If you care a lot about your data, but still can live with a few minutes of data loss in case of disasters, you can simply use RDB alone.

upvoted 1 times

nekkilodeon 6 months, 1 week ago

Selected Answer: BD

B,D is correct
Zone redundancy doesn't support AOF persistence or work with geo-replication currently", meaning it shouldn't be A because "In case of an Azure data center outage, metadata loss must be kept to a minimum.
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-zone-redundancy

upvoted 1 times

lh2022 7 months ago

https://learn.microsoft.com/en-us/training/modules/develop-for-azure-cache-for-redis/3-configure-azure-cache-redis

The Premium tier allows you to persist data in two ways to provide disaster recovery:
RDB persistence takes a periodic snapshot and can rebuild the cache using the snapshot in case of failure.
AOF persistence saves every write operation to a log that is saved at least once per second. This creates bigger files than RDB but has less data loss.

So AOF has less data loss which satisfies the requirement.

So AOF with data written is a second storage account with ZRS should do the job.

I would also go with AC

upvoted 1 times

damianadalid 5 months, 1 week ago

Can we just use that second Storage account based in a different region WITHOUT ZRS? That is the main breaking point about not choosing AOF.

upvoted 1 times

damianadalid 5 months, 1 week ago

Even more, the restriction is: "Currently, you can't use Append-only File (AOF) persistence or geo-replication with multiple replicas (more than one replica).". So, as long as we have just one replica set, can we use AOF without issues?
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-multi-replicas

upvoted 1 times

adilkhan 8 months ago

A,C chatGPT!

upvoted 1 times

wrak 8 months ago

AOF leverage on the second storage account: With AOF persistence enabled, write operations to the cache are saved to the named storage account (or accounts if you've configured a second storage account). If there's a catastrophic failure that takes down both the primary and replica cache, the stored AOF log is used to rebuild the cache.

upvoted 1 times

NombreFalso 8 months, 1 week ago

Goodluck!

upvoted 1 times

Sriniv 8 months, 2 weeks ago

Options C and D are not relevant for this scenario. Configuring a second storage account is not necessary since Redis cache already has persistence options available. Setting backup frequency to minimum value is also not necessary since AOF and RDB persistence options already ensure that data is regularly saved to disk.

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: AC

finnishr is right

upvoted 1 times

toysky731 10 months, 1 week ago

Selected Answer: AC

finnishr is correct

upvoted 1 times

edmond 10 months, 3 weeks ago

Selected Answer: BD

So many misleading guys, don't answer if you are not sure! Why do you think it's AC?

"Zone redundancy doesn't support AOF persistence or work with geo-replication currently", meaning it shouldn't be A because "In case of an Azure data center outage, metadata loss must be kept to a minimum".

And it's meaningless to use another storage account so it should not be C.

upvoted 2 times

aruni_mishra 11 months ago

B & D:
"In case of an Azure data center outage, metadata loss must be kept to a minimum": this can only be achieved by "RDB persistence". period.
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-zone-redundancy.

Backup Frequency: This interval starts counting down after the previous backup operation successfully completes. When it elapses, a new backup starts.
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-premium-persistence#which-persistence-model-should-i-choose

upvoted 3 times

OPT_001122 11 months, 1 week ago

Selected Answer: AC

A. Configure Azure Redis with AOF persistence.
C. Configure second storage account for persistence.

upvoted 2 times

OPT_001122 11 months, 3 weeks ago

the admin must correct the answers based on the highly voted discussion

upvoted 3 times

Ciupaz 4 months, 3 weeks ago

I agree.

upvoted 1 times

ms_master 1 year ago

Selected Answer: AC

correct answer is AC.

upvoted 1 times

Question #32 *Topic 5*

You are developing an Azure-based web application. The application goes offline periodically to perform offline data processing. While the application is offline, numerous Azure Monitor alerts fire which result in the on-call developer being paged.
The application must always log when the application is offline for any reason.
You need to ensure that the on-call developer is not paged during offline processing.
What should you do?

A. Add Azure Monitor alert processing rules to suppress notifications.

B. Disable Azure Monitor Service Health Alerts during offline processing.

C. Create an Azure Monitor Metric Alert.

D. Build an Azure Monitor action group that suppresses the alerts.

Correct Answer: *D*
You can use alert processing rules to add action groups or remove (suppress) action groups from your fired alerts.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-action-rules

Community vote distribution
A (92%) 8%

finnishr Highly Voted 1 year, 1 month ago

Selected Answer: A
I'd choose A.
upvoted 17 times

scrafy Highly Voted 1 year, 1 month ago

It's A 100% sure
https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-processing-rules?tabs=portal
upvoted 9 times

dddddd111 Most Recent 1 week, 6 days ago

Selected Answer: D
D. as of 18Oct23
upvoted 1 times

katrang 1 month, 2 weeks ago

Selected Answer: A
n all these cases, an alert processing rule provides an easy way to suppress notifications.
upvoted 1 times

warchoon 7 months, 4 weeks ago

Selected Answer: A
The explanation points to A
upvoted 1 times

adilkhan 8 months ago

D. Build an Azure Monitor action group that suppresses the alerts. (Chat GPT)

To ensure that the on-call developer is not paged during offline processing while still logging when the application is offline, you can create an Azure Monitor action group that suppresses the alerts during that time. This action group can be configured to temporarily suspend alerts or send them to a different channel, such as an email or a different paging system.
upvoted 1 times

warchoon 7 months, 4 weeks ago

ChatGPT is incorrect as usual
upvoted 2 times

NombreFalso 8 months, 1 week ago

The WebApp is based
upvoted 1 times

SATHTECH 11 months, 1 week ago

Selected Answer: D

ANSWER D
Suppression: This action removes all the action groups from the affected fired alerts. So, the fired alerts won't invoke any of their action groups, not even at the end of the maintenance window. Those fired alerts will still be visible when you list your alerts in the portal, Azure Resource Graph, API, or PowerShell. The suppression action has a higher priority over the Apply action groups action. If a single fired alert is affected by different alert processing rules of both types, the action groups of that alert will be suppressed.
URL : https://learn.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-processing-rules?tabs=portal#suppress-notifications-during-planned-maintenance

upvoted 1 times

Bear_Polar 10 months, 4 weeks ago

does your source refer to ans A?

upvoted 2 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: A

It's A

upvoted 3 times

gmishra88 1 year ago

You can use alert processing rules to add action groups or remove (suppress) action groups from your fired alerts.

upvoted 2 times

Question #33 *Topic 5*

You are developing an online game that includes a feature that allows players to interact with other players on the same team within a certain distance. The calculation to determine the players in range occurs when players move and are cached in an Azure Cache for Redis instance.

The system should prioritize players based on how recently they have moved and should not prioritize players who have logged out of the game.

You need to select an eviction policy.

Which eviction policy should you use?

A. allkeys-lru

B. volatile-lru

C. allkeys-lfu

D. volatile-ttl

Correct Answer: *A*

Community vote distribution

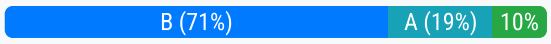

baroo1 Highly Voted 8 months, 1 week ago

Selected Answer: B

Changed my mind;
There must be a way to tell our redis that logged off users must not be prioritized.
Sample: User A moves and then automatically logs-off. With allkeys-lru we can't distinguish this particularity. With volatile-lru we can tell our redis what are good candidates to be removed using different TTL values.

upvoted 10 times

surprise0011 6 months, 3 weeks ago

totally agree with you my friend

upvoted 2 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000.

upvoted 2 times

AlexChacon99 Most Recent 2 weeks ago

Hi! I just passed my AZ 204 exam, I have a question, I don't want my employer to know my score, only that I passed, is there any way they can figure out my score if I only show them the certificate PDF ? Thank you!

upvoted 1 times

dom271219 2 weeks, 1 day ago

Selected Answer: A

There's no link with ttl and players who have logged out of the game.

In Redis, you can set the TTL (Time To Live) for keys using commands such as EXPIRE, PEXPIRE, or similar commands that allow you to specify a lifespan for a key. The TTL specifies the period for which the key will be retained in the Redis database before being automatically removed.

upvoted 1 times

mic_ 1 month ago

I think the answer is correct. We can remove the related entry from the cache after logoff or simply grab the next entry from the cache if the user is inactive. Volatile policies could even prevent caching, see
https://redis.io/docs/reference/eviction/
The policies volatile-lru, volatile-lfu, volatile-random, and volatile-ttl behave like noeviction if there are no keys to evict matching the prerequisites.

upvoted 1 times

unraval 1 month, 1 week ago

Selected Answer: B

Think it should be B : volatile-lru
volatile-lru : Removes least recently used keys with the expire field set to true.
The expire field for the logged out players can be marked true indicating Redis to evict those entries

upvoted 2 times

sundeepch1 4 months ago

ChatGPT answer option is volatile-ttl

upvoted 1 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.

Had cases:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 2 times

adilkhan 8 months ago

The "volatile-ttl" eviction policy evicts keys based on their expiration time, which is set when the key is created or updated. This policy is the most suitable for this scenario since it allows the system to prioritize recently moved players by setting a shorter expiration time for their cached location data. It also ensures that players who have logged out of the game are automatically removed from the cache when their data expires. Chat GPT

upvoted 3 times

baroo1 8 months, 1 week ago

Selected Answer: A

allkeys-lru: Keeps most recently used keys; removes least recently used (LRU) keys
volatile-lru: Removes least recently used keys with the expire field set to true.
Source: https://redis.io/docs/reference/eviction/

Not prioritizing logged off players does not mean to remove them! volatile-lru will not be the best solution at all.

upvoted 1 times

NombreFalso 8 months, 2 weeks ago

Selected Answer: A

I think it is A because I don't understand the need to be restricted to only entries with TTL.

upvoted 2 times

adilkhan 9 months ago

B looks correct b/c of log out condition.

upvoted 1 times

alexein74 9 months, 1 week ago

Selected Answer: B

B. volatile-lru

The eviction policy "volatile-lru" evicts the least recently used (LRU) key among keys with an expiring time-to-live (TTL). This policy is suitable for your scenario because it prioritizes players based on how recently they have moved and does not prioritize players who have logged out of the game.

upvoted 1 times

Osip21 9 months, 1 week ago

Selected Answer: D

"volatile-ttl" is a type of eviction policy that uses a time-to-live (TTL) value for each cache entry. This policy automatically evicts entries from the cache after their TTL expires. This can be useful for scenarios where you want to ensure that stale or outdated data is not kept in the cache for too long. In your case, you can set a shorter TTL for players who have moved recently and a longer TTL for players who have not moved recently or logged out of the game, which allow the system to prioritize players based on their activity.

upvoted 2 times

JustHereToLearn 9 months, 1 week ago

Selected Answer: B

I think the answer should be B based on the facts. "The system should prioritize players based on how recently they have moved" this means we will use LRU but not we have to decide between allkeys-lru or valatile-lru.
The statement "should not prioritize the players who logged out of game" provides a hint that there needs to be some to evict logged out users like using the expire to true.
Reference: https://redis.io/docs/reference/eviction/

upvoted 2 times

Question #34 *Topic 5*

You develop an Azure App Service web app and deploy to a production environment. You enable Application Insights for the web app.

The web app is throwing multiple exceptions in the environment.

You need to examine the state of the source code and variables when the exceptions are thrown.

Which Application Insights feature should you configure?

A. Smart detection

B. Profiler

C. Snapshot Debugger

D. Standard test

Correct Answer: *C*

Community vote distribution

C (100%)

alexein74 Highly Voted 9 months, 1 week ago

Selected Answer: C

C. Snapshot Debugger

The Snapshot Debugger is an Application Insights feature that allows you to examine the state of the source code and variables when an exception is thrown. It captures a snapshot of the call stack, local variables, and object state at the time of the exception. With this feature, you can easily identify the cause of the exception and fix the issue.

upvoted 6 times

Jhilphis Highly Voted 9 months, 3 weeks ago

Selected Answer: C

Correct answer

upvoted 5 times

Ciupaz Most Recent 1 month ago

The Snapshot Debugger in Azure is a debugging tool and feature that helps developers diagnose and troubleshoot issues in their applications running on Azure App Service without impacting the application's availability or performance. It allows you to capture snapshots of the application's state at specific points in time when an issue or exception occurs, allowing you to inspect the state of variables and the call stack to identify the root cause of problems.

upvoted 1 times

Question #35 *Topic 5*

DRAG DROP
-

You develop and deploy a Java application to Azure. The application has been instrumented by using the Application Insights SDK.

The telemetry data must be enriched and processed before it is sent to the Application Insights service.

You need to modify the telemetry data.

Which Application Insights SDK 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.

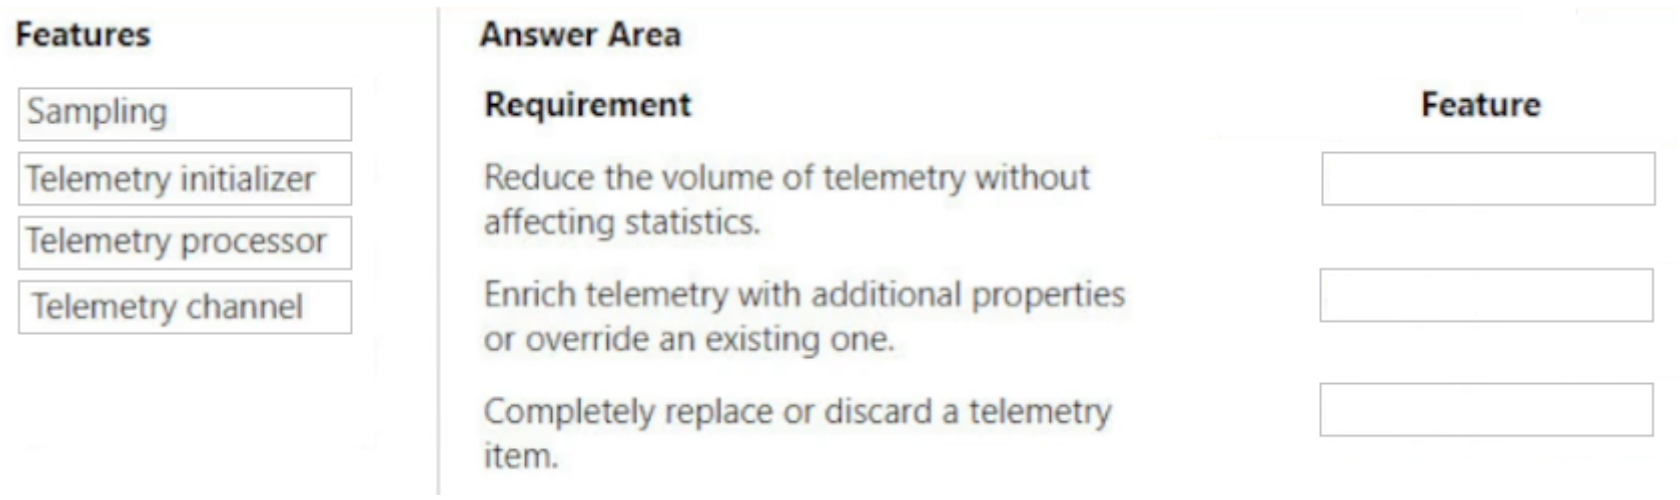

Correct Answer:

Answer Area

Requirement	Feature
Reduce the volume of telemetry without affecting statistics.	Sampling
Enrich telemetry with additional properties or override an existing one.	Telemetry initializer
Completely replace or discard a telemetry item.	Telemetry processor

whiteblack Highly Voted 9 months, 3 weeks ago

Correct answers
Link: https://learn.microsoft.com/en-us/azure/azure-monitor/app/api-filtering-sampling

upvoted 18 times

Rodikito Most Recent 2 months, 1 week ago

what is 'Case study: city and Lights'??

upvoted 1 times

Rodikito 2 months ago

sorry, I din´t read all exam topics (I didn´t have contributor yet). Case studies are in the lasts pages of this dump

upvoted 1 times

kvtborad 2 months, 3 weeks ago

I got this question on 6th August 2023. chose given answers. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.

upvoted 3 times

WH16 2 months, 3 weeks ago

Hi bro, how many questions in case study?

upvoted 2 times

Rodikito 2 months ago

Normally, 1 case study per exam , with 5 questions

upvoted 3 times

macobuzi 1 month, 3 weeks ago

5 - 12 question

upvoted 2 times

Question #36 *Topic 5*

HOTSPOT
-

You develop new functionality in a web application for a company that provides access to seismic data from around the world. The seismic data is stored in Redis Streams within an Azure Cache for Redis instance.

The new functionality includes a real-time display of seismic events as they occur.

You need to implement the Azure Cache for Redis command to receive seismic data.

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

NOTE: Each correct selection is worth one point.

Answer Area

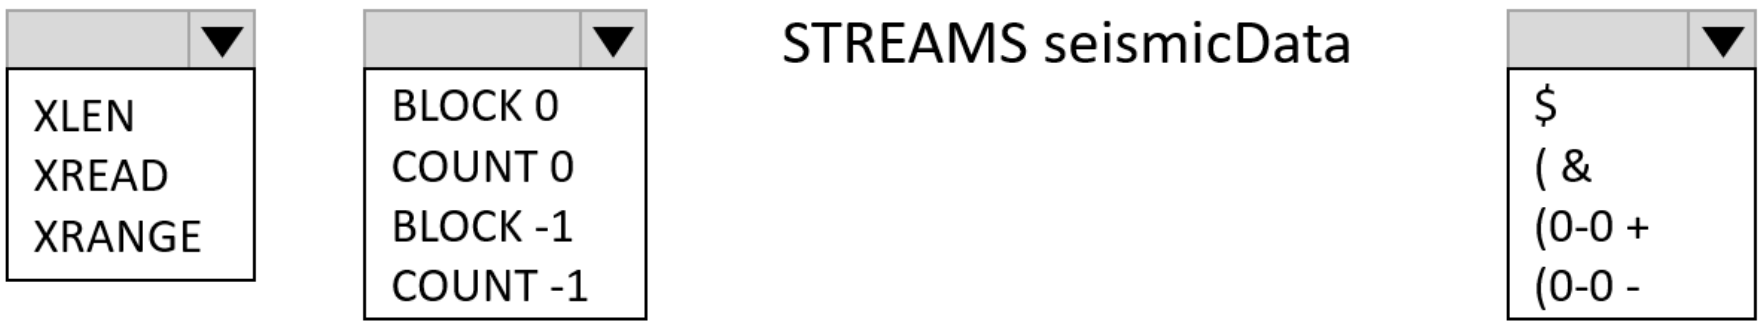

Correct Answer:

Answer Area

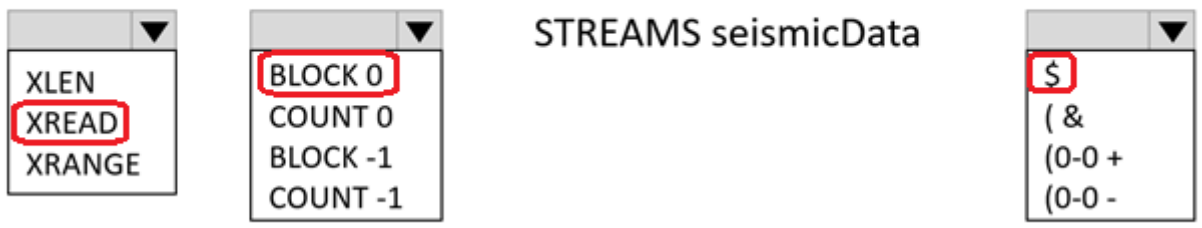

Michael2023 Highly Voted 9 months ago
puuh... knowing all this little things/ remembering all this little things for the exam sounds impossible :(
upvoted 45 times

macobuzi 2 months ago
Exactly! That's why we study the exam in question bank :)
upvoted 6 times

Ciupaz 1 month ago
Agree!
upvoted 2 times

2fa Highly Voted 9 months, 3 weeks ago
Correct.
BLOCK 0 means that we will wait infinitely (timeout 0 means it will never expire) for new items in a stream.
$ ID means that we will receive only new messages, starting from the time we started listening

https://redis.io/docs/data-types/streams-tutorial/#listening-for-new-items-with-xread
upvoted 15 times

dddddd111 Most Recent 1 week, 6 days ago
I failed on my first attempt. I only got 646/1000. I chose C#. MS has recently updated the AZ-204 (Oct 2023), so please be guided that some highly voted answers here might not be correct.
upvoted 2 times

AndySmith 5 days, 12 hours ago

Did they add any new Qs outside of this "tutorial"?
Or they just updated existing Qs - e.g. answer options, etc.

upvoted 1 times

BaoNguyen2411 3 months ago

got this question on 29/06/2023

upvoted 5 times

brwilli1 3 months, 1 week ago

This was on the test 07/14/2023

upvoted 4 times

KarthikCS 6 months, 1 week ago

These AZ exam questions are some of the most ridiculous ones ever. Looks like MS wants you to fail unless you senselessly memorize stupid things. I wonder how dumb the team that sets these questions must be.

upvoted 12 times

surprise0011 6 months, 3 weeks ago

wtf is that haha

upvoted 8 times

mmdex 8 months ago

This question is so ridiculous that I actually remember its answer better than any other.

upvoted 14 times

uffuchsi 8 months, 1 week ago

Agree with the others here - seems quite silly to be expected to remember detail like this.

Would be interested to know if anyone recieved this question in their exam

upvoted 4 times

NombreFalso 8 months, 1 week ago

Feels like anything goes and they can ask you about anything that's ever existed

upvoted 10 times

Question #37 *Topic 5*

You develop an ASP.NET Core app that uses Azure App Configuration. You also create an App Configuration containing 100 settings.

The app must meet the following requirements:

- Ensure the consistency of all configuration data when changes to individual settings occur.
- Handle configuration data changes dynamically without causing the application to restart.
- Reduce the overall number of requests made to App Configuration APIs.

You must implement dynamic configuration updates in the app.

What are two ways to achieve this goal? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Create and register a sentinel key in the App Configuration store. Set the refreshAll parameter of the Register method to true.

B. Increase the App Configuration cache expiration from the default value.

C. Decrease the App Configuration cache expiration from the default value.

D. Create and configure Azure Key Vault. Implement the Azure Key Vault configuration provider.

E. Register all keys in the App Configuration store. Set the refreshAll parameter of the Register method to false.

F. Create and implement environment variables for each App Configuration store setting.

Correct Answer: *AB*

Community vote distribution

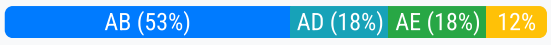

whiteblack Highly Voted 9 months, 3 weeks ago

Correct Answers
Link: https://learn.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-netfx

upvoted 8 times

halfway Highly Voted 7 months, 1 week ago

Selected Answer: AB

Quotes from the reference below:
"A sentinel key is a key that you update after you complete the change of all other keys. Your app monitors the sentinel key. When a change is detected, your app refreshes all configuration values. This approach helps to ensure the consistency of configuration in your app and reduces the overall number of requests made to your App Configuration store, compared to monitoring all keys for changes."

"You can add a call to the refreshOptions.SetCacheExpiration method to specify the minimum time between configuration refreshes. In this example, you use the default value of 30 seconds. Adjust to a higher value if you need to reduce the number of requests made to your App Configuration store."

Reference: https://learn.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core6x

upvoted 5 times

AndySmith Most Recent 4 days, 17 hours ago

It's AB for sure, since it fully follows the best practices of AzureAppConfig service:
https://learn.microsoft.com/en-us/azure/azure-app-configuration/howto-best-practices#reduce-requests-made-to-app-configuration

It can't be D since there is no any word about a secret in the Q.

upvoted 1 times

dddddd111 1 week, 6 days ago

Selected Answer: AD

Option B. Increase the App Configuration cache expiration from the default value is a valid consideration for reducing the number of requests made to App Configuration APIs, but it does not directly address the requirement of dynamic configuration updates without causing the application to restart.

In fact, the default behavior of Azure App Configuration already includes caching to minimize the number of requests to the configuration store. Adjusting the cache expiration might help reduce the frequency of requests, but it won't handle dynamic updates without application restarts.

For achieving dynamic configuration updates without restarting the application, options A (sentinel key) and D (Azure Key Vault configuration

provider) are more relevant. These options allow you to implement dynamic configuration changes and ensure consistency without requiring application restarts.

upvoted 1 times

drkpnsh 2 months ago

Selected Answer: AD

I had to go for A & D.

For B "You can add a call to the refreshOptions.SetCacheExpiration method to specify the minimum time between configuration refreshes. In this example, you use the default value of 30 seconds. Adjust to a higher value if you need to reduce the number of requests made to your App Configuration store." explanation makes no sense for me.

upvoted 1 times

Arr14487 2 months, 3 weeks ago

Selected Answer: AF

Option B, Increase the App Configuration cache expiration from the default value, is not a good solution for dynamic configuration updates.

upvoted 2 times

adilkhan 6 months, 1 week ago

A,B 100% correct answer!

upvoted 2 times

adilkhan 6 months, 1 week ago

B makes no sense A and D are correct!

upvoted 1 times

Dats1987 6 months, 2 weeks ago

Selected Answer: AB

AB , looks correct.

upvoted 2 times

adilkhan 8 months ago

A. Create and register a sentinel key in the App Configuration store. Set the refreshAll parameter of the Register method to true.
D. Create and configure Azure Key Vault. Implement the Azure Key Vault configuration provider. Chat GPT
B n C are related to Cache and has nothing to do with the question. Chat GPT

upvoted 1 times

CODE_STS 8 months ago

Selected Answer: AD

A. Create and register a sentinel key in the App Configuration store. Set the refreshAll parameter of the Register method to true: By creating a sentinel key in the App Configuration store, you can ensure that all configuration data remains consistent when changes to individual settings occur. You can register the sentinel key and set the refreshAll parameter of the Register method to true to refresh all settings in the App Configuration store.

D. Create and configure Azure Key Vault. Implement the Azure Key Vault configuration provider: The Azure Key Vault configuration provider enables you to store configuration data in Azure Key Vault. By using the Azure Key Vault configuration provider, you can implement dynamic configuration updates in your app without causing the app to restart. This approach also reduces the overall number of requests made to the App Configuration APIs.

upvoted 1 times

Sriniv 8 months, 2 weeks ago

The two ways to achieve dynamic configuration updates in an ASP.NET Core app that uses Azure App Configuration are:

A. Create and register a sentinel key in the App Configuration store. Set the refreshAll parameter of the Register method to true.

D. Create and configure Azure Key Vault. Implement the Azure Key Vault configuration provider.

upvoted 1 times

SEEKER31 8 months, 3 weeks ago

Selected Answer: AB

correct answer

upvoted 1 times

OlavG 9 months ago

Selected Answer: AB

Correct

upvoted 1 times

adilkhan 9 months ago

AB is correct!
https://learn.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core6x

upvoted 3 times

alexein74 9 months, 1 week ago

Selected Answer: AE

A. Create and register a sentinel key in the App Configuration store. Set the refreshAll parameter of the Register method to true.
E. Register all keys in the App Configuration store. Set the refreshAll parameter of the Register method to false.

upvoted 3 times

alexein74 9 months, 1 week ago

Increasing or decreasing the App Configuration cache expiration from the default value will not achieve the goal because it will not handle configuration data changes dynamically without causing the application to restart. Creating and configuring Azure Key Vault and implementing the Azure Key Vault configuration provider will not achieve the goal because it will not handle configuration data changes dynamically without causing the application to restart. Creating and implementing environment variables for each App Configuration store setting will not achieve the goal because it will not handle configuration data changes dynamically without causing the application to restart.

upvoted 2 times

alexein74 9 months, 1 week ago

Create and register a sentinel key in the App Configuration store. Set the refreshAll parameter of the Register method to true: Sentinel key is a special configuration key that you can use to control the refreshing of all other configuration keys. When a sentinel key changes, the Azure App Configuration service updates all other keys, by setting the refreshAll parameter of the Register method to true, this way you can ensure the consistency of all configuration data when changes to individual settings occur.

upvoted 2 times

adilkhan 9 months ago

A and B is correct. For B check this
You can add a call to the refreshOptions.SetCacheExpiration method to specify the minimum time between configuration refreshes. In this example, you use the default value of 30 seconds. Adjust to a higher value if you need to reduce the number of requests made to your App Configuration store.
https://learn.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core6x

upvoted 1 times

alexein74 9 months, 1 week ago

Register all keys in the App Configuration store. Set the refreshAll parameter of the Register method to false: By registering all keys in the App Configuration store, you can handle configuration data changes dynamically without causing the application to restart, by setting the refreshAll parameter of the Register method to false, you can reduce the overall number of requests made to App Configuration APIs.

upvoted 2 times

Question #38 *Topic 5*

HOTSPOT
-

You develop and deploy an Azure App Service web app that connects to Azure Cache for Redis as a content cache. All resources have been deployed to the East US 2 region.

The security team requires the following audit information from Azure Cache for Redis:

• The number of Redis client connections from an associated IP address.
• Redis operations completed on the content cache.
• The location (region) in which the Azure Cach3e for Redis instance was accessed.

The audit information must be captured and analyzed by a security team application deployed to the Central US region.

You need to log information on all client connections to the cache.

Which configuration values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Configuration value
Store log information.	Log Analytics workspace Blob Storage account Data Lake Storage Gen2 Storage account Event hub
Enable client connection logging.	Diagnostic setting Managed identity App registration Environment variable

Correct Answer:

Answer Area

Requirement	Configuration value
Store log information.	Log Analytics workspace Blob Storage account Data Lake Storage Gen2 Storage account Event hub
Enable client connection logging.	Diagnostic setting Managed identity App registration Environment variable

halfway Highly Voted 6 months, 2 weeks ago

Correct.
Only Log Analytics supports cross-region logging

Reference: https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-monitor-diagnostic-settings?tabs=basic-standard-premium#enable-connection-logging-using-the-azure-portal

upvoted 14 times

kakajos Most Recent 1 month, 1 week ago

Log Analytics Workspace - used to store log information.
Diagnostic setting - diagnostic logs are streamed to that workspace as soon as new event data is generated

upvoted 1 times

Question #39 *Topic 5*

You develop and deploy a web app to Azure App Service. The Azure App Service uses a Basic plan in a single region.

Users report that the web app is responding slow. You must capture the complete call stack to help identify performance issues in the code. Call stack data must be correlated across app instances. You must minimize cost and impact to users on the web app.

You need to capture the telemetry.

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

NOTE: Each correct selection is worth one point.

A. Restart all apps in the App Service plan.

B. Enable Application Insights site extensions.

C. Upgrade the Azure App Service plan to Premium.

D. Enable Profiler.

E. Enable the Always On setting for the app service.

F. Enable Snapshot debugger.

G. Enable remote debugging.

Correct Answer: *DEF*

Community vote distribution

BDF (50%) BDE (50%)

paunski7 Highly Voted 6 months, 1 week ago

B E F

B. Enable Application Insights site extensions: Application Insights provides end-to-end insights into the web app's performance and usage. Enabling the Application Insights site extension allows you to collect telemetry data, including call stack traces, to help identify performance issues in the code.

E. Enable the Always On setting for the app service: Enabling the Always On setting ensures that the web app is always running, even when there is no traffic. This reduces the impact of cold starts and improves the web app's response time.

F. Enable Snapshot debugger: Enabling Snapshot debugger allows you to capture the complete call stack and local variables for an exception that occurred in the web app. This helps to identify performance issues and diagnose issues in the code.

upvoted 13 times

JohnStivenson 1 month ago

E - Enable the Always On, how this will fulfill "You need to capture the telemetry."?

upvoted 1 times

JohnStivenson 1 month ago

I think, it's BDF:
- Enable Application Insights site extensions
- Enable Profiler
- Enable Snapshot debugger

upvoted 2 times

paunski7 6 months, 1 week ago

Option A, restarting all apps in the App Service plan, and option G, enabling remote debugging, are not necessary for capturing telemetry and identifying performance issues in the web app.

Option C, upgrading the Azure App Service plan to Premium, is not necessary for capturing telemetry and identifying performance issues in the web app, and it would significantly increase the cost of the App Service plan.

Option D, enabling Profiler, is not necessary as it is an outdated feature and has been replaced by the more powerful and comprehensive Snapshot debugger.

upvoted 5 times

damianadalid 5 months, 1 week ago

Hello, could you please provide the source where you found that Profiler is being replaced by Snapshot debugger?

upvoted 3 times

Ciupaz Most Recent 1 month ago

Selected Answer: BDF

For me:
- Enable Application Insights site extensions
- Enable Profiler
- Enable Snapshot debugger

upvoted 1 times

kakajos 1 month, 1 week ago

BDE is the correct answer. Profiler capture data automatically without negatively affecting users.

upvoted 1 times

Paolo12 2 months ago

Always On is available from Standard tier

upvoted 2 times

EliteAllen 2 months, 2 weeks ago

Selected Answer: BDF

B. Enable Application Insights site extensions.

This will allow you to collect telemetry from your web app, including performance metrics, failed requests, and dependencies.
D. Enable Profiler.

The Application Insights Profiler will help you identify performance issues in your app's code by capturing telemetry about the app's slowest operations.
F. Enable Snapshot debugger.

The Snapshot Debugger takes snapshots of your in-production apps when exceptions occur, allowing you to see exactly what went wrong without affecting traffic of your production application.

upvoted 2 times

Dianahu 3 months ago

D) Enable profiler
With Application Insights Profiler, you can capture and view performance traces for your application in all these dynamic situations. The process occurs automatically at scale and doesn't negatively affect your users. Profiler captures the following information so that you can easily identify performance issues while your app is running in Azure (https://learn.microsoft.com/en-us/azure/azure-monitor/profiler/profiler-overview)

E)enable always on (to enable profiler -> https://learn.microsoft.com/en-us/azure/azure-monitor/profiler/profiler)

F)Snapshot debugger - to collect call stacks for your application - https://learn.microsoft.com/en-us/azure/azure-monitor/snapshot-debugger/media/snapshot-debugger/enablement-ui.png

upvoted 3 times

JH81 4 months, 1 week ago

Selected Answer: BDE

B, D, E

Gotta enable App Insights.
Always On can solve the performance issues.
Profiler can help find performance issues.
You can eliminate Premium upgrade because it says minimize costs, snapshot debugger is for debugging exceptions and has nothing to do with performance issues, restarting app impacts users and it says minimize impact.

upvoted 3 times

the_odd_one 1 week ago

Dosent say anything about needing to solve the performance issue, for me it is B, D, F

upvoted 1 times

Dreamer999 4 months, 2 weeks ago

I think "Always On" is an option to fix the issue. But the question is about collecting metrics. So I think "Always On" should not be selected. I would choose Profiler, Snapshot debugger, Enable application Insights site extensions

upvoted 2 times

damianadalid 5 months, 1 week ago

I think it should be B,E,D :
-B and E make the system able to use Application Insights via setting the App Serviced to always On and enabling Application Insights extensions.
-D: Profiler is the most suitable one regarding finding performance issues instead of Snapshot debugger, which is more related to exceptions and source code state when errors.

upvoted 3 times

dy0917 5 months, 2 weeks ago

Should be B D E
F. Enable Snapshot debugger for debugging purposes.
D. Profiler for monitoring performance.

upvoted 4 times

sachinrikhe 6 months, 1 week ago

This is new question. It was not here yesterday

upvoted 2 times

adilkhan 6 months, 1 week ago

answers look correct!

upvoted 1 times

Question #40 *Topic 5*

You are building an application to track cell towers that are available to phones in near real time. A phone will send information to the application by using the Azure Web PubSub service. The data will be processed by using an Azure Functions app. Traffic will be transmitted by using a content delivery network (CDN).

The Azure function must be protected against misconfigured or unauthorized invocations.

You need to ensure that the CDN allows for the Azure function protection.

Which HTTP header should be on the allowed list?

A. Authorization

B. WebHook-Request-Callback

C. Resource

D. WebHook-Request-Origin

Correct Answer: *D*

Community vote distribution

D (67%) A (33%)

halfway Highly Voted 6 months, 2 weeks ago

Selected Answer: D

https://learn.microsoft.com/en-us/azure/azure-web-pubsub/howto-develop-eventhandler#upstream-and-validation

upvoted 8 times

AndySmith Most Recent 4 days, 16 hours ago

Most likely, it's A, since they ask about "Unauthorized protection" instead of "Abuse Protection", which are both supported by underlying "CloudEvents HTTP protocols" (web-hooks).
https://github.com/cloudevents/spec/blob/v1.0/http-webhook.md#4-abuse-protection

So, if we read attentively about Abuse Protection / webhook-request-origin, we can find this excerpt:
"It is important to understand is that the handshake does not aim to establish an authentication or authorization context. It only serves to protect the sender from being told to a push to a destination that is not expecting the traffic. While this specification mandates use of an authorization model, this mandate is not sufficient to protect any arbitrary website from unwanted traffic if that website doesn't implement access control and therefore ignores the Authorization header."

And this is Authorization header description in this protocol:
https://github.com/cloudevents/spec/blob/v1.0/http-webhook.md#3-authorization

upvoted 1 times

ddddd111 1 week, 6 days ago

Selected Answer: A

The Authorization header (Option A) is widely used for authentication in HTTP requests. It is a standard header for including credentials, tokens, or other authentication information. By allowing the Authorization header on the CDN, you can ensure that only requests with valid authorization tokens or credentials can invoke your Azure function, providing a secure method of protection against unauthorized invocations.

upvoted 2 times

Ciupaz 1 month, 1 week ago

Selected Answer: A

Why not A?
To ensure that the Azure Function is protected against misconfigured or unauthorized invocations when using a CDN, you should allow the "Authorization" HTTP header.

upvoted 2 times

JH81 4 months, 1 week ago

This cannot possibly be a real question the way it is written.

upvoted 2 times

adilkhan 6 months, 1 week ago

B. Enable Application Insights site extensions.
E. Enable the Always On setting for the app service.
D. Enable Profiler. chat gpt

upvoted 1 times

damianadalid 5 months, 1 week ago

wrong questions, this relates to #39

upvoted 3 times

adilkhan 6 months, 1 week ago

D is correct

upvoted 3 times

Question #41 *Topic 5*

You are developing an Azure App Service web app.

The web app must securely store session information in Azure Redis Cache.

You need to connect the web app to Azure Redis Cache.

Which three Azure Redis Cache properties should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Access key

B. SSL port

C. Subscription name

D. Location

E. Host name

F. Subscription id

Correct Answer: *ABE*

anupmaurya 1 month, 1 week ago

When connecting a web app to Azure Redis Cache, you need to specify certain properties. Here are three Azure Redis Cache properties you should use:

Hostname: This property represents the fully qualified domain name (FQDN) or the endpoint URL of the Azure Redis Cache instance. It is used to establish a connection between the web app and the cache. The hostname typically follows the format yourcache.redis.cache.windows.net.

Port: This property specifies the port number used for communication with the Azure Redis Cache instance. The default port for Redis is 6379, but you can configure it differently if needed.

Access key: This property is the primary or secondary access key for the Azure Redis Cache instance.

upvoted 3 times

AbidooKing 3 months, 3 weeks ago

Answer is correct.

https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-howto#retrieve-host-name-ports-and-access-keys-from-the-azure-portal

upvoted 3 times

laurorucker 3 months, 3 weeks ago

When connecting a web app to Azure Redis Cache, you need to specify certain properties. Here are three Azure Redis Cache properties you should use:

Hostname: This property represents the fully qualified domain name (FQDN) or the endpoint URL of the Azure Redis Cache instance. It is used to establish a connection between the web app and the cache. The hostname typically follows the format yourcache.redis.cache.windows.net.

Port: This property specifies the port number used for communication with the Azure Redis Cache instance. The default port for Redis is 6379, but you can configure it differently if needed.

Access key: This property is the primary or secondary access key for the Azure Redis Cache instance. It acts as a password for authentication purposes. The access key is used by the web app to authorize access to the cache and securely store session information.

upvoted 1 times

AbidooKing 3 months, 3 weeks ago

At least let us know when you are posting chatGPT-shit here. This time is seems to be correct.
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-howto#retrieve-host-name-ports-and-access-keys-from-the-azure-portal

upvoted 3 times

Question #42 *Topic 5*

HOTSPOT
-

You are developing several microservices to run on Azure Container Apps.

You need to monitor and diagnose the microservices.

Which features should you use? To answer, select the appropriate feature in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Feature
View console logs from a container in near real-time.	Log streaming / Container console / Azure Monitor metrics / Azure Monitor Log Analytics
Debug the microservice from inside the container.	Container console / Azure Monitor metrics / Azure Container Registry / Azure Monitor Log Analytics

Correct Answer:

Answer Area

Requirement	Feature
View console logs from a container in near real-time.	**Log streaming** / Container console / Azure Monitor metrics / Azure Monitor Log Analytics
Debug the microservice from inside the container.	**Container console** / Azure Monitor metrics / Azure Container Registry / Azure Monitor Log Analytics

jaf19f Highly Voted 2 months, 2 weeks ago
I got this question (12-Aug-2023) and I chose the given answers - 932 passed
upvoted 7 times

AbidooKing Highly Voted 3 months, 3 weeks ago
Answer is correct!

https://learn.microsoft.com/en-us/azure/container-apps/observability
upvoted 5 times

AndySmith 1 week, 4 days ago
From this page:
Log Streaming - View streaming system and console logs from a container in near real-time.
Container console - Connect to the Linux console in your containers to debug your application from inside the container.
upvoted 1 times

75daltonic 3 months, 1 week ago
Correct.
It is not deployed yet. So:

Development and test

During the development and test phase, real-time access to your containers' application logs and console is critical for debugging issues. Container Apps provides:

Log streaming: View real-time log streams from your containers.
Container console: Access the container console to debug your application.

https://learn.microsoft.com/en-us/azure/container-apps/observability#development-and-test

upvoted 3 times

laurorucker Most Recent 3 months, 3 weeks ago

Azure Monitor Logs: Azure Monitor Logs, also known as Azure Log Analytics, allows you to collect, analyze, and view log data from various sources, including containerized applications. By configuring Azure Monitor Logs for your microservices, you can stream and view console logs in near real-time. You can create log queries and use the Live Stream feature to monitor logs as they are generated, enabling you to observe the output from the console logs in real-time.

upvoted 1 times

laurorucker 3 months, 3 weeks ago

For me both would be azure monitor logs

upvoted 1 times

Case study -

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

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

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

To start the case study -
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

• Secure the website by using SSL.
• Minimize costs for data storage and hosting.
• Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).
• Distribute the website content globally for local use.
• Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.
• The website must have 99.95 percent uptime.

Retail store locations -

• Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.
• Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

• Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.
• Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -

• All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
• Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -

• You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
• Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

You need to test the availability of the corporate website.

Which two test types can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Standard

B. URL ping

C. Custom testing using the TrackAvailability API method

D. Multi-step

Correct Answer: *AC*

Community vote distribution

AC (100%)

AndySmith 4 days, 16 hours ago

A and C per this requirement: "Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification."

A - Standard test: This single request test is similar to the URL ping test. It includes TLS/SSL certificate validity, proactive lifetime check, HTTP request verb.

C - Custom TrackAvailability test: If you decide to create a custom application to run availability tests, you can use the TrackAvailability() method to send the results to Application Insights.

https://learn.microsoft.com/en-us/azure/azure-monitor/app/availability-overview#types-of-tests

upvoted 1 times

ENGs 1 month ago

Selected Answer: AC

A and C

upvoted 1 times

drkpnsh 2 months, 1 week ago

Because it is for the local use, my understanding is it is in the intranet, not internet. So url ping & multi step cannot be used.

https://learn.microsoft.com/en-us/azure/azure-monitor/app/availability-overview#types-of-tests

upvoted 1 times

lorenaizzo 2 months, 1 week ago

The answer must meet the following rule "• Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification."
So Standard Test can check if the certificate is valid. And creating a custom application and sending the results using a Custom TrackAvailabilityTest.

About the local site vs internet...
"The older classic tests, URL ping test and multi-step web test, rely on the DNS infrastructure of the public internet to resolve the domain names of the tested endpoints. If you're using private DNS, you must ensure that the public domain name servers can resolve every domain name of your test. When that's not possible, you can use custom TrackAvailability tests instead."
I understand that you can use any of them, just need to pay attention to the DNS entries (and routing)

upvoted 1 times

lyggwtwtczxnhznebw 2 months, 2 weeks ago

Selected Answer: AC

Based on https://learn.microsoft.com/en-us/azure/azure-monitor/app/availability-overview#types-of-tests the answer is correct.

upvoted 3 times

Arr14487 2 months, 3 weeks ago

given answer , i think its wrong.
it should be urlping and multi step

upvoted 2 times

Firo 3 months ago

Selected Answer: AC

The answer is correct
A-Standard
C-Custom Testing - TrackAvailability
https://learn.microsoft.com/en-us/azure/azure-monitor/app/availability-overview#types-of-tests

upvoted 4 times

Question #44 *Topic 5*

You have an Azure API Management (APIM) Standard tier instance named APIM1 that uses a managed gateway.

You plan to use APIM1 to publish an API named API1 that uses a backend database that supports only a limited volume of requests per minute. You also need a policy for API1 that will minimize the possibility that the number of requests to the backend database from an individual IP address you specify exceeds the supported limit.

You need to identify a policy for API1 that will meet the requirements.

Which policy should you use?

A. ip-filter

B. quota-by-key

C. rate-limit-by-key

D. rate-limit

Correct Answer: *C*

Community vote distribution

C (100%)

ddddd111 1 week, 6 days ago

Selected Answer: C

C is correct.

upvoted 1 times

Ciupaz 1 month, 1 week ago

Selected Answer: C

The "rate-limit-by-key" policy in Azure API Management allows you to apply rate limiting based on a specific key, which can be an IP address or any other attribute that you want to use for rate limiting.

upvoted 2 times

macobuzi 2 months ago

Selected Answer: C

C is correct:
Rate limits are usually used to protect against short and intense volume bursts. For example, if you know your backend service has a bottleneck at its database with a high call volume, you could set a rate-limit-by-key policy to not allow high call volume by using this setting.
https://learn.microsoft.com/en-us/azure/api-management/api-management-sample-flexible-throttling#rate-limits-and-quotas

upvoted 2 times

drogbix 2 months, 3 weeks ago

Selected Answer: C

Correct Answer.

https://learn.microsoft.com/en-us/azure/api-management/api-management-policies

upvoted 1 times

Topic 6 - Question Set 6

Question #1 *Topic 6*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Service Bus. Configure a topic to receive the device data by using a correlation filter.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*

A message is raw data produced by a service to be consumed or stored elsewhere. The Service Bus is for high-value enterprise messaging, and is used for order processing and financial transactions.

Reference:

https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

Community vote distribution

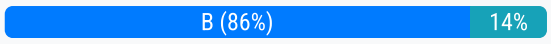

Spooky7 Highly Voted 2 years, 5 months ago

First of all - that question is not about THE BEST solution but about THE VALID solution. So responses like "Event Hub is the best choice" doesn't bring anything to the table.
That being said - is the provided solution valid? I would say NO, because:
- topics allows multiple subscribers, and here we need to process each event once
- correlation filter is for subscriptions, not topics
- even when assuming there is typo in the question and correlation filter is defined on the subscription level - it still is not a valid solution, because new stores can be opened in the future with many new device identifiers which you can't know in advance. Besides that filter make no sense in this scenario whatsoever, you just need to save data in storage account and basically partition it by device identifier.

upvoted 44 times

kabs Highly Voted 3 years, 3 months ago

Answer should be event hub

upvoted 27 times

Tmm58 2 years, 11 months ago

Why Event Hub? This is clearly about delivering messages (with the actual data) and not events.

upvoted 23 times

xofowi5140 3 years, 3 months ago

The event hub message size limit is 1MB.

upvoted 26 times

TalesFromEarth 3 years, 3 months ago

So is in service hub. Maximum message size: 256 KB for Standard tier, 1 MB for Premium tier.

upvoted 3 times

Dirk 3 years, 3 months ago

256 KB is the max size of the message queue, But this is not about the messages delivered per 24 hours.
https://stackoverflow.com/questions/22384193/azure-service-bus-message-size-technology-limit-and-pricing/22384741#:~:text=1%20Answer&text=Service%20Bus%20queues%20support%20a,messages%20held%20by%20a%20queue.
In the question it's about the size in 24 hours.

upvoted 4 times

MaheshBeeravelli 3 years, 3 months ago

I agree with xofowi5140. Given answer is correct.

upvoted 13 times

Adol 2 years, 9 months ago

Message size of 1MB has no conflict with 2MB of messages in 24 hours! So still Event hub when it comes to IoT or when many devices sending data with device id as identifier

upvoted 5 times

profesorklaus 2 years, 9 months ago

Of course it has a conflict. One store can send message as a batch with size 2MB.

upvoted 8 times

dddddd111 Most Recent 1 week, 6 days ago

Selected Answer: B

Answer is No.

upvoted 1 times

lugia4000 1 month ago

These question series came out on my exam 2023-09-29

upvoted 2 times

adilkhan 8 months ago

One approach could be to use Azure IoT Hub to receive the data from the POS devices, and then use Azure Functions or Azure Stream Analytics to process and store the data in Azure Blob storage based on the device identifier. ChatGPT

upvoted 1 times

Michael2023 9 months ago

Events vs Message
This is clearly a message (2MB) from multiple devices which needs to be stored in Azure Blob Storage... that means we need to decide between Azure Service Bus or Azure Queue Storage queues (if available option) because these are message-based solutions...

for me Azure Service Bus sounds like a good fit... take a look at the comparison page of these Azure Services

https://learn.microsoft.com/en-us/azure/event-grid/compare-messaging-services

upvoted 1 times

adilkhan 9 months ago

Here we are using Azure Blob storage. We store the data in the blob storage and name of the blob we store in the Azure service bus to process it but we don't need topic here.

upvoted 2 times

rolling_potato_ 10 months ago

This series of questions I got on my exam today (03/01/2023) I chose no for Service Bus and Event Grid and Yes for Event HUb (score 921)

upvoted 8 times

OPT_001122 10 months, 3 weeks ago

Selected Answer: B

Coorect ans - event hub (There are multiple questions with different answers)
Event Hub for this scenario - 2000 POS devices

upvoted 1 times

hubekpeter 11 months, 1 week ago

Selected Answer: A

I vote for A. It's POS terminal. We are talking about financial transactions that needs to be reliably delivered and processed at least once. All the limitations mentioned below can be easily overcome. Azure service - bus service is intended for enterprise applications that require transactions, ordering, duplicate detection, and instantaneous consistency.

upvoted 2 times

somename20221106 11 months, 2 weeks ago

Answer will be Yes - As Service Bus is the recommended for order processing and Financials
POS - is order process/finance.
https://learn.microsoft.com/en-us/azure/event-grid/compare-messaging-services

upvoted 1 times

Knightie 1 year ago

remember service bus is an always correct answer in all of the event and messaging scenario.

upvoted 2 times

Knightie 1 year ago

Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties. A common use is to match against the CorrelationId property, but the application can also choose to match against the following properties:

ContentType
Label
MessageId
ReplyTo
ReplyToSessionId
SessionId

To
any user-defined properties.
A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter. For string expressions, the comparison is case-sensitive. If you specify multiple match properties, the filter combines them as a logical AND condition, meaning for the filter to match, all conditions must match.

Answer A, nothing to stop us from grouping up the topics and you can have whosoever interested to subscribe to the message of the grouped devices.

upvoted 1 times

Knightie 1 year ago

Selected Answer: A

Yes, it can be used so... even if you use the correlation id as topic, the point is. it is still workable, right? who knows how the pub/sub model they wanna be? may be the finance, boss and account wanna be the 3 subscribers for all different topics? unless there is a limiation of topics... otherwise, nothing is stopping the correlation_id to be the topic.

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.
Event Hub would be a solution here NOT Service Bus.

upvoted 2 times

Netspud 1 year, 8 months ago

Selected Answer: B

Answer is no.

upvoted 3 times

leonidn 1 year, 9 months ago

Selected Answer: B

Filters can be applied to a subscription level. Each topic is limited to 2000 subscriptions. Since we are expecting to exceed that value in the future, then we need to have >2000 filters, hence > 2000 subscriptions. Consequently, that is not a solution.

upvoted 11 times

ReniRechner 1 year, 8 months ago

they have 2000+ producers not 2000+ consumers.

upvoted 8 times

Question #2 *Topic 6*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Event Grid. Configure event filtering to evaluate the device identifier.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Instead use an Azure Service Bus, which is used order processing and financial transactions.

Note: An event is a lightweight notification of a condition or a state change. Event hubs is usually used reacting to status changes.

Reference:

https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

Community vote distribution

B (80%) A (20%)

xRiot007 Highly Voted 3 years, 4 months ago

For the ones above me. Sensors do not send events, they send messages containing specific data that has been gathered. This makes automatically the solution incorrect, because you need a Service Bus to collect them. Event Grids and Event Hubs won't do the job here.

upvoted 44 times

muggins 1 year, 4 months ago

Except these are Point of Sale devices, not sensors.

upvoted 2 times

xRiot007 1 year, 4 months ago

It's irrelevant, you have a device that sends *messages* containing valuable business information. The service bus implementation is even used in Microsoft's own documentation as an example for POS data gathering implementation

upvoted 2 times

Nabeelcp Highly Voted 3 years, 2 months ago

Event Grid and and Event Hubs are basically for event based communication . here the scenario is more suited to Message based communication . so the answer is correct

upvoted 15 times

ferut 2 years, 5 months ago

The question doesn't mention 'message', it's just plain 'data'. Events can also be data.

upvoted 3 times

dddddd111 Most Recent 1 week, 6 days ago

Selected Answer: A

I think the answer is Yes.

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.
Event Hub would be a solution here NOT Event Grid.

upvoted 3 times

meoukg 1 year, 7 months ago

Got it on 03/2022, chose B. No

upvoted 1 times

Netspud 1 year, 8 months ago

Selected Answer: B

Answer is no.
upvoted 1 times

AratiIra 1 year, 10 months ago
What is the right ans ???????????
upvoted 3 times

chingdm 1 year, 8 months ago
I would go for Service Bus because it supports transactional data. POS usually connects to inventory at real time, so make sense for a transaction.
upvoted 1 times

ning 2 years, 2 months ago
No, for any subs, you can have max 100 event grid topics, but there are 2000 stores and 5 POS for each store, for each topic, you can have max 500 subs. The only possible way, is to multiple those, to generate 50,000 subs for 100 topics. But, this would be a horrible situation
upvoted 2 times

Knightie 1 year, 2 months ago
This is the reason for b, else Event Grid makes sense too. It works with a pub/sub topic and event handler based operation. It can work in the scenario, but too little topics for 2000stores with 5 devices...100 topics is not efficent at all.
upvoted 1 times

Ram0202 2 years, 2 months ago
max size allowed in service bus is 1 mb ,ques is for 2 mb so go for event hub
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted
upvoted 1 times

altafpatel1984 1 year, 11 months ago
max size allowed for event hub is 1 MB
upvoted 1 times

renzoku 1 year, 3 months ago
2MB for day, each message could be a lot less
upvoted 1 times

kondapaturi 2 years, 4 months ago
Answer – No, The Azure Event Grid service is used to receiving events and would not fit the purpose of the requirement.
upvoted 1 times

Hamg 2 years, 5 months ago
I think many of clarionprogrammer's answers can be very confusing
upvoted 1 times

eobo 2 years, 6 months ago
Correct Answer is : B

Reference: https://azure.microsoft.com/en-us/blog/events-data-points-and-messages-choosing-the-right-azure-messaging-service-for-your-data/
A sale recorded in a point-of-sale solution is both a financial record and an inventory tracking record, and not a mere event. It's recorded in a ledger, which will eventually be merged into a centralized accounting system, often via several integration bridges, and the information must not be lost on the way. The sales information, possibly expressed as separate messages to keep track of the stock levels at the point of sale, and across the sales region, may be used to initiate automated resupply orders with order status flowing back to the point of sale. Service Bus queues or topic subscriptions are ideal for this use-case
upvoted 8 times

paru123456789 2 years, 8 months ago
Answer: No
upvoted 3 times

agcertif 2 years, 9 months ago
An Event is a lightweight notification that indicates thet something happened.
A Message contains the data itself (not just a reference to that data)
=> Correct answer
upvoted 2 times

Cornholioz 2 years, 11 months ago
Just curious: what if IoT Hub or Event Hub was provided as a possible option.
Agree that the scenario is messages based and not event based.
But wouldn't the use case still hold good to have maybe the IoT Hub be the right pick too?
upvoted 1 times

Elecktrus 3 years, 1 month ago
Answer is correct, because EventGrid only support 25 filters as maximum. You can't filter by deviceid, they are a lot more of 25 device
upvoted 4 times

sebainones 3 years ago
As the question mentions :"Device data must be correlated based on a device identifier":
Service Bus offers "Correlation Filters" capabillity.

On the other hand, with Event Grid you could use Advanced filtering you have this limitation:
"Evnt-Grid: 5 advanced filters and 25 filter values across all the filters per event grid subscription" . So, I guess it has to be Service Bus and thus NO, Event Grid won't solve this.
reference: https://docs.microsoft.com/en-us/azure/event-grid/event-filtering

upvoted 4 times

anandk203 3 years, 2 months ago

Comparison of services
COMPARISON OF SERVICES
Service Purpose Type When to use
Event Grid Reactive programming Event distribution (discrete) React to status changes
Event Hubs Big data pipeline Event streaming (series) Telemetry and distributed data streaming
Service Bus High-value enterprise messaging Message Order processing and financial transactions

upvoted 6 times

Question #3 *Topic 6*

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:

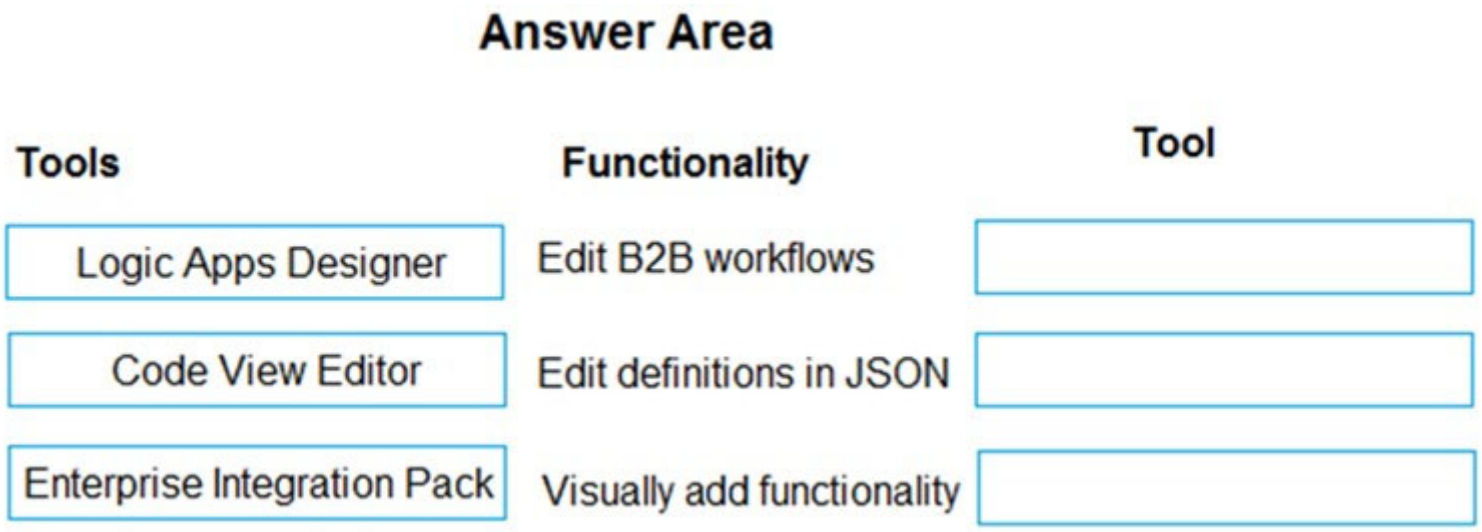

Correct Answer:

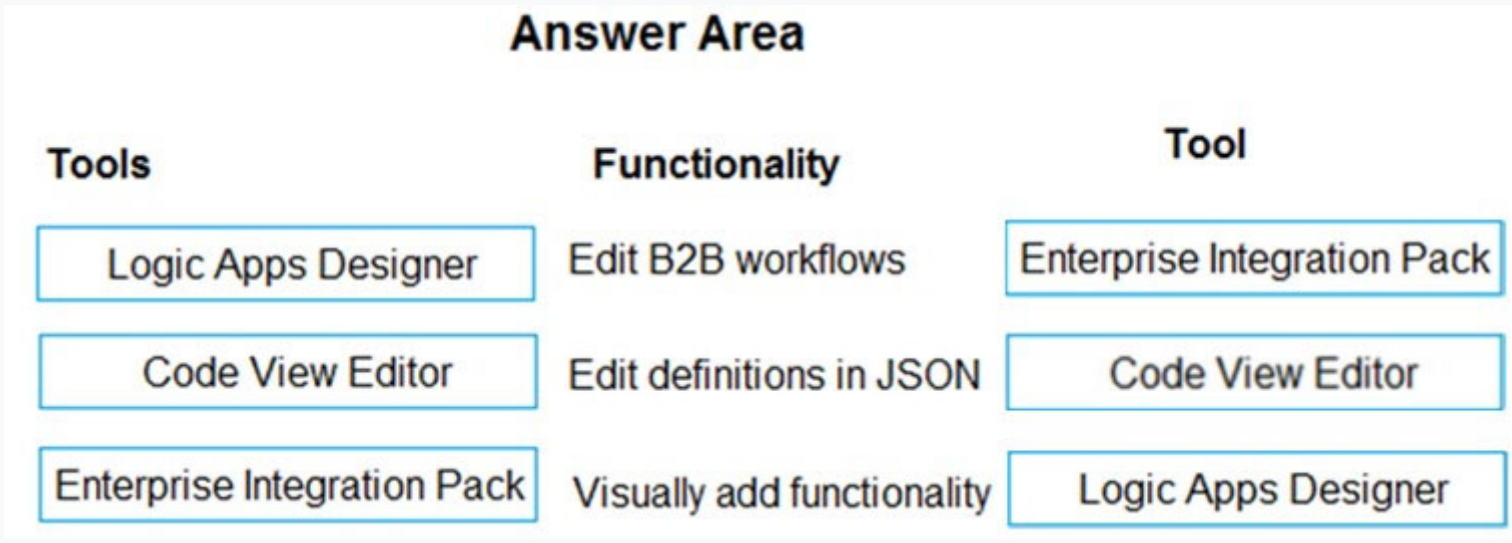

Box 1: Enterprise Integration Pack

For business-to-business (B2B) solutions and seamless communication between organizations, you can build automated scalable enterprise integration workflows by using the Enterprise Integration Pack (EIP) with Azure Logic Apps.

Box 2: Code View Editor -

Edit JSON - Azure portal -

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

Box 3: Logic Apps Designer -

Reference:

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

kimalto452 Highly Voted 2 years, 5 months ago

start from march 2021 no logic apps in the exam!

upvoted 41 times

Marusyk Highly Voted 2 years, 7 months ago

The answer is correct

upvoted 38 times

olowoyinka Most Recent 1 year, 10 months ago

Correct answer

upvoted 4 times

JSrikanth 2 years, 1 month ago

That's correct Answer

upvoted 3 times

Question #4 *Topic 6*

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?

A.

```
az servicebus queue create
--resource-group fridge-rg
--namespace-name fridge-ns
--name fridge-q
```

B.

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

C.

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

D.

```
connectionString-$)az serviceBus namespace authorization-rule keys list
--resource-group fridge-rg
--fridge-ns fridge-ns
--query primaryConnectionString -output tsv)
```

Correct Answer: *A*

A service bus instance has already been created (Step 2 below). Next is step 3, Create a Service Bus queue.

Note:

Steps:

Step 1: # Create a resource group

resourceGroupName="myResourceGroup"

az group create --name $resourceGroupName --location eastus

Step 2: # Create a Service Bus messaging namespace with a unique name namespaceName=myNameSpace$RANDOM az servicebus namespace create --resource-group $resourceGroupName --name $namespaceName --location eastus

Step 3: # Create a Service Bus queue

az servicebus queue create --resource-group $resourceGroupName --namespace-name $namespaceName --name BasicQueue

Step 4: # Get the connection string for the namespace

connectionString=$(az servicebus namespace authorization-rule keys list --resource-group $resourceGroupName --namespace-name $namespaceName --name

RootManageSharedAccessKey --query primaryConnectionString --output tsv)

Reference:

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

SnakePlissken Highly Voted 2 years, 5 months ago

I think the Service Bus has already been created and now they ask you to complete the configuration. The next step is creating the queue.
In fact, all the steps are shown:
B. Create group.
C. Create Service Bus.
A. Create Queue. <-- Correct answer.
D. Get connectionstring.

upvoted 60 times

SnakePlissken 2 years, 5 months ago

Got this question at the exam and scored 100% on Azure Storage, so I'm sure this is correct.

upvoted 20 times

SnakePlissken 2 years, 5 months ago

I'm really sorry, Service Bus falls in another section where I scored 85%, so I can't be sure if this is the correct answer.

upvoted 15 times

xRiot007 1 year, 4 months ago

Your steps are correct.
You need a resource group before creating any resource.
The service bus that you create in Step 2 will be a resource in that group.
Inside an Azure Service Bus you can create queues and topics
Inside the same Azure Service Bus, you will have the primary and secondary connection strings. You can get the connection before or after you created your queue.
In code, you will connect to the service bus and then specify the name of the topic or queue to the service client.

upvoted 2 times

RavindraDevkhile Highly Voted 2 years, 7 months ago

It Ask for create an Azure Service Bus instance by providing not the Service bus Queue
So correct Answes should be C

upvoted 22 times

damianadalid 5 months, 1 week ago

In Azure Portal explation to set up Azure Service Bus, it forces you to create a namespace as you start creating the instance: https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-portal So I would say the namespace is already there and therefore the answer is A (create the queue).

upvoted 1 times

Bear_Polar 10 months, 4 weeks ago

As C location part is incorrect (fridge-loc is not a variable neither a valid location). So, A must be the ans (assume that the queue has been created)

upvoted 2 times

Michael2023 Most Recent 9 months ago

Got this on my exam 30.1.2023 went with Highly Voted answer score 876

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

A or C? looking at question C looks correct but here A is highly voted

upvoted 1 times

student85 1 year ago

I think A is correct https://learn.microsoft.com/en-us/cli/azure/servicebus/queue?view=azure-cli-latest

upvoted 1 times

elequiel 1 year ago

Got it in exam 20/10/2022

upvoted 3 times

rcamara32 1 year, 9 months ago

C is the correct answer

upvoted 5 times

tiggertag2000 1 year, 10 months ago

C is the correct answer
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-tutorial-topics-subscriptions-cli

upvoted 2 times

tiggertag2000 1 year, 10 months ago

I take that back, the service bus is already created, so A is the right answer

upvoted 3 times

hinoue 1 year, 12 months ago

question says You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location

pricing tier(sku) can be set at the creation of namespace, not queue

So, I think answer is C.
https://docs.microsoft.com/en-us/cli/azure/servicebus/namespace?view=azure-cli-latest#az_servicebus_namespace_create
https://docs.microsoft.com/en-us/cli/azure/servicebus/queue?view=azure-cli-latest#az_servicebus_queue_create

upvoted 4 times

noro5 1 year, 9 months ago

It says also 'You need to complete the configuration.' so I assume the Service Bus instance is in place

upvoted 3 times

Franz22 2 years ago

Answer A is correct, because if we look at the documentation of Microsoft where are listed all the steps to create a Service Bus:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-portal
We can see that when we create a Service Bus, we MUST specify a namespace so the ServiceBus instance is the namespace itself and the question says "You create a service bus instance... You need to complete the configuration" --> So, in order to complete the configuration, we have to create a Queue.

upvoted 8 times

RajMasilamani 2 years, 1 month ago

The answer should be A.
connection string and the queue name. You use them to send and receive messages.
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-cli

upvoted 2 times

ning 2 years, 2 months ago

No idea what this question is asking ...

upvoted 10 times

warchoon 7 months, 4 weeks ago

It's asking which command is correct. The only one has no errors and provides requirements

upvoted 1 times

mc0re 2 years, 2 months ago

According to Microsoft, the namespace must be created before the queue. So the correct answer must be C.
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-cli

upvoted 2 times

Mr_X 2 years, 3 months ago

Commands B&C have set wrong location this will throw an error.
D does not make sense to me at all. Thats nonsense.

upvoted 6 times

Helpnosense 2 years, 4 months ago

--location fridge-loc is definitely wrong for az servicebus namespace create command. Whoever thinks C is right, just try the command to see if CLI will succeed or not.
A is correct.

upvoted 7 times

Laaptu 2 years, 4 months ago

Option A is correct, az servicebus queue create --resource-group myresourcegroup --namespace-name mynamespace --name myqueue
https://docs.microsoft.com/en-us/cli/azure/servicebus/queue?view=azure-cli-latest#az_servicebus_queue_create

upvoted 2 times

manojchavan 2 years, 4 months ago

Question is poorly worded. I think what is asked here is that service bus instance is already created and now you need to complete the configuration to start using the bus. In this case, you will need to create Queue and hence correct answer is A.

upvoted 3 times

Question #5 *Topic 6*

HOTSPOT -

You are developing an application that uses Azure Storage Queues.

You have the following code:

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse
(CloudConfigurationManager.GetSetting("StorageConnectionString"));
CloudQueueClient queueClient = storageAccount.CreateCloudQueueClient();

CloudQueue queue = queueClient.GetQueueReference("appqueue");
await queue.CreateIfNotExistsAsync();

CloudQueueMessage peekedMessage = await queue.PeekMessageAsync();
if (peekedMessage != null)
{
    Console.WriteLine("The peeked message is: {0}", peekedMessage.AsString);
}
CloudQueueMessage message = await queue.GetMessageAsync() ;
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statement	Yes	No
The code configures the lock duration for the queue.	○	○
The last message read remains in the queue after the code runs.	○	○
The storage queue remains in the storage account after the code runs.	○	○

Correct Answer:

Answer Area

Statement	Yes	No
The code configures the lock duration for the queue.	○	○
The last message read remains in the queue after the code runs.	○	○
The storage queue remains in the storage account after the code runs.	○	○

Box 1: No -

The QueueDescription.LockDuration property gets or sets the duration of a peek lock; that is, the amount of time that the message is locked for other receivers.

The maximum value for LockDuration is 5 minutes; the default value is 1 minute.

Box 2: Yes -

You can peek at the message in the front of a queue without removing it from the queue by calling the PeekMessage method.

Box 3: Yes -

Reference:

https://docs.microsoft.com/en-us/azure/storage/queues/storage-dotnet-how-to-use-queues https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicebus.messaging.queuedescription.lockduration

rajwit Highly Voted 2 years, 9 months ago

Given answer is correct

upvoted 43 times

igorole Highly Voted 2 years, 11 months ago

GetMessageAsync:
Gets a message from the queue using the default request options. This operation marks the retrieved message as invisible in the queue for the default visibility timeout period.

Only marks the message is invisible but does not delete.

upvoted 31 times

muggins 1 year, 4 months ago

Thank you. Provided answer's reasoning for 'Yes' is incorrect.

upvoted 3 times

TaoLu Most Recent 9 months ago

when you use the GetMessageAsync method to retrieve a message from an Azure Storage queue, the message is removed from the queue. If you want to retrieve a message without removing it from the queue, you can use the PeekMessageAsync method instead.

upvoted 4 times

warchoon 7 months, 4 weeks ago

This is not a simple Queue. This is another class that does not delete the message when it is read.

upvoted 2 times

OPT_001122 11 months, 1 week ago

Given answer is correct

upvoted 1 times

xRiot007 1 year, 4 months ago

No - we are not configuring anything here
Yes - a message that is retrieved from a queue and not marked as completely processed will return to the same queue after a given amount of time, no matter if you use Peek or Get.
Yes - the queue will stay there until we decide to delete it ourselves.

upvoted 3 times

Miroshi 1 year, 9 months ago

2nd answer is No

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.queue.cloudqueue.getmessageasync?view=azure-dotnet-legacy

You Take the msg after you peeked.

upvoted 2 times

ReniRechner 1 year, 8 months ago

peek: get the message, don't lock, don't delete
get: get the message, lock the message (make it invisible for some seconds)
delete: delete the message

intention is: if get would also delete the message from the queue, an error in the function would render the message unhandled. Thus first get, process than delete.
"at least handled once"

upvoted 9 times

Netspud 1 year, 8 months ago

Your link appears to imply that the message is not deleted unless additional params are included. Which would suggest the answer is as quoted, yes. (Not no as you suggest)

upvoted 3 times

Yazhu 1 year, 9 months ago

I dont understand here.. its a repetitive of 9th question.. its a copy paste question
why diff answers? can someone clarify me?

upvoted 1 times

Yazhu 1 year, 9 months ago

commented in wrong window..sorry

upvoted 1 times

nonoss 2 years ago

Box 3 : Correct answer
https://docs.microsoft.com/en-us/azure/storage/queues/storage-dotnet-how-to-use-queues?tabs=dotnet#delete-a-queue

upvoted 1 times

paru123456789 2 years, 8 months ago

Answer:
No
Yes
Yes

upvoted 12 times

bhushan_786 2 years, 9 months ago

Can someone confirm if the given answers are correct or not??

upvoted 2 times

diligent176 2 years, 9 months ago

They are correct, N, Y, Y.
The code does not configure lock duration.
The message will remain because GetMessageAsync does not remove it.
The queue will also remain after execution.

upvoted 14 times

igorole 2 years, 11 months ago

Previous comment can be ignored, sorry, the right interface is here:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.queue.cloudqueue.getmessageasync?view=azure-dotnet-legacy

upvoted 2 times

igorole 2 years, 11 months ago

LastLine, GetMessageAsync() is not even part of the interface, this code just won't compile.
https://docs.microsoft.com/en-us/dotnet/api/azure.storage.queues.queueclient?view=azure-dotnet

upvoted 1 times

Kobee 2 years, 10 months ago

You're wrong. It's CloudQueue, not QueueClient

upvoted 4 times

gematsaljoa 2 years, 11 months ago

1. X
2. O
3. O

upvoted 5 times

JVTM 2 years, 11 months ago

2) explanation comments a peekMessage (which leave message in queue). But the last command is getMessageAsync() which reads the message out of the queue. So, it is not in queue anymore.

upvoted 3 times

thomas204 2 years, 11 months ago

According to me you need to call DeleteMessageAsync() to remove the message from the queue.

upvoted 21 times

ashuyop 2 years, 11 months ago

thats right
https://docs.microsoft.com/en-us/azure/storage/queues/storage-dotnet-how-to-use-queues?tabs=dotnet#de-queue-the-next-message

upvoted 8 times

nonoss 2 years ago

Dequeue a message from a queue in two steps. When you call ReceiveMessages, you get the next message in a queue. A message returned from ReceiveMessages becomes invisible to any other code reading messages from this queue. By default, this message stays invisible for 30 seconds. To finish removing the message from the queue, you must also call DeleteMessage. This two-step process of removing a message assures that if your code fails to process a message due to hardware or software failure, another instance of your code can get the same message and try again. Your code calls DeleteMessage right after the message has been processed.

upvoted 1 times

Magnetor 2 years, 11 months ago

GetMessageAsync() does not qet out the message from the queue

upvoted 31 times

Question #6 *Topic 6*

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?

A.

```
az group create
  --name fridge-rg
  --location fridge-loc
```

B.

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

C.

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

D.

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

Correct Answer: *C*

A service bus instance has already been created (Step 2 below). Next is step 3, Create a Service Bus queue.

Note:

Steps:

Step 1: # Create a resource group

resourceGroupName="myResourceGroup"

az group create --name $resourceGroupName --location eastus

Step 2: # Create a Service Bus messaging namespace with a unique name namespaceName=myNameSpace$RANDOM az servicebus namespace create --resource-group $resourceGroupName --name $namespaceName --location eastus

Step 3: # Create a Service Bus queue

az servicebus queue create --resource-group $resourceGroupName --namespace-name $namespaceName --name BasicQueue

Step 4: # Get the connection string for the namespace

connectionString=$(az servicebus namespace authorization-rule keys list --resource-group $resourceGroupName --namespace-name $namespaceName --name

RootManageSharedAccessKey --query primaryConnectionString --output tsv)

Reference:

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

manojchavan Highly Voted 2 years, 4 months ago

Question is poorly worded. I think what is asked here is that service bus instance is already created and now you need to complete the configuration to start using the bus. In this case, you will need to create Queue and hence correct answer is C

upvoted 30 times

Franz22 Highly Voted 2 years ago

This question is equal to the question #9, the ONLY difference here is that we have to pick the Powershell answer, because it's the only that has the creation of the service bus queue. Go to the question #9 to see my answer why we should create the queue.

upvoted 13 times

AbidooKing Most Recent 3 months, 2 weeks ago

Got this 2023July13. Went with given answer. score 917

upvoted 3 times

elequiel 1 year ago

Got it in exam 20/10/2022

upvoted 3 times

Eltooth 1 year, 4 months ago

C is correct answer.

upvoted 6 times

[Removed] 1 year, 6 months ago

Agree with the accepted answer but given question 9, this seems less like it is testing your knowledge and more an attempt to trick you.

upvoted 4 times

SivajiTheBoss 1 year, 7 months ago

This is the command to create Service bus namespace.
New-AzServiceBusNamespace -ResourceGroupName ContosoRG -Name ContosoSBusNS -Location eastus

Option C is wrong New-AzureRM....

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

To create ServiceBusQueue this is the command in Powershell
New-AzServiceBusQueue -ResourceGroupName ContosoRG -NamespaceName ContosoSBusNS -Name ContosoOrdersQueue

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer C: As needs to creata instance of Queue

upvoted 1 times

edengoforit 1 year, 9 months ago

The description is almost the same as the question #9. Hence, after creating a namespace, we should create a Service Bus Queue thus provided answer is correct. Answer should be C

upvoted 3 times

Richard2021 2 years, 1 month ago

the correction answer is B.

upvoted 2 times

Illumielle 2 years, 2 months ago

This is a repeat of question 9. But it seems to have different answers.

upvoted 2 times

ReniRechner 1 year, 8 months ago

this is because there is more than 1 way to do it...

upvoted 2 times

ning 2 years, 2 months ago

No idea what question is this ...

Steps should be 1. create a resource group 2. create a service bus namespace 3. create a service bus queue 4. get connection string for the queue

upvoted 2 times

ning 2 years, 2 months ago

There is no such thing called service bus instance ...

Only Service Bus Namespace, Service Bus Queue Or Service Bus Topic

upvoted 1 times

ReniRechner 1 year, 8 months ago

namespace = instance

upvoted 1 times

warchoon 7 months, 4 weeks ago

Everything created is an instance. It's a trick

upvoted 1 times

mc0re 2 years, 2 months ago

A namespace must be created before creating a queue. So why not B or D?
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-cli

upvoted 1 times

asdasdasg2 1 year, 10 months ago

The question states that you have already created the service bus, which is what "creating a namespace" does. So therefore the next step is to create the queue

upvoted 2 times

if54uran 2 years, 4 months ago

This question does not seem to make much sense

upvoted 3 times

jokergester 2 years, 7 months ago

There is a requirement to provide an instance and yet the correct answer does not have the prior creation of the namespace. "C" would be ideal if the namespace is created prior to the script.

upvoted 2 times

Zidimirite 2 years, 6 months ago

"You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location."
They seem to ask you to do this, but in this case I think they mean that you've done that and next "You need to complete the configuration.", which is adding a queue to the just created namespace. So C is right.
It has to be right, A makes no sense and B and D both do the same thing.

upvoted 14 times

rustycables 2 years, 2 months ago

I don't know but guessed D. D seems more like a CLI command, and it also has location - which is specified in the question. C does not have location and does not look like a typical CLI command that begins AZ.

upvoted 2 times

satyadharma 2 years ago

It is a powershell command which is mentioned in question as well. C is correct answer

upvoted 2 times

Question #7 *Topic 6*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.
You have the following requirements:
- Queue size must not grow larger than 80 gigabytes (GB).
- Use first-in-first-out (FIFO) ordering of messages.
- Minimize Azure costs.

You need to implement the messaging solution.
Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application. Create an Azure Function App that uses an Azure Storage
Queue trigger.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Create an Azure Function App that uses an Azure Service Bus Queue trigger.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function

Community vote distribution

B (100%)

Laaptu Highly Voted 2 years, 4 months ago

Answer is B (NO),
Becz FIFO is supported by Service bus queue and Service bus queue should be use for data less than 80GB
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted

upvoted 23 times

xRiot007 1 year, 4 months ago

The most important thing is not capacity, but the FIFO guarantee that is provided only by the service bus and not the service queue. The link you provided is specifying that as well.

upvoted 3 times

ay_m 2 months, 1 week ago

great link, thanks a lot!

upvoted 1 times

elequiel Highly Voted 1 year ago

Got it in exam 20/10/2022

upvoted 5 times

adilkhan Most Recent 8 months ago

The solution uses Azure Storage Queue, which provides FIFO ordering of messages, and can handle a large volume of messages. By using the .Net API to add a message to the Azure Storage Queue, the mobile application can send messages to the queue reliably.

The Azure Function App with a Storage Queue trigger can process the messages in the queue as soon as they are added, ensuring that messages are processed in the order in which they are received. This helps meet the FIFO requirement.

Additionally, by using Azure Storage Queue, the solution can ensure that the queue size does not grow larger than 80 GB, as specified in the requirement. This can help minimize Azure costs as well, as storage usage is a key factor in determining Azure costs.

Overall, the proposed solution meets all the requirements and can provide a reliable, scalable, and cost-effective messaging solution for the given scenario. YES ChatGPT

upvoted 2 times

paunski7 6 months, 1 week ago

ChatGPT answered NO - 3.5V paid

upvoted 2 times

adilkhan 8 months ago

ChatGPT is wrong.B is correct as Azure Service Bus Queue can ensure that the queue size does not grow larger than 80 GB
upvoted 3 times

Bear_Polar 11 months, 1 week ago

Selected Answer: B

Correct
upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: B

B is correct answer.
upvoted 2 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.
upvoted 3 times

heisenberg33 1 year, 9 months ago

Selected Answer: B

Answer is No. Ref: https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted
upvoted 1 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with NO
upvoted 4 times

kondapaturi 2 years, 4 months ago

Answer is No
upvoted 4 times

idrisfl 2 years, 7 months ago

Shouldn't this be Yes? I'm not sure where "Service Bus" came into the picture in their answer.
upvoted 3 times

idrisfl 2 years, 7 months ago

I take back my previous comment. 'No' is the correct answer. It says 'Queue size should not grow larger than 80 GB'
upvoted 9 times

Didi123 2 years, 7 months ago

agreed, also FiFo is only guaranteed in Service Bus Queues and not in Azure Storage Queues
upvoted 11 times

Franz22 2 years ago

I think that here, the size is not the thing that should let you think "oh, it's not the right answer", because you can use storage queue even with less than 80gb of data. The answer "No" should be given because the storage queue doesn't support FIFO.
upvoted 8 times

warchoon 7 months, 4 weeks ago

And it has no Push-style API.
https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted#:~:text=Push%2Dstyle%20API-,No,-Yes%0A%0AOur%20.NET
upvoted 1 times

Question #8 *Topic 6*

DRAG DROP -

You develop software solutions for a mobile delivery service. You are developing a mobile app that users can use to order from a restaurant in their area. The app uses the following workflow:

1. A driver selects the restaurants for which they will deliver orders.
2. Orders are sent to all available drivers in an area.
3. Only orders for the selected restaurants will appear for the driver.
4. The first driver to accept an order removes it from the list of available orders.

You need to implement an Azure Service Bus solution.

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:

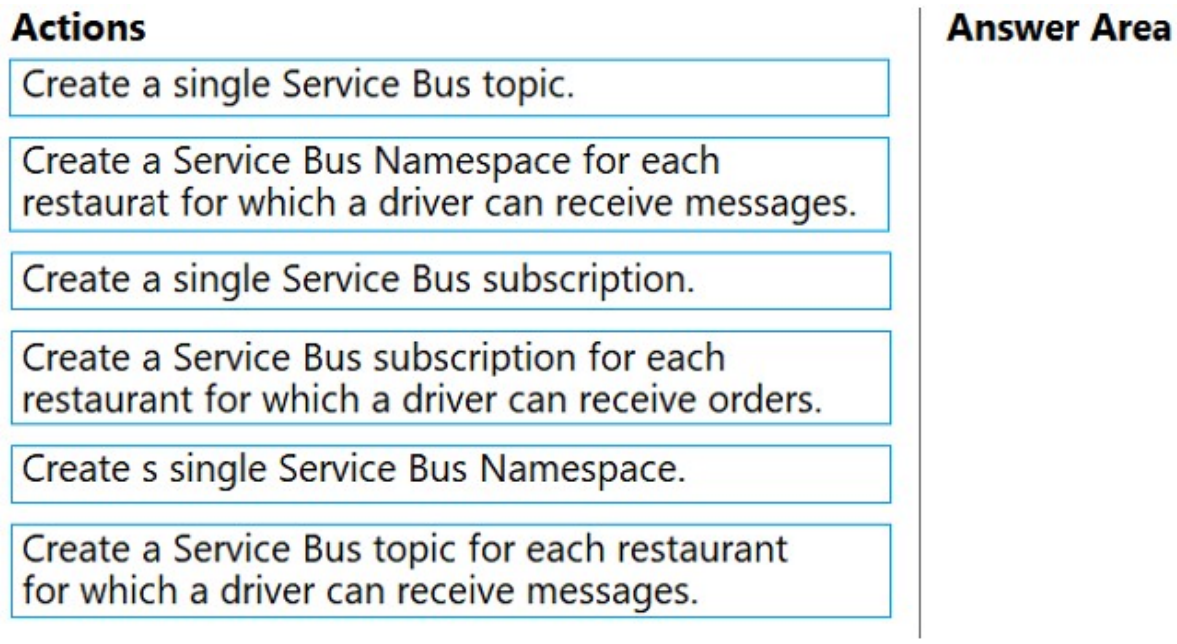

Correct Answer:

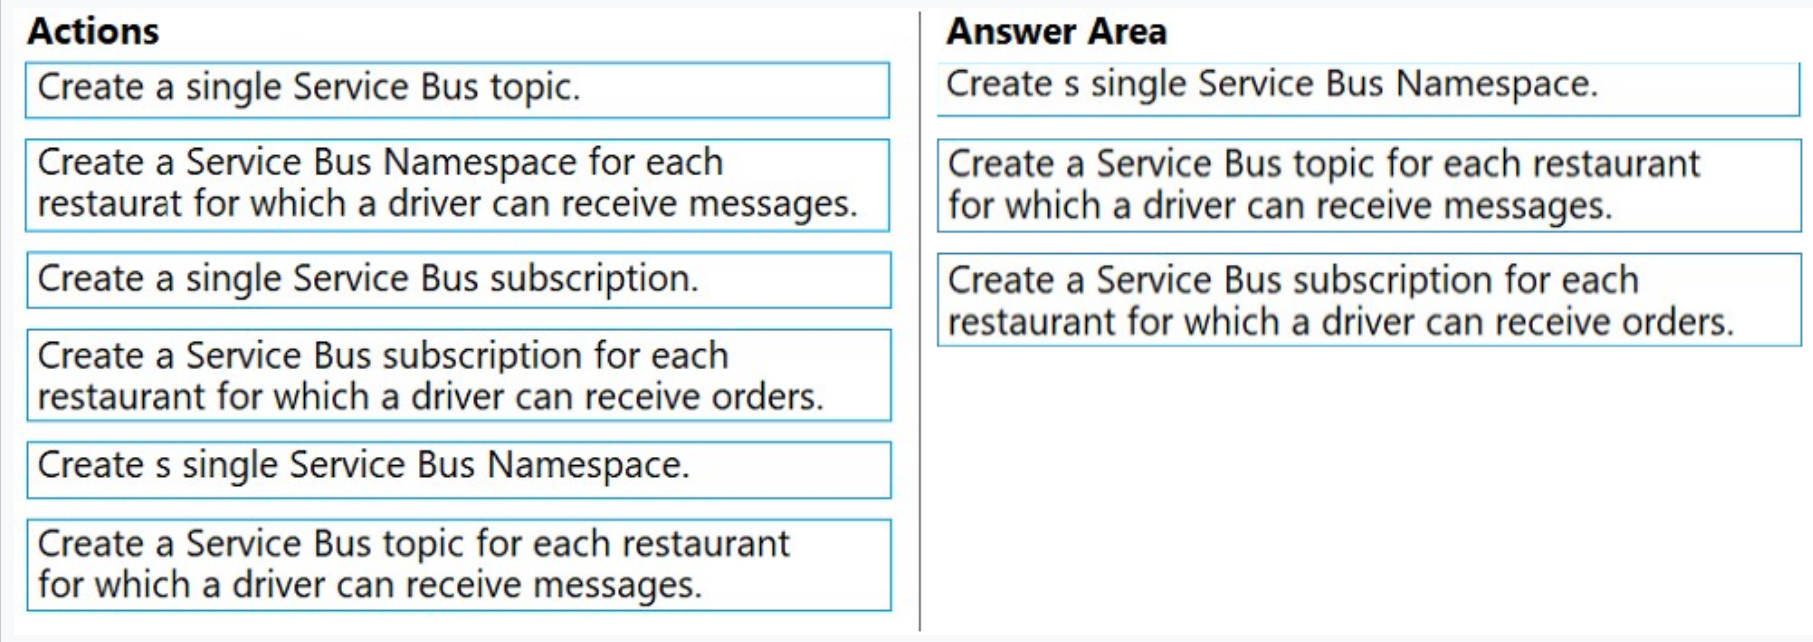

Box 1: Create a single Service Bus Namespace

To begin using Service Bus messaging entities in Azure, you must first create a namespace with a name that is unique across Azure. A namespace provides a scoping container for addressing Service Bus resources within your application.

Box 2: Create a Service Bus Topic for each restaurant for which a driver can receive messages.

Create topics.

Box 3: Create a Service Bus subscription for each restaurant for which a driver can receive orders.

Topics can have multiple, independent subscriptions.

Reference:

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messaging-overview

NH Highly Voted 2 years, 7 months ago

Create a single Service Bus Namespace.
Create a single Service Bus Topic.
Create a Service Bus subscription for each restaurant for which a driver can receive orders.

upvoted 152 times

naivecoder786 8 months ago

You are correct. Azure Service Bus supports filtering messages at the subscription level by using SQL-like syntax called "SQL filters".

So in this case, it is possible to create only one topic and have different subscriptions for each restaurant and filter the messages at the subscription level based on the restaurant name. This would allow drivers to only receive orders from the selected restaurants, without having to create a separate topic for each restaurant.

upvoted 2 times

sashasasha 2 years, 5 months ago

If you create a single Service Bus Topic, two different drivers who subscribed to the same restaurant will each receive a copy of the order that could process independently and it's in conflict with the condition: "4. The first driver to accept an order removes it from the list of available orders."

upvoted 4 times

coffecold 1 year ago

Why ? One order is placed solely with the restaurant topic applicable. No copy. Then
one subscriber locks the message, and then deletes it. Once locked it cannot be accessed by other subscribers.

upvoted 2 times

jay158 2 years, 4 months ago

In fact all the drivers, who have subscribed to restaurant will get notification.
The first driver, who accepts it, will click the order, and it shall be dequeued.

upvoted 11 times

hstml 2 years, 1 month ago

That is completely correct. This question is also in Whizlabs where the explanation is the same as jay158s.

upvoted 2 times

lugospod 1 year, 10 months ago

No it won't get dequeued because this is not a queue.
There is no option for SUBSCRIPTION PER DRIVER, only PER RESTAURANT or a SINGLE subscription. So only SINGLE listener will receive the notification and until s/he accepts it, rejects it or it times out no one else will see the message. If this were a QUEUE than it would be as you described. This question sucks on so many levels - too many open questions arise to be able to give a valid response, plus, the final one, I would rather use queues for this problem then subscriptions...

upvoted 3 times

TaoLu 9 months ago

This is obviously wrong. The examtopic answser is correct.

upvoted 2 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000.

upvoted 4 times

lugospod 1 year, 10 months ago

One thing to note in this scenario if we have 20 drivers... and 19 of them are busy (talking to their spouse, ...) it will take 19 timeouts for that one order to finally get picked up by the 20th driver. So yes, this is less evil then the resto of the options, but still I would rather use queues so that each driver SEES all of the orders... by using subscriptions the driver doesnt have a list..s/he only sees the current message that was delegated by the round robin algorithm.

upvoted 1 times

Opimer Highly Voted 2 years, 7 months ago

Correct. If you create only one topic there won't be any filtering of messages depending on the restaurant. So every driver will have access to all the the orders. With a topic by restaurant, message from one restaurant will go to only one topic, from which the driver can choose to subscribe.

upvoted 16 times

coffecold 1 year ago

I agree. A driver (subscriber) can subscribe to a topic once he is near the restaurant and un-subscribe if he is at a distance. I don't get the high vote of NH.

upvoted 2 times

Phenr 2 years, 7 months ago

I don't think that's correct. You can filter the messages on subscription, so the subscription can get a message from a specific restaurant even with just one topic created.

upvoted 8 times

rdemontis 2 years, 7 months ago

Exactly and creating one topic for each restaurant doesn't cover two of the four requirements:
- "Orders are sent to all available drivers in an area"
- "The first driver to accept an order removes it from the list of available orders"

This means that one restaurant can be server by more drivers. So how can you think to manage order acceptance from one driver if you duplicate the same in more topics? It's a big deal that cause a lot of overwork not necessary.

It's all simpler by managing all orders with a topic and using the filters for each subscription to allow drivers to view only the orders that compete with them. I also inform you that on a paid test platform such as https://www.testpreptraining.com/ there is the same question and the correct answer requires only one topic.

upvoted 7 times

kwaazaar 2 years, 6 months ago

so how would the driver remove the order if he has his own subscription? The order would still be visible in other subscriptions.
upvoted 1 times

anvimi 2 years, 6 months ago
multiple drivers can share a single restaurant subscription thus only one driver will handle an order
upvoted 3 times

mandynotmandy 1 year, 7 months ago
exactly, you cant have each driver on their own sub because that duplicates the orders into each sub queues, meaning one guy accepts the order, another guy still sees it in their own sub queue. the subs has to be on restaurant level, and you can do that by filters on the sub
upvoted 1 times

jaf19f Most Recent 2 months, 2 weeks ago
I got this question (12-Aug-2023) and I chose the given answers - 932 passed
upvoted 3 times

juanckar 3 months, 3 weeks ago
This was on the exam (July 2023). Went with single/single/subscription for each restaurant. Scored 917
upvoted 3 times

warchoon 7 months, 4 weeks ago
Correct
1: Namespace is about the solution
2: Topics are about restaurant messages
3: Subscriptions are about driver's orders
upvoted 2 times

apparaog99 10 months ago
Got in exam 12/31
upvoted 2 times

elequiel 1 year ago
Got it in exam 20/10/2022
upvoted 5 times

bipinboops 1 year, 2 months ago
The provided answer is the most correct. There is a 2000 subscription limit for topics, so you would only allow 2000 drivers TOTAL for the entire app if you chose this as a solution. You need one topic per restaurant or the application cannot scale. You can absolutely have 1 topic, that will still work as intended, it's just not the best solution. Source: https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits#service-bus-limits
upvoted 3 times

zhongzi 1 year, 4 months ago
It's the same question in ESI. The given answer is the same as highly voted.
upvoted 1 times

AZ204Cert 1 year, 6 months ago
Got this on 04/05/22
upvoted 1 times

iamstudying 1 year, 7 months ago
BUDDIES, listen listen..

1. Single Service Bus Namespace
2. Topic per restaurant (10k limit of restaurants).

(A single topic would work, but requires subscription filters or another component to decide which drivers receive which restaurant orders. But this isn't specified so imo not an option)

3. Single subscription per topic. Each drivers who wants to deliver for a restaurant will periodically peek the subscription for messages (no lock, message not removed). Once they decide they want to deliver the order, they dequeue the message and away they go, other drivers who attempt to do so will not be able to.

(At first, I also thought multiple subscriptions - but this introduces another complexity of syncing the orders across all the subscriptions if someone accepts the order delivery first... how can we notify other subscriptions? we can't.)

All in all, agree the solution sucks a$$ but this is the best bet.
upvoted 6 times

coffecold 1 year ago
Multi subscriptions (drivers) per topic. How can we notify other subscriptions?> the message is locked.
https://learn.microsoft.com/en-us/rest/api/servicebus/peek-lock-message-non-destructive-read
upvoted 1 times

petitbilly 1 year, 7 months ago
Got it in exam 03/22
upvoted 1 times

leonidn 1 year, 9 months ago

Single topic is less complex and more flexible solution from the client applications which submit orders. A single topic is less complex from the perspective of infrastructure management effort.

Create a single Service Bus Namespace.
Create a single Service Bus Topic.
Create a Service Bus subscription for each restaurant for which a driver can receive orders.

upvoted 3 times

altafpatel1984 1 year, 11 months ago

I guess options are incomplete. There should be option to have subscription for each driver. So in that case topic for each restaurant with subscription for each driver will make sense to have pub-sub architecture.

upvoted 4 times

altafpatel1984 1 year, 11 months ago

I guess options are incomplete. There should be option to have subscription for each driver. So in that case topic for each restaurant with subscription for each driver will make sense to have pub-sub architecture.

upvoted 1 times

altafpatel1984 1 year, 11 months ago

I guess options are incomplete. There should be option to have subscription for each driver. So in that case topic for each restaurant with subscription for each driver will make sense to have pub-sub architecture.

upvoted 1 times

debanjan10 2 years ago

Single Bus Namespace > Single Bus Topic > Service Bus subscription for each restro

upvoted 2 times

Question #9 *Topic 6*

HOTSPOT -

You develop a news and blog content app for Windows devices.

A notification must arrive on a user's device when there is a new article available for them to view.

You need to implement push notifications.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
string notificationHubName = "contoso_hub";
string notificationHubConnection = "connection_string";
[ ▼ ] hub =
```

- NotificationHubClient
- NotificationHubClientSettings
- NotificationHubJob
- NotificationDetails

```
 [ ▼ ] (notificationHubConnection, notificationHubName);
```

- NotificationHubClient
- NotificationHubClientSettings
- NotificationHubJob
- NotificationDetails

```
 [ ▼ ]
```

- GetInstallation
- CreateClientFromConnectionString
- CreateOrUpdateInstallation
- PatchInstallation

```
string windowsToastPayload =
 @"<toast><visual><binding template=""ToastText01""><text id=""1"">" +
 @"New item to view" + @"</text></binding></visual></toast>";
try
{
var result =
  await hub. [ ▼ ] (windowsToastPayload);
  ...
}
catch (System.Exception ex)
{
 ...
}
...
```

- SendWindowsNativeNotificationAsync
- SubmitNotificationHubJobAsync
- ScheduleNotificationAsync
- SendAppleNativeNotificationAsync

Correct Answer:

Answer Area

```
string notificationHubName = "contoso_hub";
string notificationHubConnection = "connection_string";
[ NotificationHubClient ▾ ] hub =
    NotificationHubClient
    NotificationHubClientSettings
    NotificationHubJob
    NotificationDetails
[ NotificationHubClient ▾ ] . [ CreateClientFromConnectionString ▾ ]
    NotificationHubClient
    NotificationHubClientSettings
    NotificationHubJob
    NotificationDetails
    GetInstallation
    CreateClientFromConnectionString
    CreateOrUpdateInstallation
    PatchInstallation
 (notificationHubConnection, notificationHubName);
string windowsToastPayload =
 @"<toast><visual><binding template=""ToastText01""><text id=""1"">" +
 @"New item to view" + @"</text></binding></visual></toast>";
try
{
var result =
  await hub. [ ▾ ] (windowsToastPayload);
    SendWindowsNativeNotificationAsync
    SubmitNotificationHubJobAsync
    ScheduleNotificationAsync
    SendAppleNativeNotificationAsync
  ...
}
catch (System.Exception ex)
{
 ...
}
...
```

Box 1: NotificationHubClient -

Box 2: NotificationHubClient -

Box 3: CreateClientFromConnectionString

// Initialize the Notification Hub

NotificationHubClient hub = NotificationHubClient.CreateClientFromConnectionString(listenConnString, hubName);

Box 4: SendWindowsNativeNotificationAsync

Send the push notification.

var result = await hub.SendWindowsNativeNotificationAsync(windowsToastPayload);

Reference:

https://docs.microsoft.com/en-us/azure/notification-hubs/notification-hubs-push-notification-registration-management

https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/app-service-mobile/app-service-mobile-windows-store-dotnet-get-started-push.md

sumitg Highly Voted 2 years, 7 months ago

Answer is correct.

upvoted 41 times

Tom87 Highly Voted 2 years, 6 months ago

The answer is correct, but the first reference is not very helpful and the second one doesn't even exists. Here are references of the methods used in the answer:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.notificationhubs.notificationhubclient.createclientfromconnectionstring?view=azure-dotnet#Microsoft_Azure_NotificationHubs_NotificationHubClient_CreateClientFromConnectionString_System_String_System_String_
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.notificationhubs.notificationhubclient.sendwindowsnativenotificationasync?view=azure-dotnet#Microsoft_Azure_NotificationHubs_NotificationHubClient_SendWindowsNativeNotificationAsync_System_String_

upvoted 18 times

Sivaramakrishnan Most Recent 1 year, 8 months ago

This the GitHub link
https://github.com/uglide/azure-content/blob/master/articles/app-service-mobile/app-service-mobile-windows-store-dotnet-get-started-push.md

upvoted 4 times

florodevops 1 year, 8 months ago

1- NotificationHubClient 2 - NotificationHubClient 3- CreateClientFromConnectionString 4- SendWindowsNativeNotificationAsync

upvoted 5 times

Question #10 *Topic 6*

You are developing an Azure messaging solution.

You need to ensure that the solution meets the following requirements:

- Provide transactional support.
- Provide duplicate detection.
- Store the messages for an unlimited period of time.

Which two technologies will meet the requirements? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Azure Service Bus Topic

B. Azure Service Bus Queue

C. Azure Storage Queue

D. Azure Event Hub

Correct Answer: *AB*

The Azure Service Bus Queue and Topic has duplicate detection.

Enabling duplicate detection helps keep track of the application-controlled MessageId of all messages sent into a queue or topic during a specified time window.

Incorrect Answers:

C: There is just no mechanism that can query a Storage queue and find out if a message with the same contents is already there or was there before.

D: Azure Event Hub does not have duplicate detection

Reference:

https://docs.microsoft.com/en-us/azure/service-bus-messaging/duplicate-detection

Community vote distribution

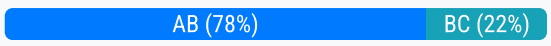

matejka Highly Voted 2 years, 8 months ago

Answer is correct. Queue Storage does not provide transactional support and Event Hub can't be configured to store events for infinite time.

upvoted 38 times

coffecold 1 year ago

Agree, answer is correct, but I think duplicate detection is the key phrase.

upvoted 3 times

mvr 2 years, 8 months ago

Event Hub can be configured to store events, with Event Hub Capture.
Makes the answer still correct though.

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-capture-overview

upvoted 1 times

cbn 2 years, 8 months ago

Event hubs (with capture) are not supposed to retain data forever.
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-faq#what-is-the-maximum-retention-period-for-events

upvoted 3 times

ferut 2 years, 5 months ago

The stress is on 'messaging' which Event Hub is not (it's eventing)

upvoted 6 times

dddddd111 Most Recent 1 week, 6 days ago

Selected Answer: AB

Given answers are correct. ChatGPT apologized for A and D after I asked to justify the answer.

upvoted 1 times

mkahmann 1 month, 3 weeks ago

Selected Answer: AB

Agreed

upvoted 1 times

OPT_001122 11 months, 1 week ago

Selected Answer: AB

A. Azure Service Bus Topic
B. Azure Service Bus Queue

upvoted 1 times

RochaG2 1 year, 1 month ago

Selected Answer: BC

correct

upvoted 2 times

> **RochaG2** 1 year, 1 month ago
>
> Sorry I meant AB!!
>
> upvoted 2 times

[Removed] 1 year, 9 months ago

Does C not become an option because of unlimited storing?
Because < 80gb queue storage can only be used.

upvoted 3 times

> **[Removed]** 1 year, 9 months ago
>
> * > 80gb (more then 80gb)
>
> upvoted 1 times

> **Knightie** 11 months, 3 weeks ago
>
> it's unlimited period of time, not storage size.
>
> upvoted 1 times

Lucario95 1 year, 9 months ago

Selected Answer: AB

Answer is correct

upvoted 4 times

Question #11 *Topic 6*

DRAG DROP -

You develop a gateway solution for a public facing news API.

The news API back end is implemented as a RESTful service and hosted in an Azure App Service instance.

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

Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value 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:

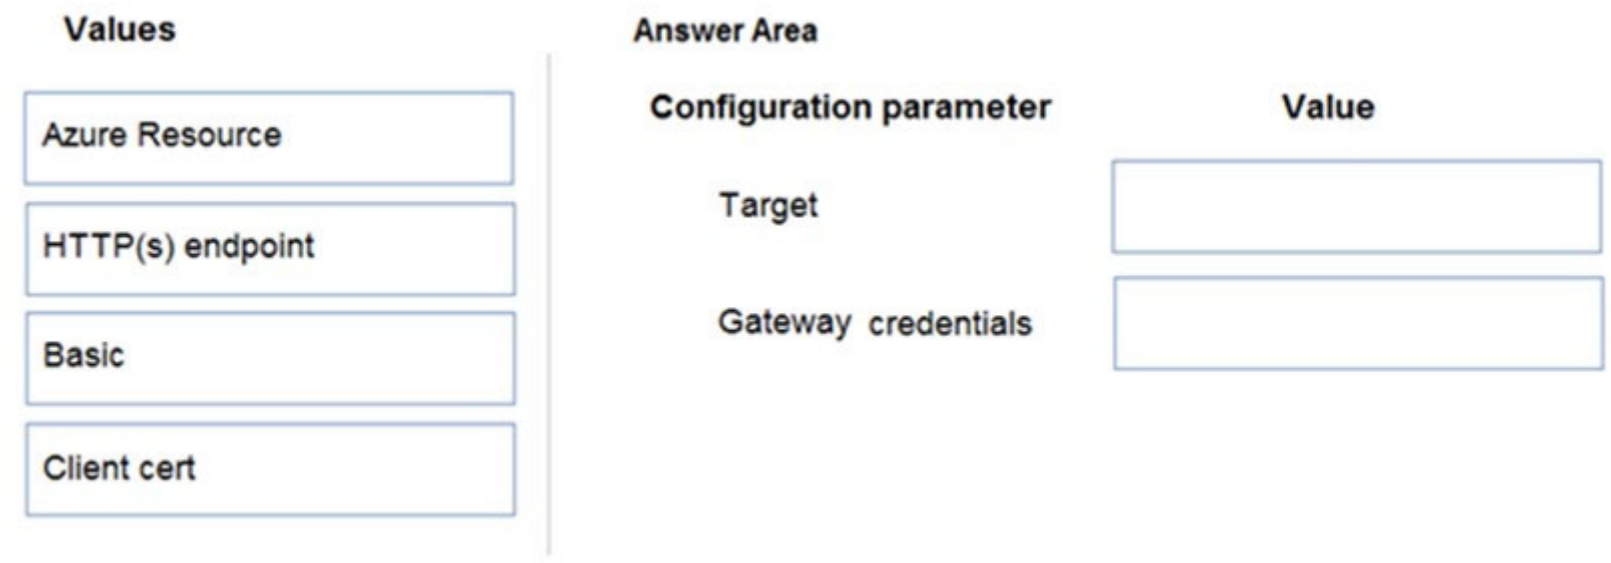

Correct Answer:

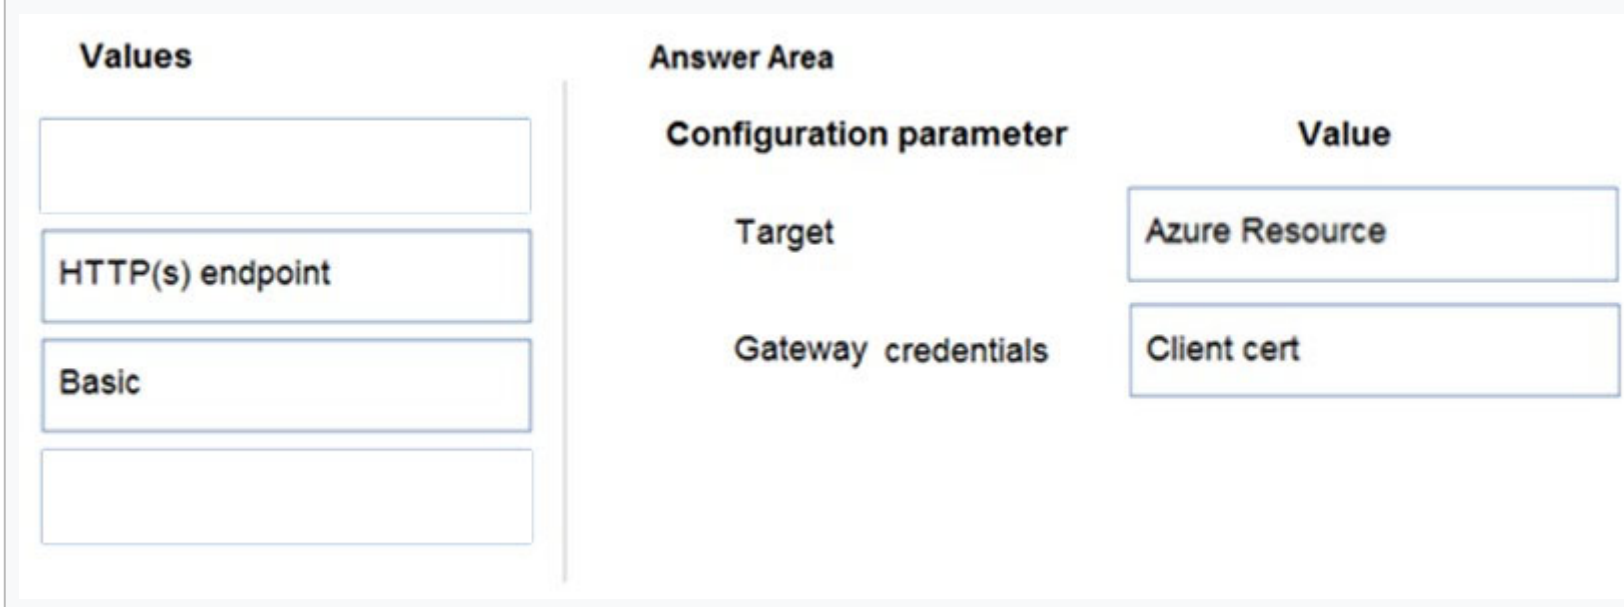

Box 1: Azure Resource -

Box 2: Client cert -

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

Reference:

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

Kuna_Lambo Highly Voted 2 years, 7 months ago

#1 Http ?

upvoted 42 times

tevivi8222 2 years, 7 months ago

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

upvoted 21 times

Molte 1 year, 10 months ago

but why cert and not basic?

upvoted 4 times

adilkhan 9 months ago

configure back-end authentication and to do that use cert

upvoted 2 times

Knightie 1 year, 2 months ago

Public Facing.. so not all incoming request will have their known credentials. So just cert to ensure the caller is the rightful app not the user right.

upvoted 1 times

Laaptu 2 years, 4 months ago

agreed
1. Https(s) endpoint,
2. Client cert

upvoted 15 times

AzureAz204Fan 2 years, 5 months ago

Yes. Target has only two options: 1) Azure Logic App 2) HTTP(s) endpoint. In the question it is given that backend is a REST API which means Target has to be HTTP

upvoted 9 times

ning 2 years, 2 months ago

On the API level there are two targets, Logic App and HTTPS endpoint, on the operation level two targets are Azure Resource and HTTPS endpoint. So sure what level is this question asking, if API level then HTTPS, if operation level that will be azure resource, since it is hosted in Azure.

upvoted 5 times

Gregoryhouse2020 Highly Voted 2 years, 4 months ago

Answer :
1. Https(s) endpoint,
2. Client cert

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

upvoted 31 times

unraval Most Recent 1 month, 1 week ago

Reference for Target has only two options: 1) Azure Logic App 2) HTTP(s) endpoint ??

upvoted 1 times

applepie 3 months ago

got this question today - 7/30/2023

upvoted 2 times

adilkhan 9 months ago

Https and client certificate. 100%
https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication

upvoted 3 times

OPT_001122 11 months, 1 week ago

1. Https(s) endpoint,
2. Client cert

upvoted 5 times

POOOJAAAAAAAAAA 10 months, 1 week ago

am always searching for this opt_001122 person's comment to know the correct answer.. hope he/she's answers are correct.

upvoted 3 times

Ciupaz 4 months, 3 weeks ago

Me too.

upvoted 1 times

gmishra88 1 year ago

Both Azure resource and Http(s) endpoint can be used because it is an azure resource Rest api. Wonder what Microsoft expects.
I'm clueless seeing Target and Gateway. No idea what they are till I saw the portal. Do I have to remember the unfortunate name the Microsoft guys chose for these UI elements?

upvoted 2 times

Pize 1 year, 3 months ago

1.Http
2. Basic or Client cert would fit

upvoted 1 times

Eltooth 1 year, 4 months ago

HTTPs) endpoint
Client cert

upvoted 3 times

Sandeep12093 1 year, 4 months ago

It's definitely since target has only options https and logic app
1. Https(s) endpoint,
2. Client cert

upvoted 3 times

Azprep 1 year, 6 months ago

1.Http End point
2. Client certificate

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

upvoted 1 times

[Removed] 1 year, 8 months ago

1. Https(s) endpoint,
2. Client cert
https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates

upvoted 1 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with http and certificate. but from my perspective basic should also fit.

upvoted 7 times

debanjan10 2 years ago

Target: HTTP endpoints
Gateway creds: Client cert

upvoted 3 times

ning 2 years, 2 months ago

This question is valid? It is so poorly worded. Not sure what is this asking for?
I guess for the first on HTTP(S), for the second one, either basic or certificate is accepted by azure. Nothing is mentioned in the question which one is preferred.

upvoted 7 times

cool_tool 2 years, 2 months ago

1. HTTPS
2. CLIENT CERT
https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication

upvoted 5 times

kondapaturi 2 years, 4 months ago

1. Https(s) endpoint,
2. Client cert

upvoted 5 times

Question #12 *Topic 6*

HOTSPOT -

You are creating an app that uses Event Grid to connect with other services. Your app's event data will be sent to a serverless function that checks compliance.

This function is maintained by your company.

You write a new event subscription at the scope of your resource. The event must be invalidated after a specific period of time.

You need to configure Event Grid.

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

NOTE: Each correct selection is worth one point.

Hot Area:

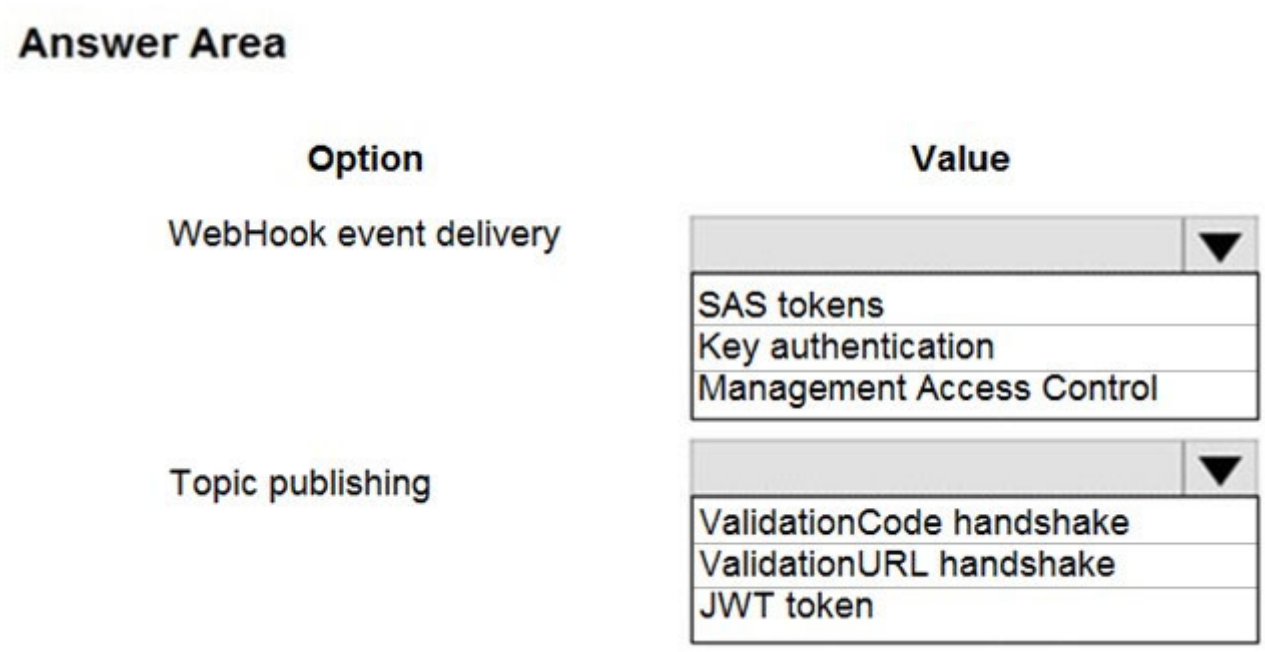

Correct Answer:

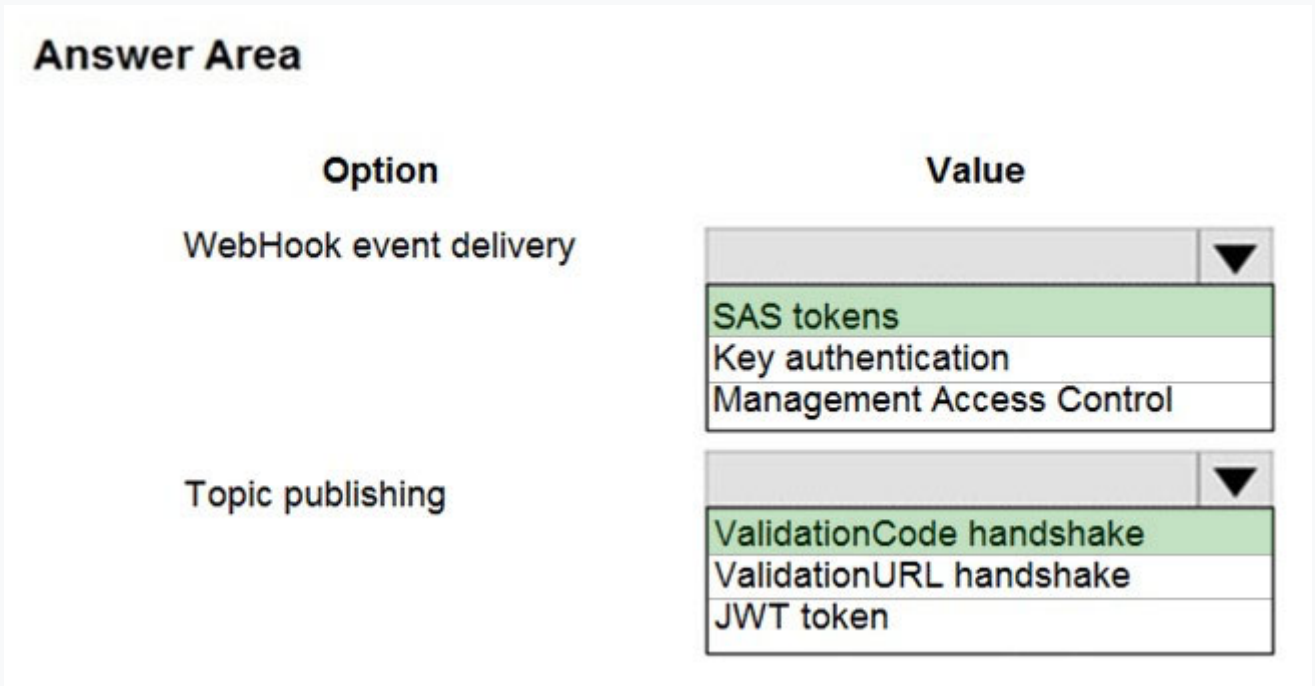

Box 1: SAS tokens -

Custom topics use either Shared Access Signature (SAS) or key authentication. Microsoft recommends SAS, but key authentication provides simple programming, and is compatible with many existing webhook publishers.

In this case we need the expiration time provided by SAS tokens.

Box 2: ValidationCode handshake -

Event Grid supports two ways of validating the subscription: ValidationCode handshake (programmatic) and ValidationURL handshake (manual).

If you control the source code for your endpoint, this method is recommended.

Incorrect Answers:

ValidationURL handshake (manual): In certain cases, you can't access the source code of the endpoint to implement the ValidationCode handshake. For example, if you use a third-party service (like Zapier or IFTTT), you can't programmatically respond with the validation code.

Reference:

https://docs.microsoft.com/en-us/azure/event-grid/security-authentication

jay158 Highly Voted 2 years, 4 months ago

Dropdowns placement in Answer area is incorrect
WebHook Event Delivery ---ValidationCode handsShake
Topic publishing --- SAS Tokens

https://docs.microsoft.com/en-us/azure/event-grid/concepts
https://docs.microsoft.com/en-us/azure/event-grid/webhook-event-delivey
upvoted 46 times

applepie 3 months ago
got this question today, go with this answer - 7/30/2023, score 895/1000
upvoted 5 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000.
upvoted 4 times

Alex_Wackoo 2 years, 4 months ago
Agreed and your last link is not working so here:
https://docs.microsoft.com/en-us/azure/event-grid/webhook-event-delivery
upvoted 7 times

NightshadeRC Most Recent 3 months, 1 week ago
Had this question in today's exam: 2023-07-26.
I think the drop-downs were in the order they are here, not 100% on it though. I went with the recommended answer anyway. Passed 940.
upvoted 4 times

warchoon 7 months, 3 weeks ago
"The event must be invalidated after a specific period of time." - Asynchronous handshake => validationUrl
https://learn.microsoft.com/en-us/azure/event-grid/webhook-event-delivery#:~:text=Asynchronous%20handshake
upvoted 2 times

18Marks 10 months, 3 weeks ago
Got this in exam. Use given answer
upvoted 3 times

jvyas 2 years ago
Below link for Authenticating clients that publish events to Azure Event Grid topics using access key or Shared Access Signature (SAS) token.
https://docs.microsoft.com/en-us/azure/event-grid/authenticate-with-access-keys-shared-access-signatures
upvoted 2 times

warchoon 7 months, 3 weeks ago
So what's the answer? Keys or tokens?
Keys are from SAS too.
SAS token can be expired after "specific period of time".
upvoted 1 times

Simbetie 2 years, 2 months ago
Topic publishing to be matched with the 1st dropdown....the answer must be Management Access Control. Azure Event Grid allows you to control the level of access given to different users to do various management operations such as list event subscriptions, create new ones, and generate keys. Event Grid uses
Azure's role-based access control (RBAC).
upvoted 2 times

ning 2 years, 2 months ago
This question is worded poorly, the handshake only happens when subscription is setup ...
If you are talking about how subscription should be authenticated when receiving messages, then webhook should use jwt
upvoted 1 times

ning 2 years, 2 months ago
For Topic Publishing https://docs.microsoft.com/en-us/azure/event-grid/authentication-overview
upvoted 3 times

ning 2 years, 2 months ago
For web hook delivery https://docs.microsoft.com/en-us/azure/event-grid/secure-webhook-delivery
upvoted 2 times

Question #13 *Topic 6*

HOTSPOT -

You are working for Contoso, Ltd.

You define an API Policy object by using the following XML markup:

```
<set-variable name= "bodySize" value="@(context.Request.Headers["Content-Length"] [0])"/>
<choose>
 <when condition= "@(int.Parse(context.Variables.GetValueOrDefault<string> ("bodySize"))<512000)">
</when>
<otherwise>
     <rewrite-uri template= "/put"/>
     <set-backend-service base-url= "http://contoso.com/api/9.1/"/>
</otherwise>
</choose>
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	○	○
If the body size is >256k, an error will occur.	○	○
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	○	○

Correct Answer:

Answer Area

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	○	○
If the body size is >256k, an error will occur.	○	○
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	○	○

Box 1: Yes -

Use the set-backend-service policy to redirect an incoming request to a different backend than the one specified in the API settings for that operation. Syntax:

<set-backend-service base-url="base URL of the backend service" />

Box 2: No -

The condition is on 512k, not on 256k.

Box 3: No -

The set-backend-service policy changes the backend service base URL of the incoming request to the one specified in the policy.

Reference:

https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies

paru123456789 Highly Voted 2 years, 8 months ago

Answer:
Yes
No
No
upvoted 70 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000.
upvoted 3 times

KingChuang 10 months, 1 week ago
On my exam 2022-12-26.
Chose：Y/N/N
upvoted 5 times

clarionprogrammer 2 years, 6 months ago
This is correct.
upvoted 4 times

Cornholioz Highly Voted 2 years, 10 months ago
It's doing nothing in the When condition. So if msg is <512, it does nothing. The otherwise gets executed only for >512 ??? So it will set backend service only for msgs >512? Question doesn't talk about this case.

Also, like someone said in the other discussion for this question:
technically >256k can be >512k

Poorly framed question!!!
upvoted 13 times

MiraA 2 years, 1 month ago
I think that no action in <when> means nothing needs to be changed (for requests with body < 512000) so the back-end service URL remains as specified in the policy.
And only large requests (> 512000) will be redirected to API 9.1 using <set-backend-service> in <otherwise> branch.

https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies#SetBackendService
upvoted 8 times

Knightie 1 year, 2 months ago
not framed, just tricked.. if body > 512k wther you use 1024 base of 1000 base, nevermind, it just change the base URL and use /put instead.. it does not give an error, that's what otherwise for.
upvoted 2 times

NightshadeRC Most Recent 3 months, 1 week ago
Had this question in today's exam: 2023-07-26
upvoted 4 times

AlexeyG 8 months, 2 weeks ago
Got this in 16/02/2023
upvoted 2 times

edengoforit 1 year, 9 months ago
I think the first question should be backend instead of inbound?
Backend service: <!-- statements to be applied before the request is forwarded to
the backend service go here -->
upvoted 3 times

gabavo 1 year, 7 months ago
rewrite-uri policy can appear only in inbound scope, so the answer is correct
https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies#usage-8
upvoted 3 times

mattvasc 1 year, 8 months ago
Makes sense and if we check the set-var docs on Microsoft we can see that it can be used under inbound, outbound or backend scope:

https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication
upvoted 1 times

lugospod 1 year, 9 months ago
Got this one 01/2022. Went with yes no no
upvoted 5 times

Pomphard 2 years, 6 months ago
For statements 2 & 3, I think they try to emphasize whether something will happen or may happen
>256K may be larger than 512K, so it may happen but there's no guarantee as the size might be in between these values
The base URL may be retained if the condition is met, but there's no guarantee as it might be rewritten
So I'll go with yes, no, no

upvoted 3 times

Araneus 2 years, 11 months ago

I would say "it depends" for the third question. Since the set-backend-service policy is in the otherwise block, it would only rewrite the URL for requests with a body size of >=512000. Or am I missing a clue here?

upvoted 9 times

Cornholioz 2 years, 11 months ago

I think Araneus is right. Easy to miss. Looks like a basic WHEN / OTHERWISE condition to me. If bodysize limit is not hit, it will never reach the OTHERWISE scope and hence will not retain a higher version. I think.

upvoted 4 times

Zidimirite 2 years, 6 months ago

You mean will retain a higher version, right? Since the OTHERWISE scope doesn't get executed hence the set-backend-service isn't set to 9.1? 9.1 is the lower version.

upvoted 6 times

iiiihhhh 2 years, 10 months ago

There is another discussion for the question: https://www.examtopics.com/discussions/microsoft/view/23300-exam-az-204-topic-5-question-8-discussion/

upvoted 3 times

ahadjithoma 2 years, 10 months ago

Use the set-backend-service policy to redirect an incoming request to a different backend than the one specified in the API settings for that operation. This policy changes the backend service base URL of the incoming request to the one specified in the policy.

Ref: https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies

So it looks like why don't even care if any of the WHEN / OTHERWISE conditions are met, since the 9.2 is not specified anywhere in the given code snippet

upvoted 4 times

deborahQ 1 year ago

to my understanding, if the body size is less then 512K, it could use higher version (if it has and requested, which is not mentioned here).

upvoted 1 times

Question #14 *Topic 6*

You are developing a solution that will use Azure messaging services.
You need to ensure that the solution uses a publish-subscribe model and eliminates the need for constant polling.
What are two possible ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Service Bus

B. Event Hub

C. Event Grid

D. Queue

Correct Answer: *AC*
It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber

Community vote distribution
AC (100%)

Amankothari Highly Voted 3 years, 3 months ago
Correct Ans is A & C
https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber#issues-and-considerations
upvoted 50 times

coffecold 1 year ago
Agree, but publish/subscribe is a feature of event-hub as well, but event-based receiving is not. That is why option B falls off.
upvoted 3 times

rrongcheng 3 years, 3 months ago
"In Azure, consider using Service Bus or Event Grid. "
upvoted 8 times

lion2k6 Highly Voted 2 years, 8 months ago
The article has been updated, the new text says "In Azure, consider using Service Bus, Event Hubs or Event Grid."

So A, B and C should be all correct.

https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber#issues-and-considerations
upvoted 49 times

coffecold 1 year ago
Event Hubs : event-based receiving is not possible, you still need to poll. So B falls off
upvoted 1 times

juanckar Most Recent 3 months, 3 weeks ago
This was on the exam (July 2023). Went with AC. Scored 917
upvoted 2 times

OPT_001122 11 months, 3 weeks ago
Selected Answer: AC
A & C Correct
upvoted 2 times

elequiel 1 year ago
Got it in exam 20/10/2022
upvoted 5 times

AZAdam22 1 year, 4 months ago
Selected Answer: AC
A and C as both service bus and event grid implement publish-subscribe and don't require polling.
upvoted 2 times

renzoku 1 year, 4 months ago

exactly, Event HUB needs to poll and process an event

upvoted 1 times

liuliangzhou 1 year, 4 months ago

A, B and C should be all correct.
It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus, Event Hubs or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka.
https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber

upvoted 2 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose A. Service Bus & C. Event Grid

upvoted 3 times

heisenberg33 1 year, 9 months ago

Selected Answer: AC

Correct answer is A and C Ref: https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber#issues-and-considerations.
Wrong answer B (eliminate the need for constant polling)Ref: https://docs.microsoft.com/en-us/azure/connectors/connectors-create-api-azure-event-hubs#trigger-polling-behavior

upvoted 7 times

wk__ 1 year, 9 months ago

It seems that the selected answers are correct: "Existing technologies. It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus, Event Hubs or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka."

https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber

upvoted 2 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 3 times

kondapaturi 2 years, 4 months ago

Service bus
Event Grid

upvoted 3 times

mprokopy 2 years, 5 months ago

A & C
Triggers of Event hub are long-polling, so B is not correct.
https://docs.microsoft.com/en-us/azure/connectors/connectors-create-api-azure-event-hubs#trigger-polling-behavior

upvoted 6 times

Frakandel 2 years, 5 months ago

I believe MS is rather picky about the distinction between events and messages, that why I go for Service Bus and Queues... However Event Hubs and Grid are also capable of delivering the requested functionality (as far as I understand them... but I'm still studying on them :-))

upvoted 5 times

ferut 2 years, 5 months ago

Agree, see:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services?toc=https%3A%2F%2Fdocs.microsoft.com%2Fen-us%2Fazure%2Fservice-bus-messaging%2Ftoc.json&bc=https%3A%2F%2Fdocs.microsoft.com%2Fen-us%2Fazure%2Fbread%2Ftoc.json

upvoted 2 times

iamstudying 1 year, 7 months ago

Buddies, although there is a distinction between events and messages, all of these services (Service Bus Topics, Service Bus Queues, Event Grid, Event Hubs) are under the "Azure Messaging" umbrella.

Queues do not satisfy the pub-sub model and Event Hub is long polling https://docs.microsoft.com/en-us/azure/connectors/connectors-create-api-azure-event-hubs#trigger-polling-behavior

upvoted 2 times

wtkwsk 2 years, 6 months ago

A, B and C are all correct:

Existing technologies. It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus, Event Hubs or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka.
(https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber)

upvoted 4 times

Dnyaneshwar 2 years, 7 months ago

A & C correct answer

upvoted 2 times

vb3d 2 years, 7 months ago

The link provided in the answer says that Event Grid, Event Hub and Service bus all support pub/sub

upvoted 5 times

Question #15 *Topic 6*

A company is implementing a publish-subscribe (Pub/Sub) messaging component by using Azure Service Bus. You are developing the first subscription application.
In the Azure portal you see that messages are being sent to the subscription for each topic. You create and initialize a subscription client object by supplying the correct details, but the subscription application is still not consuming the messages.
You need to ensure that the subscription client processes all messages.
Which code segment should you use?

A. await subscriptionClient.AddRuleAsync(new RuleDescription(RuleDescription.DefaultRuleName, new TrueFilter()));

B. subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName);

C. await subscriptionClient.CloseAsync();

D. subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);

Correct Answer: *D*
Using topic client, call RegisterMessageHandler which is used to receive messages continuously from the entity. It registers a message handler and begins a new thread to receive messages. This handler is waited on every time a new message is received by the receiver.
subscriptionClient.RegisterMessageHandler(ReceiveMessagesAsync, messageHandlerOptions);
Reference:
https://www.c-sharpcorner.com/article/azure-service-bus-topic-and-subscription-pub-sub/

Community vote distribution
D (100%)

Ritesh073 Highly Voted 3 years, 2 months ago
Agreed on D
Clientsubcriber object initialization is done, so B incorrect.
Here Nothing to do with Rule and closeclient so A and C incorrect
upvoted 61 times

KingChuang 10 months, 1 week ago
On my exam 2022-12-26.
Chose： D
upvoted 3 times

paru123456789 Highly Voted 2 years, 8 months ago
Answer: D
upvoted 7 times

Vmwarevirtual Most Recent 5 months, 1 week ago
Appeared in my exam on 27 May 2023.
I chose 'D'
upvoted 3 times

AlexeyG 8 months, 2 weeks ago
Got this in 16/02/2023
upvoted 2 times

OPT_001122 11 months, 3 weeks ago
Selected Answer: D
D is correct
upvoted 4 times

coffecold 1 year ago
Old library : Microsoft.Azure.ServiceBus v5.2.0
I think nowadays (Azure.Messaging.ServiceBus v7.11.0) you simply would use :
processor = client.CreateProcessor(queueName, new ServiceBusProcessorOptions());
processor.ProcessMessageAsync += MessageHandler;
Still : the anwer is correct
upvoted 3 times

Perplex 1 year, 3 months ago
Selected Answer: D
D is correct
upvoted 4 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose D. subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);

upvoted 4 times

vulht 1 year, 8 months ago

Agreed on D when we use Microsoft.Azure.ServiceBus package.

Furthermore, with the new version (Azure.Messaging.ServiceBus), we use a processor to receive/handle messages.
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-dotnet-get-started-with-queues#add-the-code-to-receive-messages-from-the-queue

upvoted 3 times

Question #16 *Topic 6*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.
You have the following requirements:
- Queue size must not grow larger than 80 gigabytes (GB).
- Use first-in-first-out (FIFO) ordering of messages.
- Minimize Azure costs.

You need to implement the messaging solution.
Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application. Create an Azure VM that is triggered from Azure Storage
Queue events.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Don't use a VM, instead create an Azure Function App that uses an Azure Service Bus Queue trigger.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function

Community vote distribution

B (100%)

Lkk51 Highly Voted 3 years, 4 months ago

the reason that the solution is service bus because of this message "Use first-in-first-out (FIFO) ordering of messages."

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted#foundational-capabilities
Storage account does not have ordering guarantee

upvoted 51 times

xRiot007 3 years, 4 months ago

I also think the explanation should be corrected. Size <80GB + FIFO guarantee is what makes Service Bus a good option, not the existence of a VM or not.

upvoted 19 times

lautaro_sw 3 years, 3 months ago

I think the reason for not using a VM is the 'Minimize Azure costs' requirement

upvoted 17 times

rdemontis 2 years, 7 months ago

I totally agree with you all. Only a doubt for VM usage: In addition to save costs I thing it's not possible to interact with a service bus queue. in than case i believe you have to use a storage queue.

upvoted 1 times

edengoforit 1 year, 9 months ago

Absolutely correct. There are so many questions related to this. FIFO and business transactions. If anything related, it seems to be the Service Bus

upvoted 2 times

Nabeelcp Highly Voted 3 years, 2 months ago

Answer is correct . but explanation is confusing ..

upvoted 7 times

Codenoob 3 years, 1 month ago

It is confusing indeed. But the reason the explanation pointed out about VMs is because of VMs are expensive and won't match with the "Minimize azure costs" requirement. Azure Functions is the best way to handle Service Bus Queue events by utilizing its trigger functionality and providing minimal costs to it.

upvoted 8 times

PoTermin 2 years, 2 months ago

Your explanation sounds good.

upvoted 2 times

applepie Most Recent 3 months ago

got this question today, answer N - 7/30/2023, score 895/1000

upvoted 2 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with NO. Scored 917

upvoted 3 times

OPT_001122 11 months, 1 week ago

Selected Answer: B

B. No - is the correct

upvoted 1 times

elequiel 1 year ago

Got it in exam 20/10/2022

upvoted 2 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.

NO

upvoted 2 times

Udoyen 2 years, 2 months ago

This is the kind of question we need ...

upvoted 4 times

paru123456789 2 years, 8 months ago

Answer: NO

upvoted 1 times

svaza 2 years, 8 months ago

I think the given option of using Azure Storage queue is right, refer this document
https://docs.microsoft.com/en-us/azure/storage/queues/scalability-targets#scale-targets-for-queue-storage
- You can store upto 500TB of data in queue, queue works in FIFO.

The approach of using VM for processing the queue messages is overkill. So given quesition the answer is B.
Still its a stilly question

upvoted 1 times

Tom87 2 years, 6 months ago

Queues in general should be FIFO. But not Azure Storage Queue. Ordering is not guaranteed. See this:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted#foundational-capabilities

upvoted 2 times

_nma 2 years, 11 months ago

Answer: A. Yes
1.Your solution requires the queue to provide a guaranteed first-in-first-out (FIFO) ordered delivery.
2.Your queue size won't grow larger than 80 GB.
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted#consider-using-service-bus-queues

upvoted 1 times

_nma 2 years, 11 months ago

Apologies, wrong discussion. This was meant for the net question.

upvoted 1 times

Question #17 *Topic 6*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.
You have the following requirements:
- Queue size must not grow larger than 80 gigabytes (GB).
- Use first-in-first-out (FIFO) ordering of messages.
- Minimize Azure costs.

You need to implement the messaging solution.
Solution: Use the .Net API to add a message to an Azure Service Bus Queue from the mobile application. Create an Azure Windows VM that is triggered from
Azure Service Bus Queue.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*
Don't use a VM, instead create an Azure Function App that uses an Azure Service Bus Queue trigger.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function

Community vote distribution
B (100%)

cbn Highly Voted 2 years, 8 months ago
Using Service bus is fine, however having a Windows VM does not address the cost requirement.

Answer is 'NO'
upvoted 42 times

Udoyen 2 years, 2 months ago
So if I may ask, if the cost wasn't an issue would Windows VM be okay?
upvoted 6 times

dohnalr 1 year, 12 months ago
Yes. It would be ok
upvoted 3 times

coffecold 1 year ago
I think every solution would be OK then.
upvoted 2 times

neosri Highly Voted 3 years, 5 months ago
answer explanation is wrong. It should be azure service bus queues and not storage queues trigger for Azure functions
upvoted 25 times

applepie Most Recent 3 months ago
got this question today, answer N - 7/30/2023, score 895/1000
upvoted 2 times

OPT_001122 11 months, 1 week ago
Selected Answer: B
Azure Service Bus Queue trigger
upvoted 1 times

hubekpeter 11 months, 1 week ago
Selected Answer: B
Queue size must not grow larger than 80 gigabytes (GB) - yes that's fine
Use first-in-first-out (FIFO) ordering of messages. - yes, you can get it with service bus

Minimize Azure costs & Create an Azure Windows VM that is triggered from Azure Service Bus Queue - Firstly there's nothing like an AzureVM triggering, but Azure Functions triggering instead. Secondly - using vm with queue is more expensive, but of course it depends on multiple factors.

upvoted 3 times

Eltooth 1 year, 4 months ago

Selected Answer: B

B is correct answer.
NO

upvoted 3 times

Dev666 1 year, 6 months ago

Selected Answer: B

I choose B as think A will be expensive

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose B. No

upvoted 1 times

altafpatel1984 1 year, 11 months ago

Here I guess following line meant to mention that VM will unnecessarily increase cost.

"Messages may not be sent to the service consistently."

upvoted 2 times

altafpatel1984 1 year, 11 months ago

If B is answer then question is confusing because theriotically VM is not bad idea since there is no mentioned of cost saving etc.

upvoted 1 times

grappigekerel 1 year, 8 months ago

Requirement clearly says minimize Azure costs

upvoted 1 times

Idkhow 2 years, 4 months ago

why do we need the Azure Windows VM for? Service bus is already fine

upvoted 2 times

silva_831 2 years, 4 months ago

The given answer is incorrect. Service Bus queue meets the requirements

upvoted 1 times

Cholo981 2 years, 4 months ago

StorageQueue doens't guarantees FIFO. Must be ServiceBus.

upvoted 1 times

paru123456789 2 years, 8 months ago

Answer: NO

upvoted 2 times

_nma 2 years, 11 months ago

Answer: A. Yes
1.Your solution requires the queue to provide a guaranteed first-in-first-out (FIFO) ordered delivery.
2.Your queue size won't grow larger than 80 GB.
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted#consider-using-service-bus-queues

upvoted 2 times

Cornholioz 2 years, 10 months ago

Wrong because of the VM use in the solution which is an overkill and increases cost.

upvoted 10 times

javayung 2 years, 12 months ago

Answer is No:
As a solution architect/developer, you should consider using Storage queues when:
Your application must store over 80 GB of messages in a queue.

upvoted 2 times

Dirk 3 years, 3 months ago

Yes - I think explanation is wrong, but answer is correct.
See this link: https://www.serverless360.com/blog/azure-storage-queue-vs-service-bus-queue#:~:text=Message%20Size,incoming%20messages%20throwing%20an%20exception.

upvoted 3 times

cloud_exam1 2 years, 11 months ago

I think the answer is wrong. In the requirement, said 'Queue size must not grow larger than 80 gigabytes (GB)'. And FIFO. So should be Service Bus Queue.

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted
upvoted 1 times

Cornholioz 2 years, 11 months ago

In the famous words of Gandalf, "cloud_exam1, you shall not pass"!
upvoted 7 times

Question #18 *Topic 6*

DRAG DROP -

You are developing a REST web service. Customers will access the service by using an Azure API Management instance.

The web service does not correctly handle conflicts. Instead of returning an HTTP status code of 409, the service returns a status code of 500. The body of the status message contains only the word conflict.

You need to ensure that conflicts produce the correct response.

How should you complete the policy? 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:

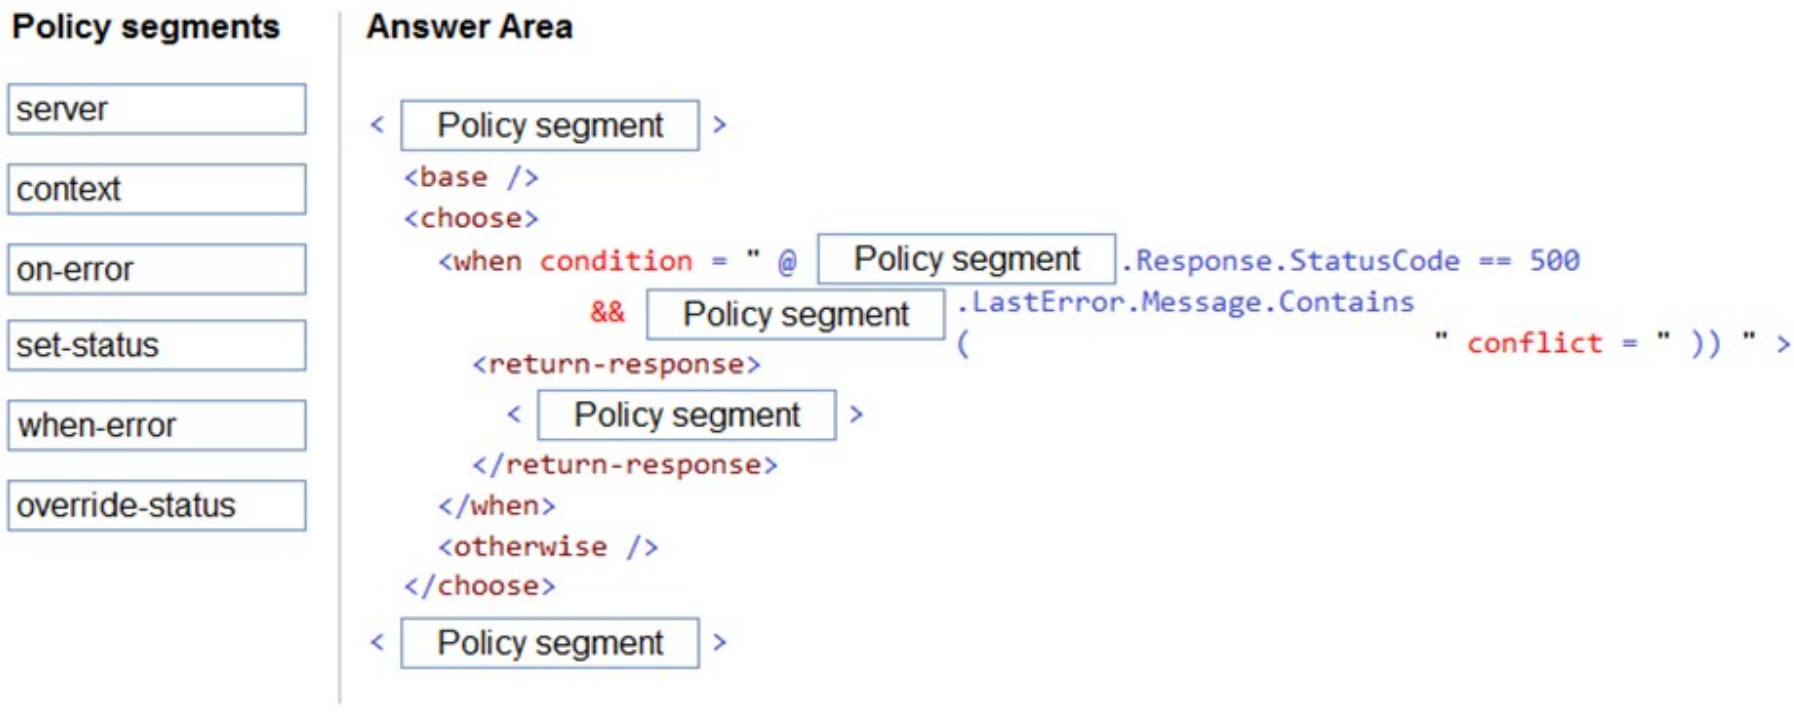

Correct Answer:

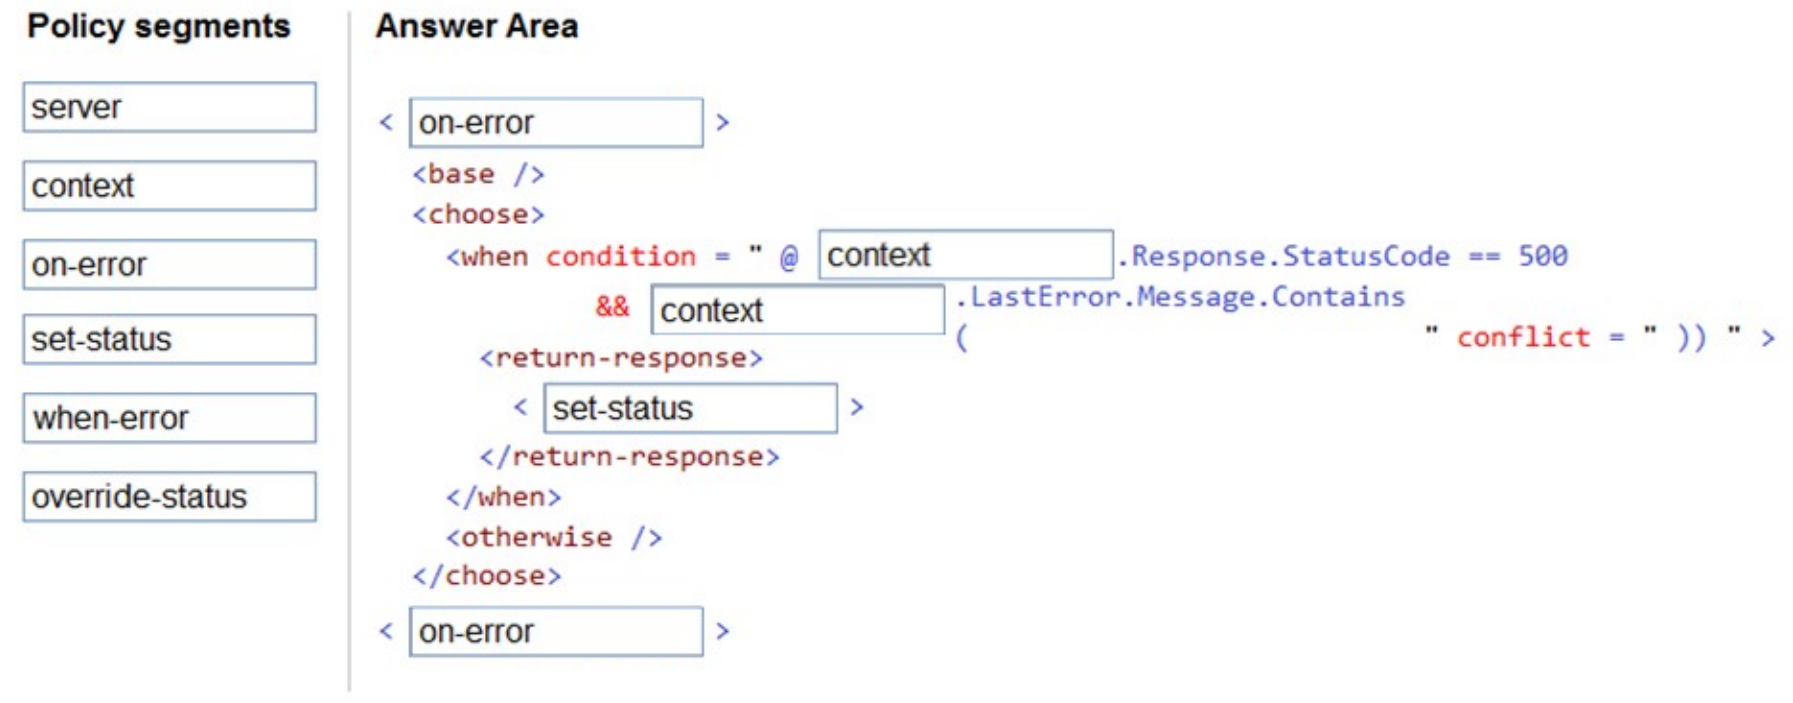

Box 1: on-error -

Policies in Azure API Management are divided into inbound, backend, outbound, and on-error.

If there is no on-error section, callers will receive 400 or 500 HTTP response messages if an error condition occurs.

Box 2: context -

Box 3: context -

Box 4: set-status -

The return-response policy aborts pipeline execution and returns either a default or custom response to the caller. Default response is 200 OK with no body.

Custom response can be specified via a context variable or policy statements.

Syntax:

<return-response response-variable-name="existing context variable">

<set-header/>

<set-body/>
<set-status/>
</return-response>

Box 5: on-error -
Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-error-handling-policies https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies

jokergester Highly Voted 2 years, 7 months ago
Just want to add that "set-status" also has code and reason (e.g. set-status code="409" reason="Conflict").

https://docs.microsoft.com/en-us/azure/api-management/api-management-advanced-policies#SetStatus
upvoted 39 times

SnakePlissken 2 years, 5 months ago
Thanks. set-status on it's own seemed a bit strange.
Minor detail: Last on-error must be preceded by a slash.
upvoted 9 times

armoniyem Highly Voted 1 year, 12 months ago
THE ANSWER IS CORRECT.
upvoted 14 times

jaf19f Most Recent 2 months, 2 weeks ago
I got this question (12-Aug-2023) and I chose the given answers - 932 passed
upvoted 2 times

BaoNguyen2411 3 months ago
got this question on 29/06/2023
upvoted 2 times

AbidooKing 3 months, 2 weeks ago
Got this 2023July13. Went with given and scored 917.
upvoted 2 times

JH81 4 months ago
I got this on 6/29/2023 and passed with 850. Answer is correct.
upvoted 1 times

Indhumathi_Saravanan 7 months ago
Question was on 19-MAR-2023 exam
upvoted 2 times

seila1331 10 months, 3 weeks ago
Did my exam on 8/12/2022. This question was on it.
upvoted 4 times

TheExamMaster2020 11 months, 2 weeks ago
Did my exam on 15th November 2022. This question was on it.
upvoted 3 times

at_sayali 11 months, 3 weeks ago
Got this in 11/2022
Went with on-error, context, context, set-status, override-status, on-error
One more blank space after <otherwise />, I selected override-status
upvoted 5 times

elequiel 1 year ago
Got it in exam 20/10/2022
upvoted 4 times

Question #19 *Topic 6*

DRAG DROP -

You are a developer for a Software as a Service (SaaS) company. You develop solutions that provide the ability to send notifications by using Azure Notification

Hubs.

You need to create sample code that customers can use as a reference for how to send raw notifications to Windows Push Notification Services (WNS) devices.

The sample code must not use external packages.

How should you complete the code segment? 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:

Correct Answer:

Box 1: windows -

Example code:

var request = new HttpRequestMessage(method, $"{resourceUri}?api-version=2017-04"); request.Headers.Add("Authorization", createToken(resourceUri, KEY_NAME, KEY_VALUE));

request.Headers.Add("X-WNS-Type", "wns/raw");

request.Headers.Add("ServiceBusNotification-Format", "windows"); return request;

Box 2: application/octet-stream -

Example code capable of sending a raw notification:

string resourceUri = $"https://{NH_NAMESPACE}.servicebus.windows.net/{HUB_NAME}/messages/"; using (var request = CreateHttpRequest(HttpMethod.Post, resourceUri))

{

request.Content = new StringContent(content, Encoding.UTF8, "application/octet-stream");

request.Content.Headers.ContentType.CharSet = string.Empty;

var httpClient = new HttpClient();

var response = await httpClient.SendAsync(request);

```
Console.WriteLine(response.StatusCode);
}
```

Reference:

https://stackoverflow.com/questions/31346714/how-to-send-raw-notification-to-azure-notification-hub/31347901

markra Highly Voted 2 years, 5 months ago

Answer is correct.
https://docs.microsoft.com/en-us/rest/api/notificationhubs/send-wns-native-notification

upvoted 45 times

lugospod Highly Voted 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 11 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 2 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 2 times

Indhumathi_Saravanan 7 months ago

Got this question - 19-Mar-2023

upvoted 1 times

at_sayali 11 months, 3 weeks ago

Answer is correct.
Got this question in 10/2022

upvoted 2 times

gmishra88 1 year ago

Notification hub and logic apps are removed from scope

upvoted 4 times

Mev4953 1 year, 9 months ago

about second box:
Set to application/json;charset=utf-8 or application/xml.
If the notification type (X-WNS-Type) is wns/raw, set to application/octet-stream

upvoted 7 times

Question #20 *Topic 6*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce
2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*

Reference:

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-programming-guide

Community vote distribution

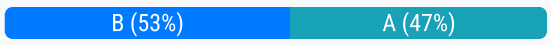

edengoforit Highly Voted 1 year, 9 months ago

When the question comes in the exam, then we need to pray to God

upvoted 46 times

huislaw 1 year, 8 months ago

Let's pray the question doesn't appear.

upvoted 9 times

jay158 Highly Voted 2 years, 4 months ago

Answer is NO
How many partitions event hub can have? NOT MORE than 1024
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-quotas

upvoted 27 times

TaoLu 9 months ago

Come on guys. The examtopics answer is correct.
Remember partition key is different than partition.

upvoted 3 times

troy89 1 year, 8 months ago

But when you use Event Hubs cluster then you can purchase up to 20 Capacity Units and every Capacity Unit can have 2000 partitions.

upvoted 2 times

AndresCH 2 years, 1 month ago

"does not reach goal". See topic5 / question 5

upvoted 2 times

Mev4953 1 year, 9 months ago

Question 5 is not the same. Asking about "Provision an Azure Notification Hub. Register all devices with the hub."

upvoted 1 times

jkes80 2 years, 3 months ago

Indeed, strong reasoning! On a dedicated plan it is max 1024 partitions, on lower plans even less. Because you want to partition per device id, and you have 2000 stores with 1-5 devices per store, you need at least 2000 partitions. So Event Hub is not gonna work. Answer: NO.

upvoted 10 times

dom271219 Most Recent 1 day, 13 hours ago

Selected Answer: B

Azure Event Hub is primarily designed for handling high-throughput, real-time event streaming, such as telemetry data and event-driven architectures. It's not the best choice for storing large volumes of static data like 2 MB of POS device data per device per day

upvoted 2 times

nardk 1 month, 2 weeks ago

Selected Answer: B

Jay158 is rigth

upvoted 1 times

Dianahu 3 months ago

https://www.examtopics.com/discussions/microsoft/view/22503-exam-az-204-topic-6-question-2-discussion/ is not the solution
https://www.examtopics.com/discussions/microsoft/view/25308-exam-az-204-topic-6-question-29-discussion/ also not the solution
so it must be the hub

upvoted 1 times

EliteAllen 3 months ago

Selected Answer: A

Azure Event Hub is suitable for the given scenario because it is designed to handle large-scale data ingestion from numerous devices, allowing efficient correlation based on device identifiers, and can automatically capture and store the data in Azure Blob storage.

upvoted 1 times

adilkhan 8 months ago

Yes is correct

upvoted 1 times

tmuras 8 months, 2 weeks ago

This question was on exam on 17/02/23.
Answered Yes. Score 860

upvoted 4 times

Sulzirsha 8 months, 2 weeks ago

Partition key is consistently used to assign related events to a partition.Partition key doesn't create unique partitions.So answer is 'Yes'.

upvoted 2 times

Twiddling5623 8 months, 3 weeks ago

Selected Answer: B

Answer is No. According to the solution each machine identifier is going to be a partition key and with 1 to 5 machines this will create at least 2000 partitions. The max is 1024 for the biggest price tier.
Ref: https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-quotas#basic-vs-standard-vs-premium-vs-dedicated-tiers

upvoted 1 times

ldenouter 9 months, 2 weeks ago

Selected Answer: A

1000 Partition Key doesn't lead to 1000 partitions.

upvoted 2 times

charthouse 9 months, 3 weeks ago

Is there a difference between device identifier (Solution) and machine identifier (Requirement)?

upvoted 1 times

edmond 10 months, 2 weeks ago

Selected Answer: A

1000 Partition Key doesn't lead to 1000 partitions. As hromadkam mentioned, it's simply something like hashing.

upvoted 1 times

hromadkam 1 year, 2 months ago

The answer is YES.

You don't need as many partitions as there are devices. Partition key is used in function to determine partition (like hashing function). So you could have 2000 devices and for example 10 partitions.

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#publishing-an-event

upvoted 1 times

FatSeal 1 year, 2 months ago

the solution says "Configure the machine identifier as the partition key and enable capture." so you can't really do it your way.

upvoted 1 times

Marek79 11 months, 3 weeks ago

The answer is YES (A).
You can use a partition key to map incoming event data into specific partitions for the purpose of data organization. The partition key is a sender-supplied value passed into an event hub. It is processed through a static hashing function, which creates the partition assignment. If you don't specify a partition key when publishing an event, a round-robin assignment is used.
https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features#mapping-of-events-to-partitions

upvoted 1 times

AZAdam22 1 year, 4 months ago

Selected Answer: A

Yes, because event hubs capture allows you to store event data in blob storage, and by making the machine identifier the partition key the data can be consumed and stored according to to the machine identifier.

upvoted 2 times

Snekerus 1 year, 7 months ago

Selected Answer: B

I think the answer is NO, based on this article:
https://azure.microsoft.com/en-gb/blog/events-data-points-and-messages-choosing-the-right-azure-messaging-service-for-your-data/

"Service Bus queues or topic subscriptions are ideal for this use-case, where the core of the business application lives in the cloud or even an on-site datacenter, branch-offices, work-sites, or service tenants spread across the world. This model is particularly popular with SaaS providers in health care, tax and legal consulting, restaurant services, and retail."

upvoted 3 times

JohnKK 1 year, 8 months ago

Answer is NO - already answered as service bus in #6.

upvoted 1 times

FatSeal 1 year, 2 months ago

that's not really a good reason, the first part says more solutions can be right
"Some question sets might have more than one correct solution, while others might not have a correct solution."

upvoted 2 times

Question #21 *Topic 6*

DRAG DROP -

You are developing an Azure solution to collect inventory data from thousands of stores located around the world. Each store location will send the inventory data hourly to an Azure Blob storage account for processing.

The solution must meet the following requirements:

- Begin processing when data is saved to Azure Blob storage.
- Filter data based on store location information.
- Trigger an Azure Logic App to process the data for output to Azure Cosmos DB.
- Enable high availability and geographic distribution.
- Allow 24-hours for retries.
- Implement an exponential back off data processing.

You need to configure the solution.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Correct Answer:

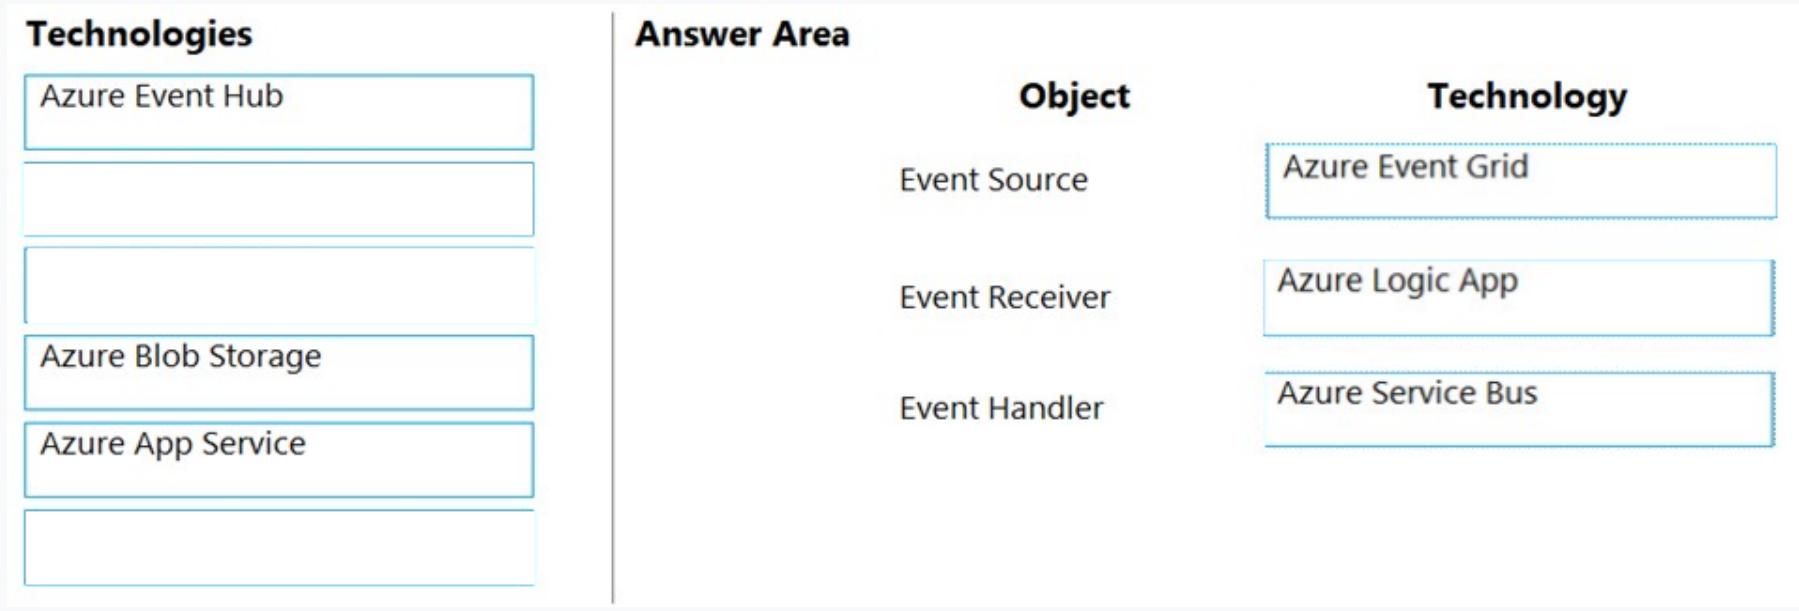

Box 1: Azure Event Grid -

Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener. Event Grid provides reliable event delivery to your applications through rich retry policies and dead-lettering.

Box 2: Azure Logic App -

Event Grid uses event subscriptions to route event messages to subscribers. This image illustrates the relationship between event publishers, event subscriptions, and event handlers.

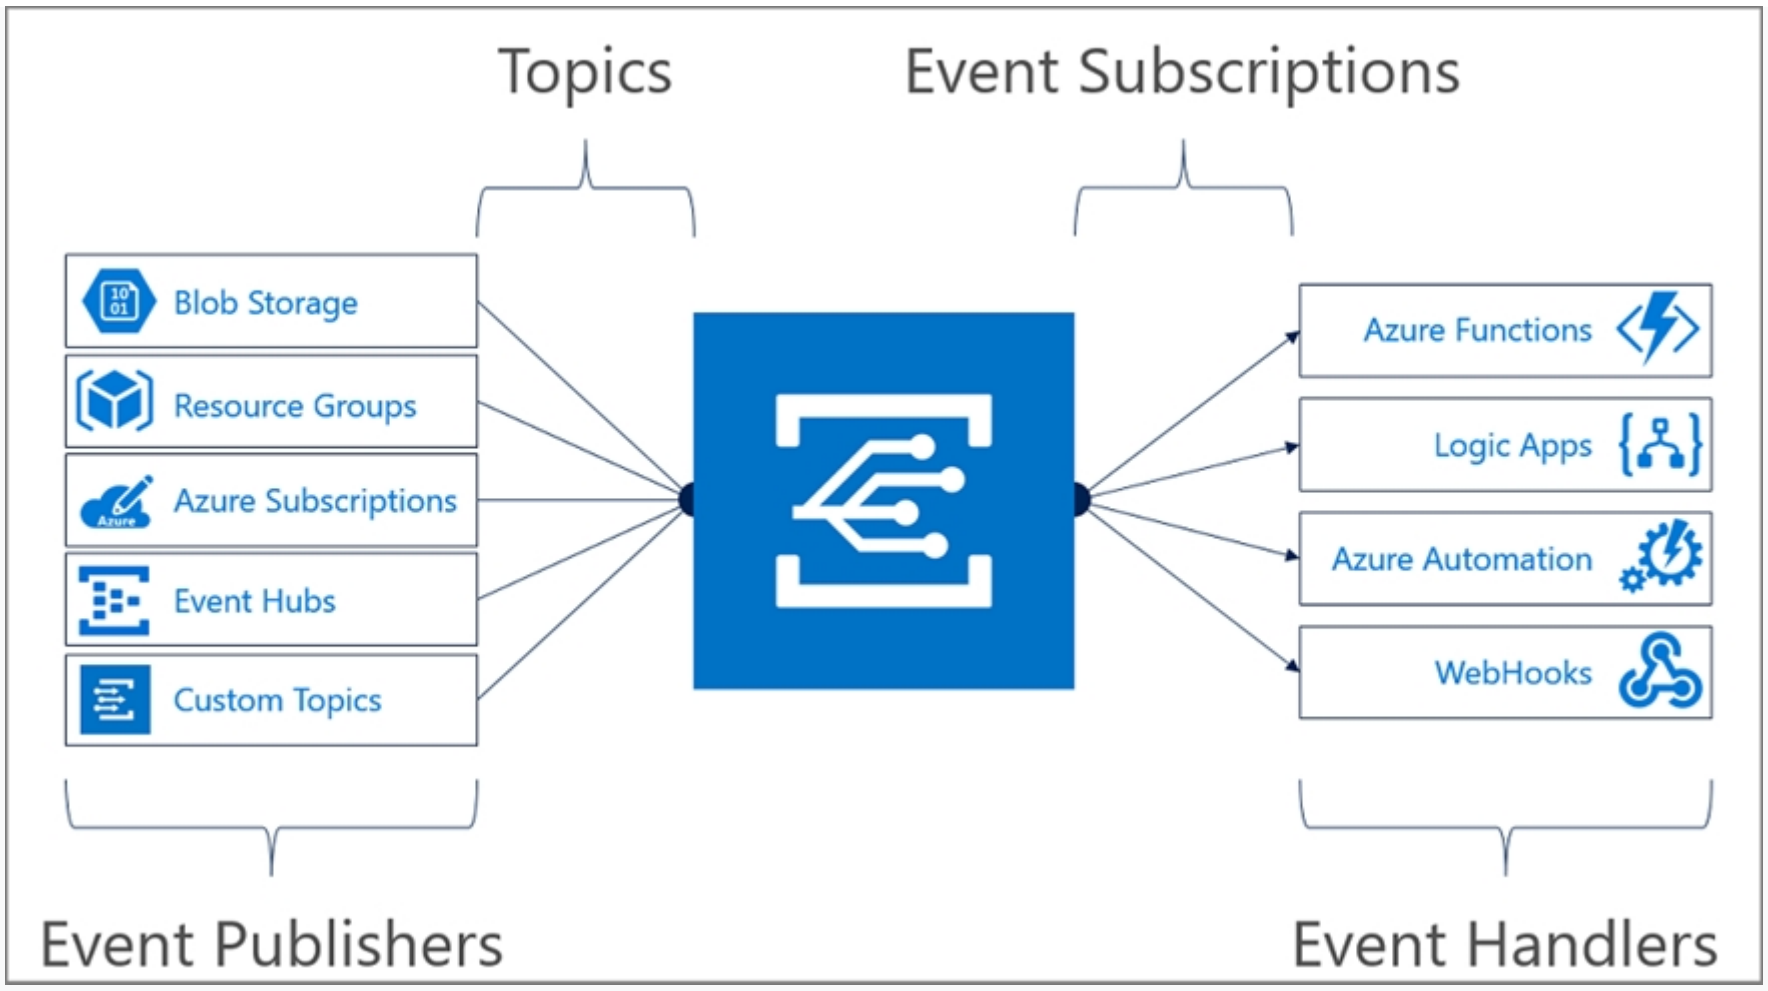

Box 3: Azure Service Bus -
The Event Grid service doesn't store events. Instead, events are stored in the Event Handlers, including ServiceBus, EventHubs, Storage Queue, WebHook endpoint, or many other supported Azure Services.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview https://docs.microsoft.com/en-us/java/api/overview/azure/messaging-eventgrid-readme

Jurgen1234 Highly Voted 2 years, 4 months ago
Source -> blob storage
Receiver -> event grid
Handler -> logic app
upvoted 112 times

paunski7 6 months, 1 week ago
Event Source: Azure Blob Storage
Event Receiver: Azure Event Grid
Event Handler: Azure Logic App

ChatGPT 3.5 paid
upvoted 1 times

jay158 2 years, 3 months ago
Receiver ??
https://docs.microsoft.com/en-us/azure/event-grid/overview
upvoted 6 times

PhilLI 1 year, 9 months ago
Indeed when you look at it from the perspective of the Event Grid the question seems answerable. Your link is good. Receiver is implicitly mentioned here:
"Event sources.
Currently, the following Azure services support sending events to Event Grid."
upvoted 1 times

TonyMel 7 months, 1 week ago
correct, in 2023Mar24, score: 904/1000.
upvoted 2 times

adilkhan 8 months ago
correct ChatGPT
upvoted 1 times

argoth Highly Voted 2 years, 3 months ago
I'm always wondering who is the big brain behind the examtopics question solutions...
upvoted 45 times

JH81 Most Recent 4 months ago
I got this on 6/29/2023 and passed with 850. Went with highly voted answer.

upvoted 2 times

aragones 5 months, 2 weeks ago

Got this 2023-05-12.

FYI:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 1 times

Nhiendo 10 months ago

I thiinh blob storage,event grid, logic app

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, chose Source: Blob storage, Receiver: Event Grid, Handler: Logic App.

upvoted 7 times

vruizrob 11 months, 3 weeks ago

And your response was ok?

upvoted 2 times

MikeAWS 1 week, 4 days ago

how can he knows that? when you pass/fail an exam, all you get is total number of points and some statistics (in %) with chapters!

upvoted 2 times

edengoforit 1 year, 10 months ago

Event Grid is a fully-managed event routing service and the first of its kind. Azure Event Grid greatly simplifies the development of event-based applications and simplifies the creation of serverless workflows. Using a single service, Azure Event Grid manages all routing of events from any source, to any destination, for any application.

upvoted 1 times

d_programmer 1 year, 10 months ago

Receiver should be event hub as the message needs to be retained for 24 hours

upvoted 2 times

koolexam 1 year, 10 months ago

Event Grid -> Reliability - 24-hour retry with exponential backoff to make sure events are delivered.
https://docs.microsoft.com/en-us/azure/event-grid/overview#capabilities

upvoted 3 times

jvyas 1 year, 11 months ago

I think the given answer tries to meet all the requirements solution must meet in your responses.

upvoted 1 times

Franz22 2 years ago

I think that the right solution should be:
Source -> blob storage
Receiver -> event grid
Handler -> logic app
Because:
1)the source of the event, which means WHO has created the event and who is publishing this event, is the Blob Storage
2)The receiver of this event, which means who should receive this event, is the Event Grid. In the event Grid we can create Topic. Topic is where all publishers write messages --> so the Blob Storage write into a topic which is inside an Event Grid.
3)Handler of this event, which means who should handle this event, is the Logic App by creating a subscription in the event grid and connecting the Logic App to that subscription.

upvoted 17 times

catalene 2 years ago

Why not "Event Hub" like Event Receiver?

upvoted 4 times

ning 2 years, 2 months ago

No idea what is event receiver? No where in documentation mentioned that!
Source --> Blob Storage, Handler --> Logic App

upvoted 1 times

ranjitklive 2 years, 2 months ago

https://www.cognizantsoftvision.com/blog/azure-event-grid-vs-event-hubs/

upvoted 3 times

7ack 2 years, 4 months ago

According to the diagram, shouldn't the source be Blob storage, the receiver Event Grid, and the handler Logic app?

upvoted 16 times

Question #22 *Topic 6*

You are creating an app that will use CosmosDB for data storage. The app will process batches of relational data.
You need to select an API for the app.
Which API should you use?

A. MongoDB API

B. Table API

C. SQL API

D. Cassandra API

Correct Answer: *Incorrect Answer:*
For relational data you will need the SQL API
A: The MongoDB API is not used for relational data.
B: The Table API only supports data in the key/value format
D: The Cassandra API only supports OLTP (Online Transactional Processing) and not batch processing.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api

Community vote distribution
C (100%)

Ciupaz 1 month ago

Selected Answer: C
Why not Table API?
The Table API in Azure Cosmos DB is specifically designed for storing semi-structured or NoSQL data in a schema-less fashion, not for storing traditional relational data. Here are some reasons why the Table API is not a good choice for handling relational data.
upvoted 1 times

hicloud 2 months, 1 week ago

Finally a question I can answer without any knowledge of Azure!
upvoted 4 times

escame 8 months, 2 weeks ago

Selected Answer: C
C is correct
upvoted 3 times

mark_1234 8 months, 3 weeks ago

Selected Answer: C
SQL is the only relational data storage type there.
upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: C
C - correct answer
upvoted 3 times

Chief2110 1 year, 1 month ago

Selected Answer: C
C should be the answer
upvoted 3 times

Frefren 1 year, 1 month ago

Selected Answer: C
Relational data
upvoted 4 times

DongKG 1 year, 6 months ago

Selected Answer: C
Because the app will process batches of relational data
upvoted 4 times

Question #23 *Topic 6*

HOTSPOT -

You are developing a .NET application that communicates with Azure Storage.

A message must be stored when the application initializes.

You need to implement the message.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse
(CloudConfigurationManager.GetSetting ("StorageConnectionString"));

[ ▼ ] pVar1 = storageAccount. [ ▼ ] ();

[ ▼ ] pVar2 = pVar1. [ ▼ ] ("contoso-storage");

try
{
  await pVar2.CreateIfNotExistsAsync();
}
catch (StorageException x)
{
  throw;
}
CloudQueueMessage cloudQueueMessage = new CloudQueueMessage("App Launch: <iUserID>");
await pVar2.AddMessageAsync(cloudQueueMessage);
```

CloudQueueClient	CreateCloudQueueClient
CloudTableClient	CreateCloudTableClient
CloudQueue	GetQueueReference
CloudTable	GetTableReference

Correct Answer:

Answer Area

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse
(CloudConfigurationManager.GetSetting ("StorageConnectionString"));

[ CloudQueueClient | CloudTableClient | CloudQueue | CloudTable ] pVar1 = storageAccount.
[ CreateCloudQueueClient | CreateCloudTableClient | GetQueueReference | GetTableReference ] ();

[ CloudQueueClient | CloudTableClient | CloudQueue | CloudTable ] pVar2 = pVar1.
[ CreateCloudQueueClient | CreateCloudTableClient | GetQueueReference | GetTableReference ] ("contoso-storage");

try
{
  await pVar2.CreateIfNotExistsAsync();
}
catch (StorageException x)
{
  throw;
}
CloudQueueMessage cloudQueueMessage = new CloudQueueMessage("App Launch: <iUserID>");
await pVar2.AddMessageAsync(cloudQueueMessage);
```

Reference:

https://docs.microsoft.com/en-us/azure/storage/queues/storage-dotnet-how-to-use-queues?tabs=dotnetv11

BogdanG Highly Voted 1 year, 6 months ago

Correct, just that this is deprecated
For .NET V12 SDK is as follows :
string connectionString = ConfigurationManager.AppSettings["StorageConnectionString"];
QueueClient queueClient = new QueueClient(connectionString, queueName);
queueClient.CreateIfNotExists();

upvoted 13 times

juanckar Highly Voted 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 6 times

huislaw Most Recent 1 year, 6 months ago

Answer is correct:
1. Create the queue client
2. Retrieve a reference to a queue

upvoted 5 times

Question #24 *Topic 6*

HOTSPOT -

A software as a service (SaaS) company provides document management services. The company has a service that consists of several Azure web apps. All

Azure web apps run in an Azure App Service Plan named PrimaryASP.

You are developing a new web service by using a web app named ExcelParser. The web app contains a third-party library for processing Microsoft Excel files.

The license for the third-party library stipulates that you can only run a single instance of the library.

You need to configure the service.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
Set-AzAppServicePlan `
  -ResourceGroupName $rg `
  -Name "PrimaryASP" `
```

[▼]
- NumberOfSites 1
- PerSiteScaling $true
- TargetWorkerCount = 1
- MaxNumberOfWorkers = 1
- SiteConfig.NumberOfWorkers = 1

```
$app = Get-AzWebApp `
  -ResourceGroupName $rg `
  -Name "ExcelParser"
$app.
```

[▼]
- NumberOfSites 1
- PerSiteScaling $true
- TargetWorkerCount = 1
- MaxNumberOfWorkers = 1
- SiteConfig.NumberOfWorkers = 1

```
Set-AzWebApp $app
```

Correct Answer:

Answer Area

```
Set-AzAppServicePlan `
  -ResourceGroupName $rg `
  -Name "PrimaryASP" `
```

[▼]
- NumberOfSites 1
- **PerSiteScaling $true**
- TargetWorkerCount = 1
- MaxNumberOfWorkers = 1
- SiteConfig.NumberOfWorkers = 1

```
$app = Get-AzWebApp `
  -ResourceGroupName $rg `
  -Name "ExcelParser"
$app.
```

[▼]
- NumberOfSites 1
- PerSiteScaling $true
- TargetWorkerCount = 1
- MaxNumberOfWorkers = 1
- **SiteConfig.NumberOfWorkers = 1**

```
Set-AzWebApp $app
```

Reference:

https://docs.microsoft.com/en-us/azure/app-service/manage-scale-per-app

huislaw Highly Voted 1 year, 6 months ago

Answer is correct, -PerSiteScaling $true; $app.SiteConfig.NumberOfWorkers = 1

upvoted 10 times

WH16 Most Recent 2 months, 3 weeks ago

This was on the exam (9th August 2023). Went with proposed.

upvoted 2 times

70PineApple 8 months, 1 week ago

Got this in exam today..20/02/23

upvoted 1 times

rotimislaw 10 months ago

Why PerSiteScaling $true is needed?

upvoted 3 times

unraval 1 month, 1 week ago

Enable per-app scaling with an existing App Service Plan by passing in the -PerSiteScaling $true parameter to the Set-AzAppServicePlan cmdlet.
https://learn.microsoft.com/en-us/azure/app-service/manage-scale-per-app

upvoted 1 times

Heighte 6 months, 3 weeks ago

i guess not to impact other apps on the service plan

upvoted 3 times

OPT_001122 10 months, 3 weeks ago

Correct
PerSiteScaling $true
$app.SiteConfig.NumberOfWorkers = 1

upvoted 2 times

LGWJ12 11 months, 2 weeks ago

yes, the Answer is correct.

upvoted 3 times

Question #25 *Topic 6*

DRAG DROP -

You have an application that provides weather forecasting data to external partners. You use Azure API Management to publish APIs.

You must change the behavior of the API to meet the following requirements:

- Support alternative input parameters
- Remove formatting text from responses
- Provide additional context to back-end services

Which types of policies should you implement? To answer, drag the policy types to the correct requirements. Each policy 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:

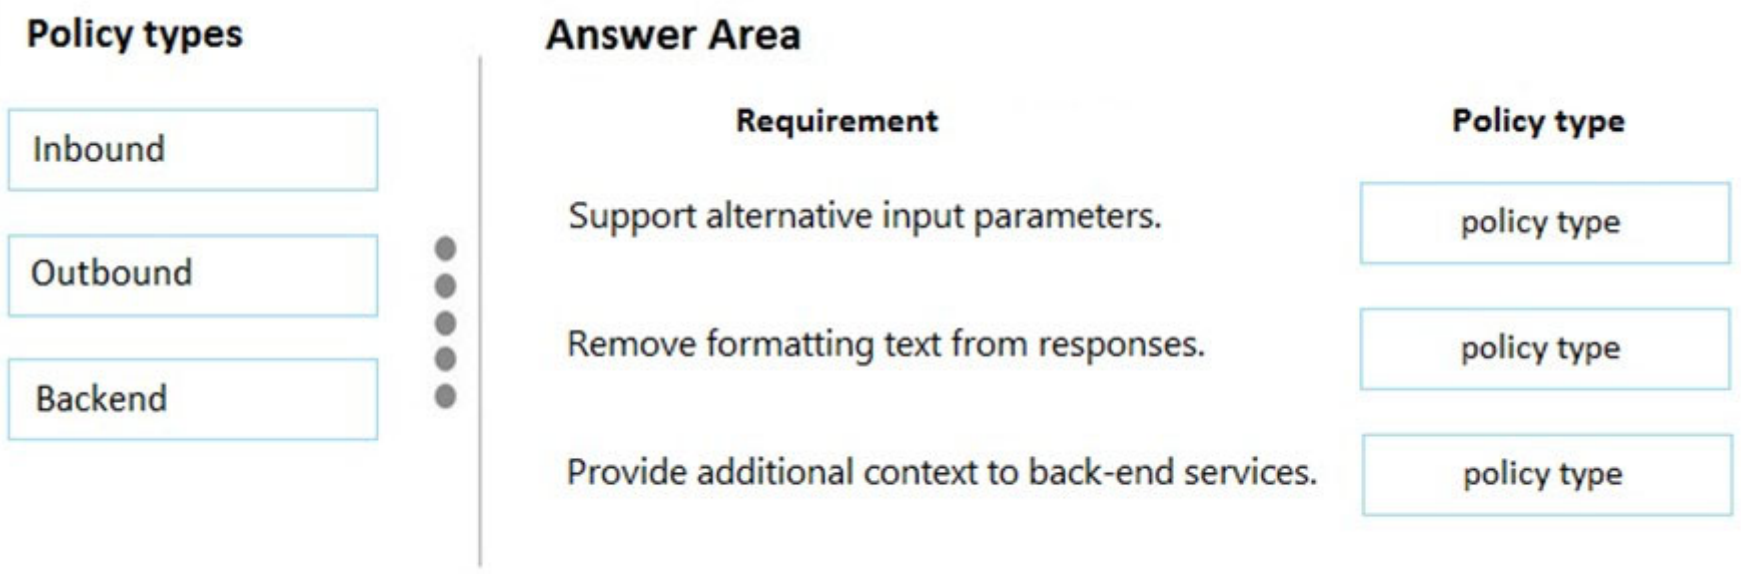

Correct Answer:

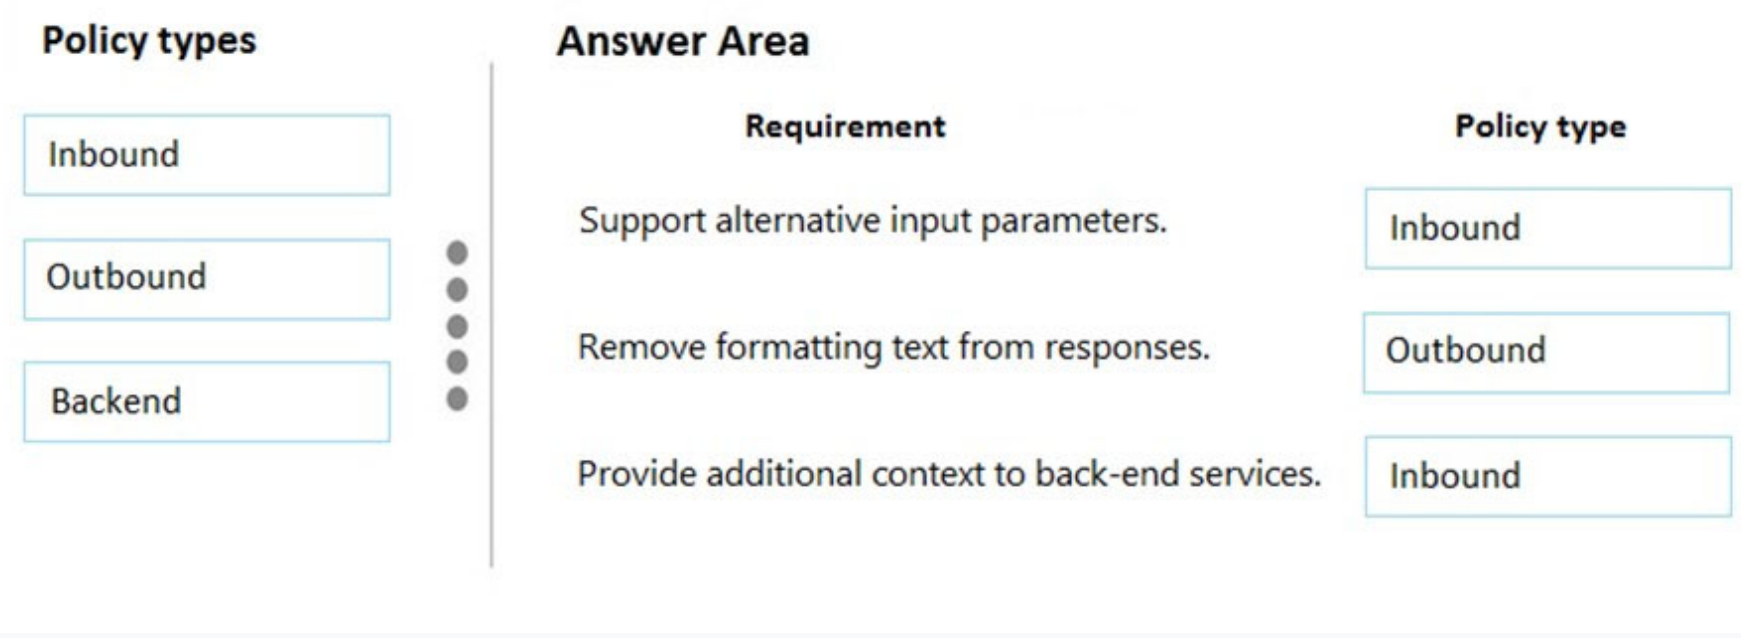

Reference:

https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-policies https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies#forward-context-information-to-the-backend-service

avocado0211 Highly Voted 1 year, 6 months ago

Inbound
Outbound
Backend

upvoted 27 times

g2000 9 months, 2 weeks ago

last one is backend
https://learn.microsoft.com/en-us/azure/api-management/forward-request-policy

upvoted 4 times

uffuchsi 8 months, 1 week ago

Last one has to be Inbound

Below is the sample snippet of the policy to send user context to backend service.

<policies>
<inbound>

<base />
<!-- Forward the name of the product associated with the subscription key in the request to the backend service. -->
<set-query-parameter name="x-product-name" exists-action="override">
<value>@(context.Product.Name)</value>
</set-query-parameter>
<!-- Forward the user id associated with the subscription key in the request as well as the region where the proxy processing the request is hosted. -->
<set-header name="x-request-context-data" exists-action="override">
<value>@(context.User.Id)</value>
<value>@(context.Deployment.Region)</value>
</set-header>
</inbound>

The configuration is part of inbound section.

https://docs.microsoft.com/en-us/azure/api-management/policies/send-request-context-info-to-backend-service

upvoted 1 times

uncledana 3 months, 1 week ago

Backend - https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-policies#-understanding-policy-configuration

upvoted 1 times

srinz Highly Voted 1 year, 1 month ago

Based on https://docs.microsoft.com/en-us/azure/api-management/policies/send-request-context-info-to-backend-service
Inbound is correct answer for the third question

<policies>
<inbound>
<base />
<!-- Forward the name of the product associated with the subscription key in the request to the backend service. -->
<set-query-parameter name="x-product-name" exists-action="override">
<value>@(context.Product.Name)</value>
</set-query-parameter>

<!-- Forward the user id associated with the subscription key in the request as well as the region where the proxy processing the request is hosted.
-->
<set-header name="x-request-context-data" exists-action="override">
<value>@(context.User.Id)</value>
<value>@(context.Deployment.Region)</value>
</set-header>
</inbound>

upvoted 8 times

coffecold 1 year ago

It is adding a parameter and a header in the original request.
<backend>: statements to be applied before the request is forwarded to
the backend service go here
<inbound> statements to be applied to the request go here.
So I think it can be both, the place that makes mores sense is backend.

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

"statements to be applied before the request is forwarded to the backend service go here"
this is not mentioned anywhere in the question hence the confusion for inbound and backend.

upvoted 1 times

nekkilodeon Most Recent 6 months, 1 week ago

The given answer is correct. Inbound, outbound, inbound

upvoted 1 times

adilkhan 6 months, 3 weeks ago

the provided answers are correct! Inbound
Outbound, Inbound https://learn.microsoft.com/en-us/azure/api-management/set-backend-service-policy

upvoted 1 times

hubekpeter 11 months, 1 week ago

I think for the third question, you can use both inbound and backend. Context variable is available in both and it should not matter at what stage you'll add it. But in this article, they're using inbound policy. One newer knows how these exams are evaluated, therefore I'm sticting with inbound as an answer. https://learn.microsoft.com/bs-latn-ba/azure/api-management/policies/send-request-context-info-to-backend-service

upvoted 1 times

le129 1 year, 1 month ago

Answer is correct. https://docs.microsoft.com/en-us/azure/api-management/policies/send-request-context-info-to-backend-service

upvoted 2 times

ivan0590 1 year, 6 months ago

I'm not 100% sure, but based on the following statement I think that the third option should be backend:
"statements to be applied before the request is forwarded to the backend service go here"

Source:
https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-policies#-understanding-policy-configuration

upvoted 4 times

sghaha 1 year, 6 months ago

Korean Page:
https://docs.microsoft.com/ko-kr/azure/api-management/api-management-transformation-policies#forward-context-information-to-the-backend-service

upvoted 3 times

Question #26 *Topic 6*

You are developing an e-commerce solution that uses a microservice architecture.
You need to design a communication backplane for communicating transactional messages between various parts of the solution. Messages must be communicated in first-in-first-out (FIFO) order.
What should you use?

A. Azure Storage Queue

B. Azure Event Hub

C. Azure Service Bus

D. Azure Event Grid

Correct Answer: *A*
As a solution architect/developer, you should consider using Service Bus queues when:
☞ Your solution requires the queue to provide a guaranteed first-in-first-out (FIFO) ordered delivery.
Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted

Community vote distribution
C (100%)

rahulrai19 Highly Voted 2 years, 12 months ago
Answer Should be C, even explanation is saying so
upvoted 159 times

warchoon 7 months, 3 weeks ago
"Transactional" is also a key
upvoted 4 times

TaoLu 9 months ago
Absolutely
upvoted 1 times

Juanlu 2 years, 8 months ago
Yes, I Agree !
upvoted 3 times

cloud_exam1 2 years, 11 months ago
I think so.

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted
upvoted 4 times

Ash111 Highly Voted 2 years, 11 months ago
Choose Service Bus queues if:
You need an At-Most-Once delivery guarantee.
You need a FIFO guarantee.
You need to group messages into transactions.
You want to receive messages without polling the queue.
You need to provide a role-based access model to the queues.
You need to handle messages larger than 64 KB but less than 256 KB.
Your queue size will not grow larger than 80 GB.
You would like to be able to publish and consume batches of messages.
Queue storage isn't quite as feature-rich, but if you don't need any of those features, it can be a simpler choice. In addition, it's the best solution if your app has any of the following requirements.

Choose Queue storage if:
You need an audit trail of all messages that pass through the queue.
You expect the queue to exceed 80 GB in size.
You want to track progress for processing a message inside of the queue.
upvoted 53 times

Ash111 2 years, 11 months ago
C will be the answer...
upvoted 5 times

dddddd111 Most Recent 1 week, 6 days ago

C is the right answer.
upvoted 1 times

nekkilodeon 6 months, 1 week ago

Selected Answer: C

Service bus
https://learn.microsoft.com/en-us/training/modules/discover-azure-message-queue/2-choose-queue-solution
upvoted 1 times

at_sayali 11 months, 3 weeks ago

Correct Answer: C
Got this one 10/2022
upvoted 2 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: C

Azure Service Bus
upvoted 2 times

Satish_Babu 1 year, 2 months ago

Selected Answer: C

C is the correct answer
upvoted 1 times

Pize 1 year, 3 months ago

Selected Answer: C

Another wrong answer on this site
Correct answer C
upvoted 1 times

AZAdam22 1 year, 3 months ago

Selected Answer: C

C - Because service bus holds a FIFO data structure and is used to send business data such as transactions.
upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: C

C is correct answer.
upvoted 1 times

AZAdam22 1 year, 4 months ago

Selected Answer: C

Because service bus guarantees first in first out delivery and is used to send messages between components, for example transferring order processing data.
upvoted 1 times

zhongzi 1 year, 4 months ago

Selected Answer: C

Answer should be C. This question show in the ESI, the answer given C.
upvoted 1 times

ivan0590 1 year, 6 months ago

Selected Answer: C

Admin, please correct the given answer.
It's C.
Even the explanation says is C.
upvoted 1 times

ivan0590 1 year, 6 months ago

Selected Answer: C

Admin, please correct this answer.
C is the correct answer. The explanation in the solution clearly says it's C.
upvoted 1 times

[Removed] 1 year, 6 months ago

Agree with C.
upvoted 1 times

lschuster 1 year, 7 months ago

Selected Answer: C

Highly voted seems correct : C
upvoted 1 times

Baskman 1 year, 8 months ago

Selected Answer: C

FIFO -> Service Bus

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 blobs to Archive tier after they have not been modified within 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:

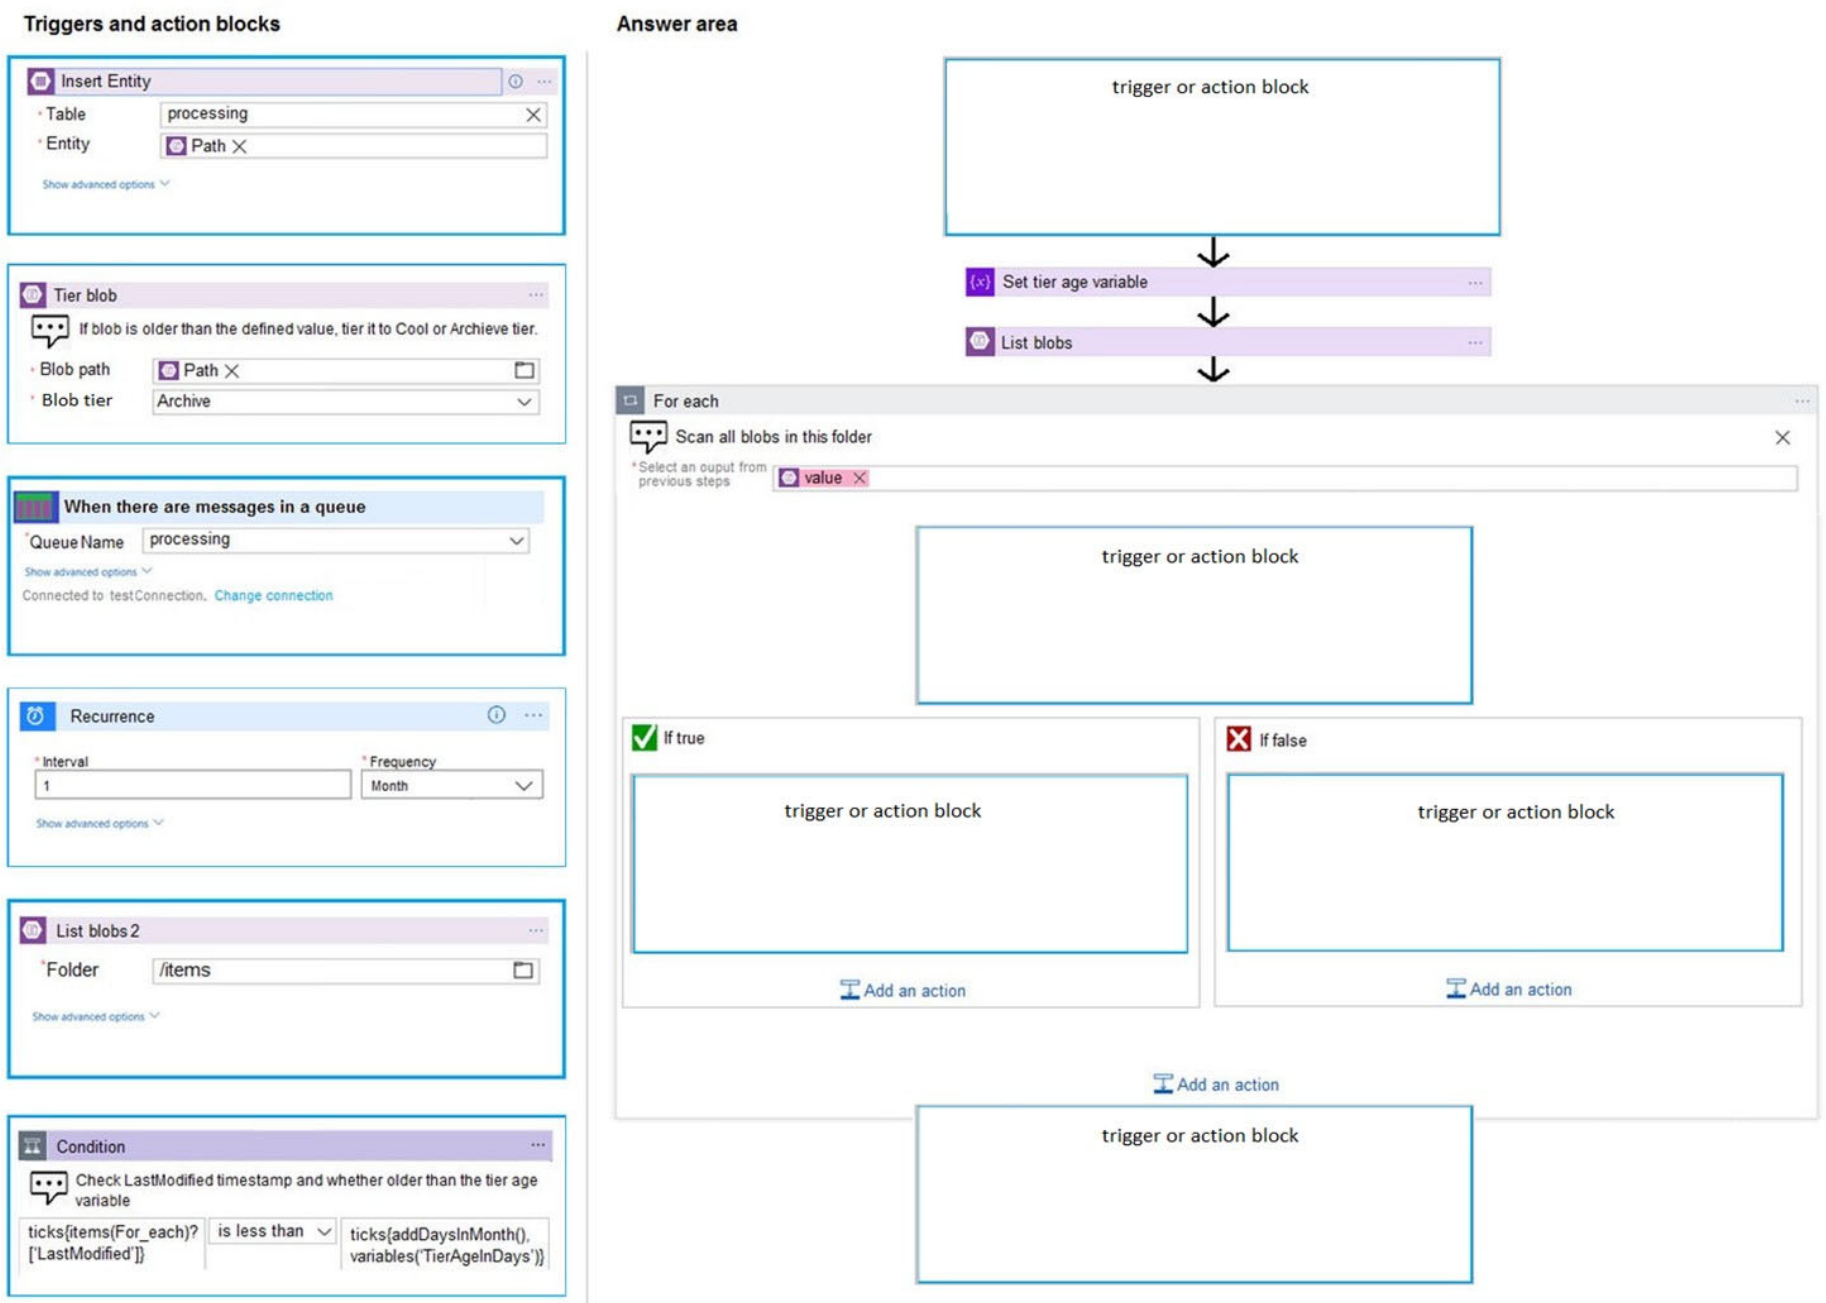

Correct Answer:

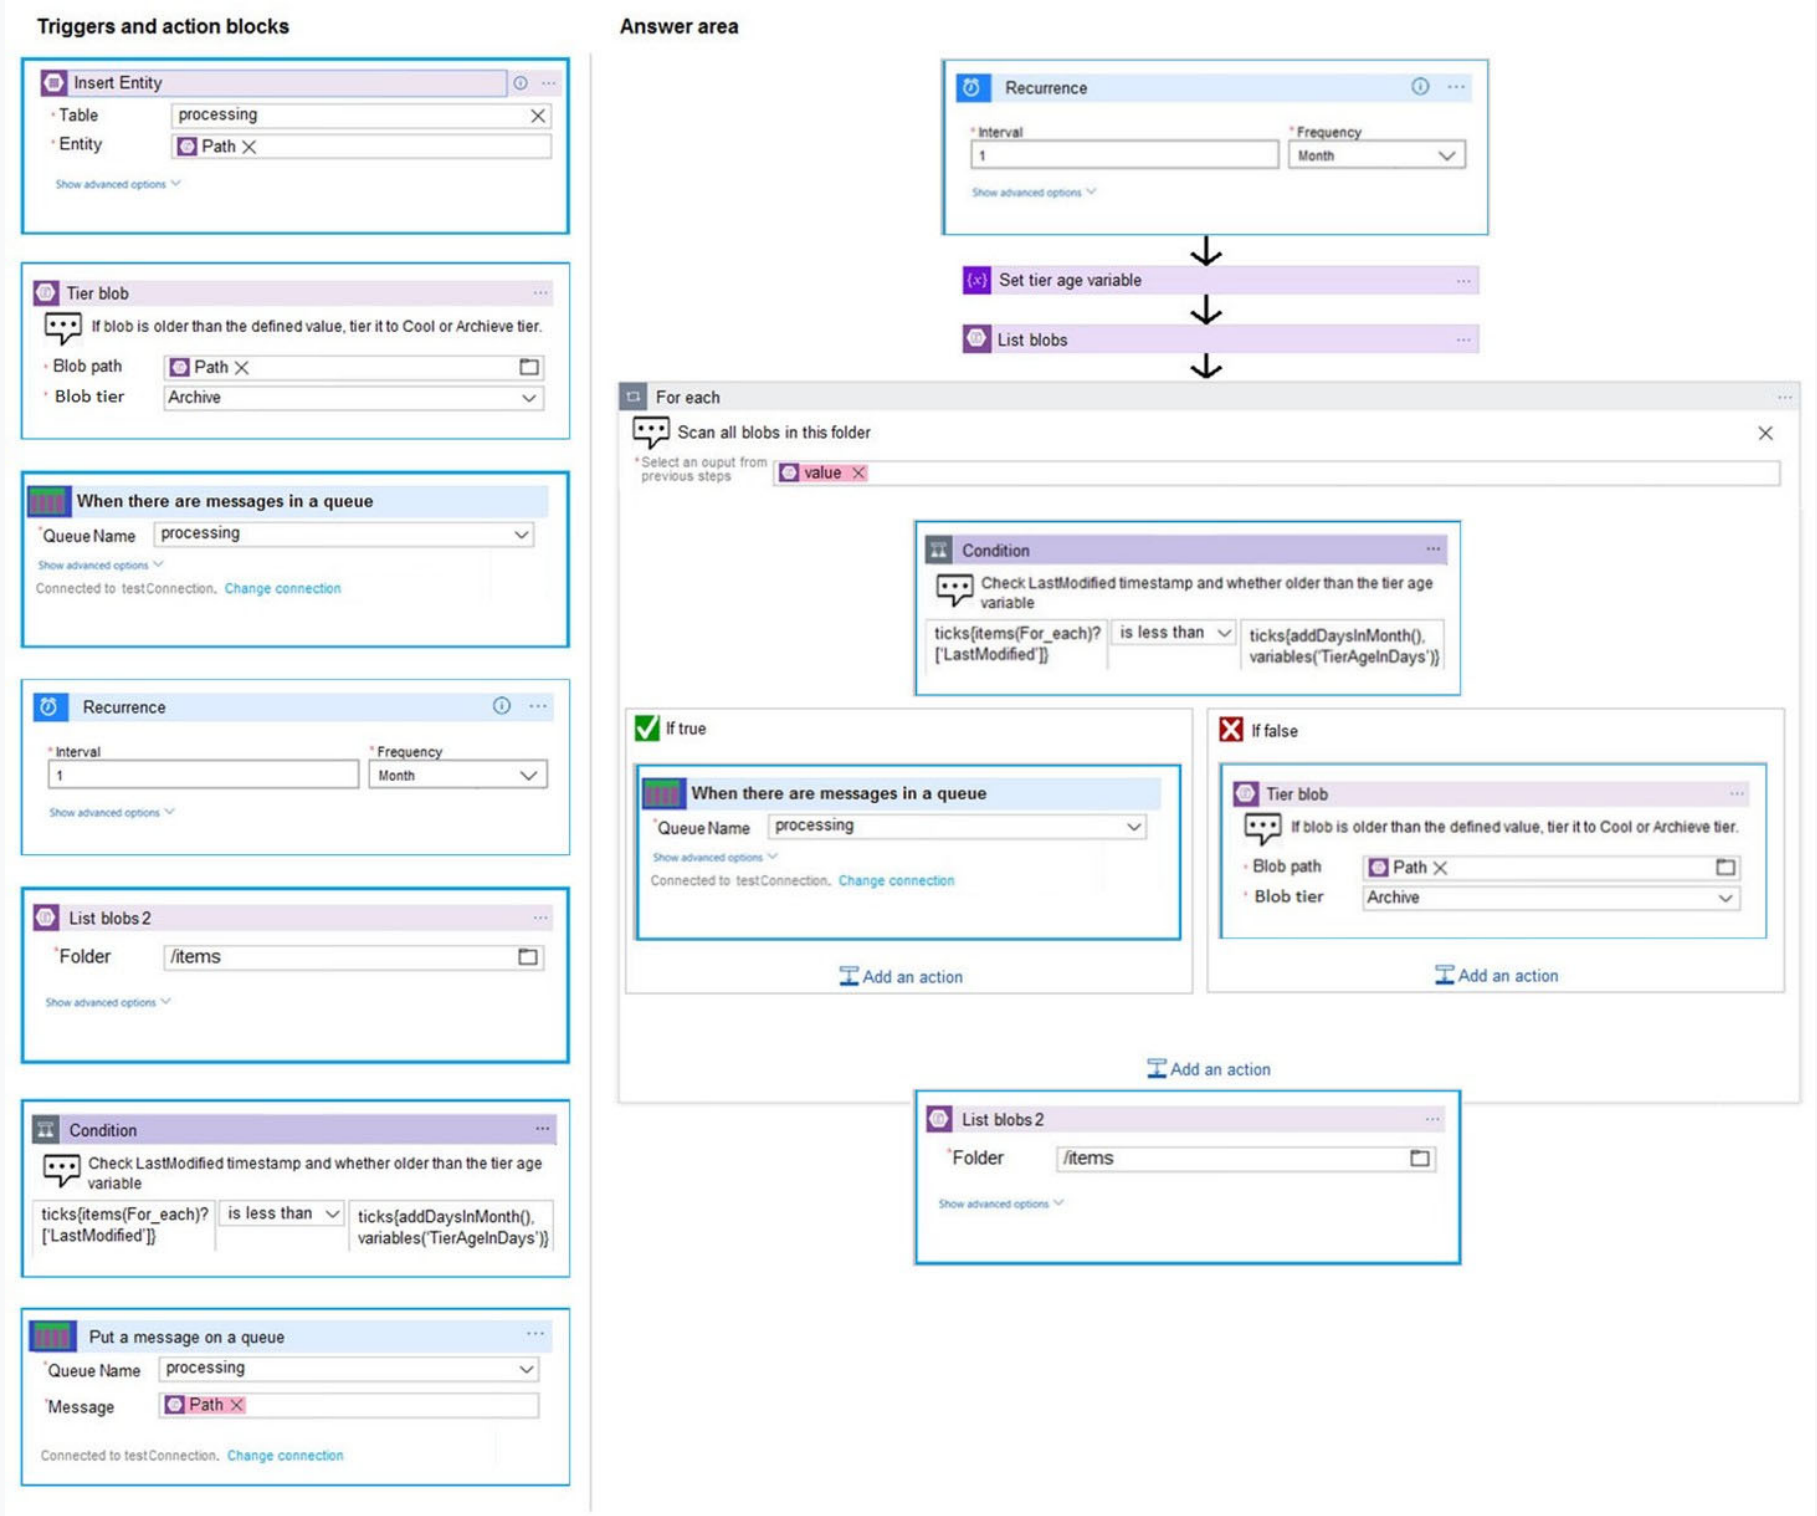

Box 1: Reoccurance..

To regularly run tasks, processes, or jobs on specific schedule, you can start your logic app workflow with the built-in Recurrence - Schedule trigger. You can set a date and time as well as a time zone for starting the workflow and a recurrence for repeating that workflow.

Set the interval and frequency for the recurrence. In this example, set these properties to run your workflow every week.

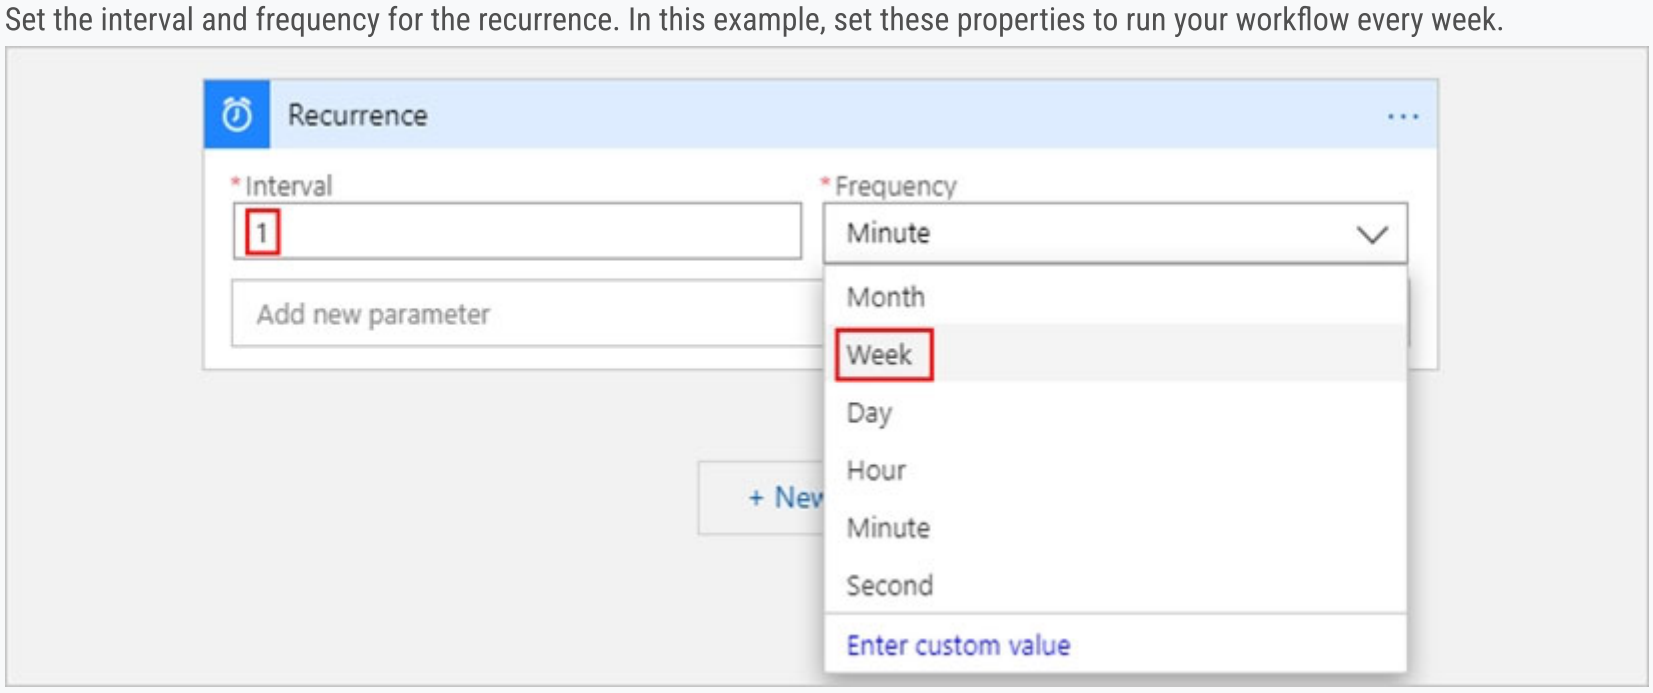

Box 2: Condition..

To run specific actions in your logic app only after passing a specified condition, add a conditional statement. This control structure compares the data in your workflow against specific values or fields. You can then specify different actions that run based on whether or not the data meets the condition.

Box 3: Put a message on a queue -

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

Note: Under If true and If false, add the steps to perform based on whether the condition is met.

Box 4: ..tier it to Cool or Archive tier.

Archive item.

Box 5: List blobs 2 -
Reference:
https://docs.microsoft.com/en-us/azure/connectors/connectors-native-recurrence https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-control-flow-loops https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-control-flow-conditional-statement

Cricketer Highly Voted 2 years, 6 months ago
Logic App is not part of 204 exam since March 2021
upvoted 118 times

if54uran 2 years, 4 months ago
In the official document logic apps are still included.

https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B

Can you provide a source?
upvoted 2 times

nocap 2 years, 4 months ago
In the linked document it actually confirms Logic Apps are no longer part of the criteria in the comparison of this section "Connect to and consume Azure services and
third-party services"
upvoted 6 times

jay158 2 years, 3 months ago
See Notes from 'fabulousethiopia': we still get questions about logic app
https://www.examtopics.com/exams/microsoft/az-204/
upvoted 2 times

SnakePlissken 2 years, 5 months ago
You're right!! Let's not waste our time with questions about Logic Apps anymore!
Please upvote Cricketer!
upvoted 7 times

MariusN Highly Voted 2 years, 12 months ago
For Box 3, refer to the answer's text that says "Put a message on a queue". The answer's graphic implies "When there are messages in a queue" forBox 3 by mistake.
upvoted 40 times

ivan0590 Most Recent 1 year, 6 months ago
Admin, please correct this answer.
The box 3 shown in the image (When there are messages in a queue) differs from the solution (Put a message on a queue).
"Put a message on a queue" is the right answer!
upvoted 3 times

ning 2 years, 2 months ago
No idea, what is the last box ...

For the condition, tick less than is older than ...
upvoted 2 times

crepatata 2 years, 6 months ago
Probably, you should use "Insert Entity" in case it should be archived and in the 5th step archive all the entities in that path?
upvoted 2 times

Pomphard 2 years, 6 months ago
It is clear that :
box 1 should be recurrence
box 2 should be the condition

However, there is confusion about the condition, and rightly so:
the right half, ticks(addDaysInMonth(), variables("TierAgeInDays")) has a few problems,
first off: addDaysInMonth() doesn't exist as a function at all
Second, ticks does not accept two parameters as is done here.
(see reference:https://docs.microsoft.com/en-us/azure/logic-apps/workflow-definition-language-functions-reference#date-and-time-functions)

So, the condition block is bogus. However, they probably wanted to write something like: addDays(utcNow(), variables("TierAgeInDays")) which IS valid.
The less-than will return true for anything older, which leaves the next boxes to be:
box 3: Tier blob
box 4: Message queue

box 5: optional, they probably want to show the results with list blob2
upvoted 12 times

BrettusMaximus 2 years, 6 months ago

Totally agree, however Box 5 should be "Insert Entity" as it occurs for every blob. A good developer would want to track "Processing" and record the state of an archive transaction or not.

upvoted 1 times

XYZ2 2 years, 7 months ago

About condition
"the Ticks property to display the number of ticks that have elapsed since the beginning of the twenty-first century"
Basically, if left side has a bigger value, then it's been updated less then 180 days ago, and should not be archived
Condition says "if left value less than right value", then we have:
If yes (true), it should be archived (left value less than right value, less ticks, older than 180 days)
If no (false) - do not put it to archive (left value greater than right value, more ticks, newer than 180 days)
So, current answer is
Box 1 Recurrence
Box 2 Condition
Box 3 Tier blob (item last modified < current date - 180 days)
Box 4 Put a message on a queue
Box 5 .... still not sure...

upvoted 5 times

gchen 2 years, 7 months ago

BOX 1: Recurrence
BOX 2: Condition
BOX 3: Tier blob ==> condition is older than 180 days
BOX 4: Put a message on a queue
BOX 5: List blob2 ==> check the next one

upvoted 6 times

pieronegri 2 years, 8 months ago

I think the false branch needs "The path for any item that is not archived must be placed in an existing queue." as it is about items which are not archived (last modified is >= adddays(-180)).
the true branch needs archiving.
Last action, listing, is not needed per se. Just to show the results.

upvoted 1 times

minsma 2 years, 8 months ago

I think this is the correct answer: http://prntscr.com/zqmpjl

upvoted 5 times

cbn 2 years, 9 months ago

Box 1 Recurrence
Box 2 Condition
Box 3 Tier blob (item last modified < current date - 180 days)
Box 4 Put a message on a queue
Box 5 When there are messages in a queue / blank

upvoted 18 times

atomicicebreaker 2 years, 6 months ago

Box 3 - you need to tier blob when modified > currentDate - 180 days, and the condition is set to "less than" so true -> blobs to ignore, false -> blobs to tier

upvoted 3 times

Juanlu 2 years, 8 months ago

I thing, after analyse all comments, this is the good one !

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

This is the right answer.

upvoted 1 times

ning 2 years, 2 months ago

No, you cannot do multiple trigger in the designer
In code view, you can

upvoted 1 times

Archimedes 2 years, 9 months ago

The condition is "LastModified" "less than" (Date - 180). (TierAgeInDays is a negative value of -180). If this statement is true, it means that the blob has been modified before 180 days from now. Hence it should be archived. So boxes 3 and 4 should be swapped. Correct me if I am wrong. Box 5 is immaterial. Not covered in the question scope. So, I am Ok with listing the blobs for Box 5.

upvoted 7 times

jvyas 2 years, 1 month ago

"You must automatically move blobs to Archive tier after they have not been modified within 180 days. " If condition block returns true meaning the block has been modified within 180 days and should not be archived.

upvoted 1 times

jvyas 2 years, 1 month ago

Sorry I did not realize it was -180.

upvoted 1 times

Kibb 2 years, 9 months ago

Aren't 3 and 4 swapped?
The question sais place in queue if the messages AREN'T archived (so not older than 180 days)

upvoted 5 times

rajwit 2 years, 9 months ago

there is different discussion
https://www.examtopics.com/exams/microsoft/az-203/view/7/
I believe
Box 1 Recurrence
Box 2 Condition
Box 3 Put a message on a queue
Box 4 Tier blob
Box 5 List Blob

upvoted 8 times

xofowi5140 2 years, 9 months ago

But the workflow is not the same

upvoted 1 times

sleepyboy777 2 years, 9 months ago

It might make sense to List blobs 2 to see what files are there at the end of the operation, but the question does not specify that

upvoted 1 times

reggina 2 years, 10 months ago

for the condition, True means the blob is older than the limit so 3 and 4 should be switched.
https://docs.microsoft.com/en-us/azure/logic-apps/workflow-definition-language-functions-reference#ticks

upvoted 15 times

borntolearn 2 years, 9 months ago

I disagree because if you read the condition, it is saying that if the item is less than age(180) which is different from the condition's description.

upvoted 3 times

[Removed] 2 years, 9 months ago

No, it says "if the item was modified at a time more than 180 days ago". Notice that "TierAgeInDays" is a negative value of -180. So if you add it to the current date it is 180 days ago. In Pesudcode this would be : If("LastModified" < Today-180d)

upvoted 9 times

internalK 2 years, 9 months ago

Yes, condition is Ticks(item.LastModified)<Ticks(today-180) --> Box3 & Box4 should be switched. If the condition is true means that the item must be archived.

upvoted 7 times

Member777 2 years, 10 months ago

Box 1 Recurrence
Box 2 Condition
Box 3 Put a message on a queue
Box 4 Tier blob
Box 5 When there are messages in a queue

upvoted 11 times

Cornholioz 2 years, 10 months ago

Wrong for 4 & 5. The WHEN is not a condition like in a WHEN block of code. It gets processed after the IF condition. So in the IF condition for true, you want to put into the queue already. Box 5 should either be empty or (just for fun) list the blobs.

upvoted 3 times

xofowi5140 2 years, 9 months ago

Cornholioz what is the correct order for you?

upvoted 1 times

cbn 2 years, 9 months ago

I think it should be
Box 3 Tier blob (item last modified < current date - 180 days)
Box 4 Put a message on a queue

upvoted 9 times

altafpatel1984 1 year, 11 months ago

Please don't misguide people.

upvoted 2 times

Question #28 *Topic 6*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.
You have the following requirements:
- Queue size must not grow larger than 80 gigabytes (GB).
- Use first-in-first-out (FIFO) ordering of messages.
- Minimize Azure costs.

You need to implement the messaging solution.
Solution: Use the .Net API to add a message to an Azure Service Bus Queue from the mobile application. Create an Azure Function App that uses an Azure
Service Bus Queue trigger.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *A*
You can create a function that is triggered when messages are submitted to an Azure Storage queue.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function

Community vote distribution
A (90%) 10%

gunencali Highly Voted 2 years, 10 months ago
Yes, is correct answer
upvoted 34 times

kemtin Highly Voted 2 years, 11 months ago
i enjoys questions on azure functions a lot :)
upvoted 17 times

betepah987 2 years, 8 months ago
Probably, cuz it's common to use so easy to answer.
upvoted 3 times

uffuchsi Most Recent 8 months, 1 week ago
Selected Answer: A
Solution meetss the goal!
upvoted 1 times

nvtienanh 11 months ago
On exam December 2, 2022
upvoted 3 times

gmishra88 1 year ago
How is the first requirement satisfied? How will the queue size not exceed 80 GB with Service bus queue and functions? But I guess the majority of people attempting this exam are already in the Microsoft stable and do not find anything missing with these pattern of questions. Sympathies
upvoted 1 times

coffecold 1 year ago
If messages exceed 80 GB in the queue , additional incoming messages will be rejected and the caller receives an exception.
upvoted 1 times

Perplex 1 year, 3 months ago
Selected Answer: A
A is correct
upvoted 2 times

AZAdam22 1 year, 3 months ago

Selected Answer: B

No - Because you should use storage queues instead, as this provides FIFO ordered delivery and is cheaper than service bus, therefore minimising costs.

upvoted 1 times

Knightie 1 year, 2 months ago

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted

Storage Queue no FIFO see the comparison table.

upvoted 2 times

scrafy 1 year, 1 month ago

Some people get high before leaving a comment

upvoted 4 times

Eltooth 1 year, 4 months ago

Selected Answer: A

A is correct answer.

upvoted 2 times

xRiot007 1 year, 7 months ago

The answer is Yes, but the link is wrong. It is not important how the implementation looks. What is important is that you satisfy those 3 conditions in the problem, which you can find in a much better link, here: https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-service-bus-trigger

upvoted 1 times

xRiot007 1 year, 7 months ago

Correction. Link is here: https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose A. Yes

upvoted 2 times

Freidrich 1 year, 8 months ago

Selected Answer: A

The correct answer is A: Yes.

upvoted 2 times

Netspud 1 year, 8 months ago

Selected Answer: A

It is correct

upvoted 2 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with YES

upvoted 4 times

MiraA 2 years, 1 month ago

The URL link in the answer mentions Azure Function and Storage Queue.
This link is better - Azure Function and Service Bus Queue:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-service-bus-trigger

upvoted 3 times

fbbbbbb 2 years, 5 months ago

Correctomundo

upvoted 2 times

paru123456789 2 years, 8 months ago

Answer: Yes

upvoted 2 times

priya1 2 years, 9 months ago

A. Yes is correct

upvoted 7 times

Question #29 *Topic 6*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Notification Hub. Register all devices with the hub.
Does the solution meet the goal?

A. Yes

B. No

Correct Answer: *B*

Instead use an Azure Service Bus, which is used order processing and financial transactions.

Reference:

https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

Community vote distribution

B (100%)

Justing_Gao Highly Voted 3 years, 3 months ago

Azure event Hub is the best choice

upvoted 35 times

Genere 3 years, 2 months ago

Azure Event Hub is for telemetry and distributed data broadcasting, while Azure Service Bus can be used for order processing and financial transactions.
What is clearly specified in the statement: "You are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores ...". It's Service Bus

upvoted 46 times

quokka 3 years, 2 months ago

Event Hub won't work as message size is currently limited to 1MB (unless each device will upload more than once a day or you put an API in front to pre-process data to two 1MB blocks).

upvoted 2 times

Adol 2 years, 9 months ago

Yes but we're talking about 2 MB in 24 hours and not every message. So still Azure Event Hub is the best choice and should be the primary choice when it comes to IoT or data from many devices

upvoted 5 times

clarionprogrammer 2 years, 6 months ago

Event Hub is incorrect. Event Hub is for streaming real-time data.
https://azure.microsoft.com/en-us/services/event-hubs/

upvoted 1 times

clarionprogrammer 2 years, 6 months ago

Service Bus is designed for financial transactions.
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

upvoted 1 times

Nabeelcp 3 years, 2 months ago

I think this scenario is more suited on message based not event based notification

upvoted 8 times

Knightie 1 year, 2 months ago

It is Event Hub, this is a main differentiator between Event Hub and Service Bus, technically Event Hub is pushing and Service Bus is pulling, so Event Hub is more for Publisher oriented use, but assuming the subscriber can digest the message in time, else exception thrown. Service bus is pulling, so messages will stack until the subscriber pick up the next job. This question has a huge demand of data pushing.. and it didn't mention on subscriber, so Event Hub, see the hybrid model for the use of both. Answer is correct
https://www.linkedin.com/pulse/azure-event-grid-vs-service-bus-eklavya-gupta/

upvoted 3 times

Knightie 1 year, 2 months ago

I mean, the answer is A, it is correct

upvoted 1 times

Knightie 1 year, 2 months ago

My apology, event hub is out because the size is limited to 1 Mb.

upvoted 1 times

Knightie 1 year, 2 months ago

but it said the data size is over 24 hours.... I am not sure.. but to be that's why we have event hub.

upvoted 1 times

RaviNikkam 1 year, 5 months ago

Answer is B, but however it is does not seem like Event Hubs because of the fact, they mentioned "Device data must be correlated based on a device identifier." Correlation id is a property of ServiceBus. So it should be Azure Service Bus.

upvoted 1 times

adilkhan 9 months ago

https://learn.microsoft.com/en-us/dotnet/api/azure.messaging.eventhubs.eventdata.correlationid?view=azure-dotnet

upvoted 1 times

Camios Highly Voted 3 years, 3 months ago

Explanation should be that notification hub is for pushing to mobile devices?

upvoted 25 times

AnkanG 3 years, 3 months ago

True, notification hub is for pushing the data, to a mobile device, not for collecting the data.

upvoted 9 times

Eltooth Most Recent 1 year, 4 months ago

Selected Answer: B

B is correct answer.
Event Hub would be a solution here NOT Notification Hub.

upvoted 2 times

Evo_Morales 1 year, 6 months ago

IMO, not a well-worded question, but since Svc Bus (Prem) is limited to 1MB/message and they specify each device may send up to 2mb / day, but don't specify whether uploads are done at the end of the day or streamed throughout the day in smaller chunks, you have to assume the former and so it must be Event Hub. It's dated, but here is a good article on the two services for general knowledge:
https://www.serverless360.com/blog/azure-event-hubs-vs-service-bus

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose B. No

upvoted 3 times

Adediwura 1 year, 7 months ago

Azure notification hun is used to push messages to

upvoted 1 times

Netspud 1 year, 8 months ago

Selected Answer: B

Answer is no

upvoted 1 times

Mev4953 1 year, 9 months ago

Selected Answer: B

B is better choice

https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

upvoted 1 times

kondapaturi 2 years, 4 months ago

Answer is NO

upvoted 2 times

rdemontis 2 years, 7 months ago

In my opinion EventHub is the best choice for the following reason:
1. Faster in elaborating millions of data every day.
2. No problem for future expansions of POS as assumed in the question
3. It is a slender system than Service bus
4. It is not required advanced features like transactions, ordering, duplicate detection, and instantaneous consistency as service bus provide
5. The requirement here is to manage huge quantity of streaming data

In a similar question (i don't remember whether on udemy or testpreptraining web site) the correct answer is Event Hub.

upvoted 3 times

deborahQ 1 year ago

I think most of POS data are transactions and should detect duplication. and event it has more than 15000 devices, the volume is still not comparable with most of telemetry data. I prefer Service bus

upvoted 1 times

Chulbul_Pandey 2 years, 7 months ago

Az event hub is a better choice

upvoted 3 times

paru123456789 2 years, 8 months ago

Answer: No

upvoted 2 times

aditya_2016 2 years, 9 months ago

https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services#comparison-of-services

upvoted 5 times

tom999 2 years, 9 months ago

No - Notification Hub is for outbound communication

upvoted 7 times

Question #30 *Topic 6*

You are building a loyalty program for a major snack producer. When customers buy a snack at any of 100 participating retailers the event is recorded in Azure
Event Hub. Each retailer is given a unique identifier that is used as the primary identifier for the loyalty program.
Retailers must be able to be added or removed at any time. Retailers must only be able to record sales for themselves.
You need to ensure that retailers can record sales.
What should you do?

A. Use publisher policies for retailers.

B. Create a partition for each retailer.

C. Define a namespace for each retailer.

Correct Answer: *A*
Event Hubs enables granular control over event publishers through publisher policies. Publisher policies are run-time features designed to facilitate large numbers of independent event publishers. With publisher policies, each publisher uses its own unique identifier when publishing events to an event hub.
Incorrect:
Not C: An Event Hubs namespace is a management container for event hubs (or topics, in Kafka parlance). It provides DNS-integrated network endpoints and a range of access control and network integration management features such as IP filtering, virtual network service endpoint, and Private Link.
Reference:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features

Community vote distribution
A (100%)

finnishr Highly Voted 1 year, 1 month ago
Given answer is correct!
upvoted 8 times

KingChuang 10 months, 1 week ago
On my exam 2022-12-26.
Chose： A
upvoted 5 times

paunski7 Most Recent 6 months, 1 week ago
ChatGPT:
Sorry, but the correct answer is B.

Creating a partition for each retailer ensures that each retailer only sees the events that pertain to them, as each partition in Event Hubs is an independent stream of events. This also allows for easy addition and removal of retailers, as you can simply add or remove partitions accordingly.

Using publisher policies for retailers or defining a namespace for each retailer are not the correct solutions for this scenario.
upvoted 1 times

macobuzi 2 months ago
ChatGPT is not a reliable source for Azure.
upvoted 3 times

nekkilodeon 6 months, 1 week ago
B is incorrect.
The number of partitions is specified at the time an Event Hubs is created and can't be changed.
https://learn.microsoft.com/en-us/training/modules/azure-event-hubs/2-event-hubs-overview
upvoted 2 times

OPT_001122 11 months, 1 week ago
Selected Answer: A
A. Use publisher policies for retailers.
upvoted 2 times

Question #31 *Topic 6*

DRAG DROP -

You develop and deploy a web app to Azure App Service in a production environment. You scale out the web app to four instances and configure a staging slot to support changes.

You must monitor the web app in the environment to include the following requirements:

- Increase web app availability by re-routing requests away from instances with error status codes and automatically replace instances if they remain in an error state after one hour.
- Send web server logs, application logs, standard output, and standard error messaging to an Azure Storage blob account.

You need to configure Azure App Service.

Which values should you use? To answer, drag the appropriate configuration value to the correct requirements. Each configuration value 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:

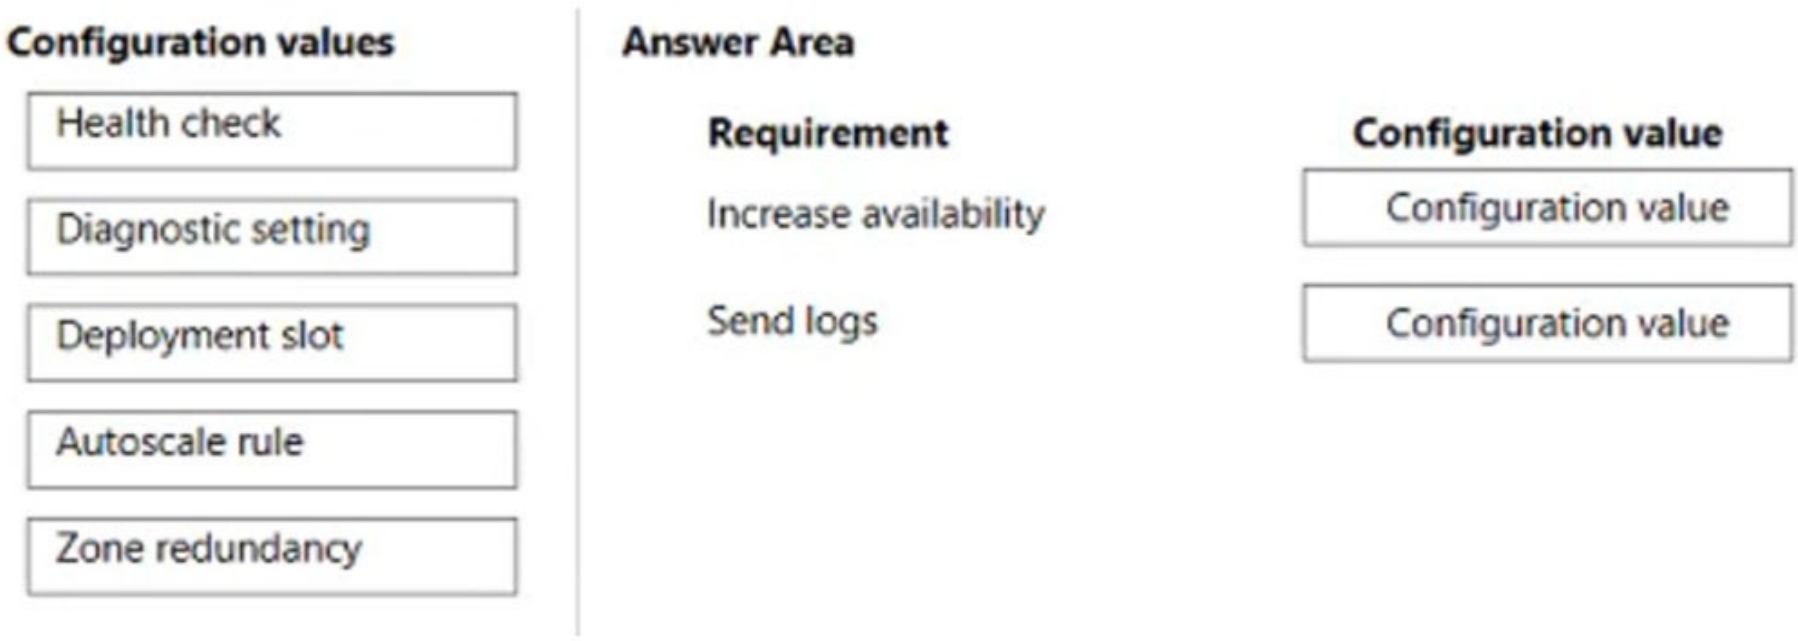

Correct Answer:

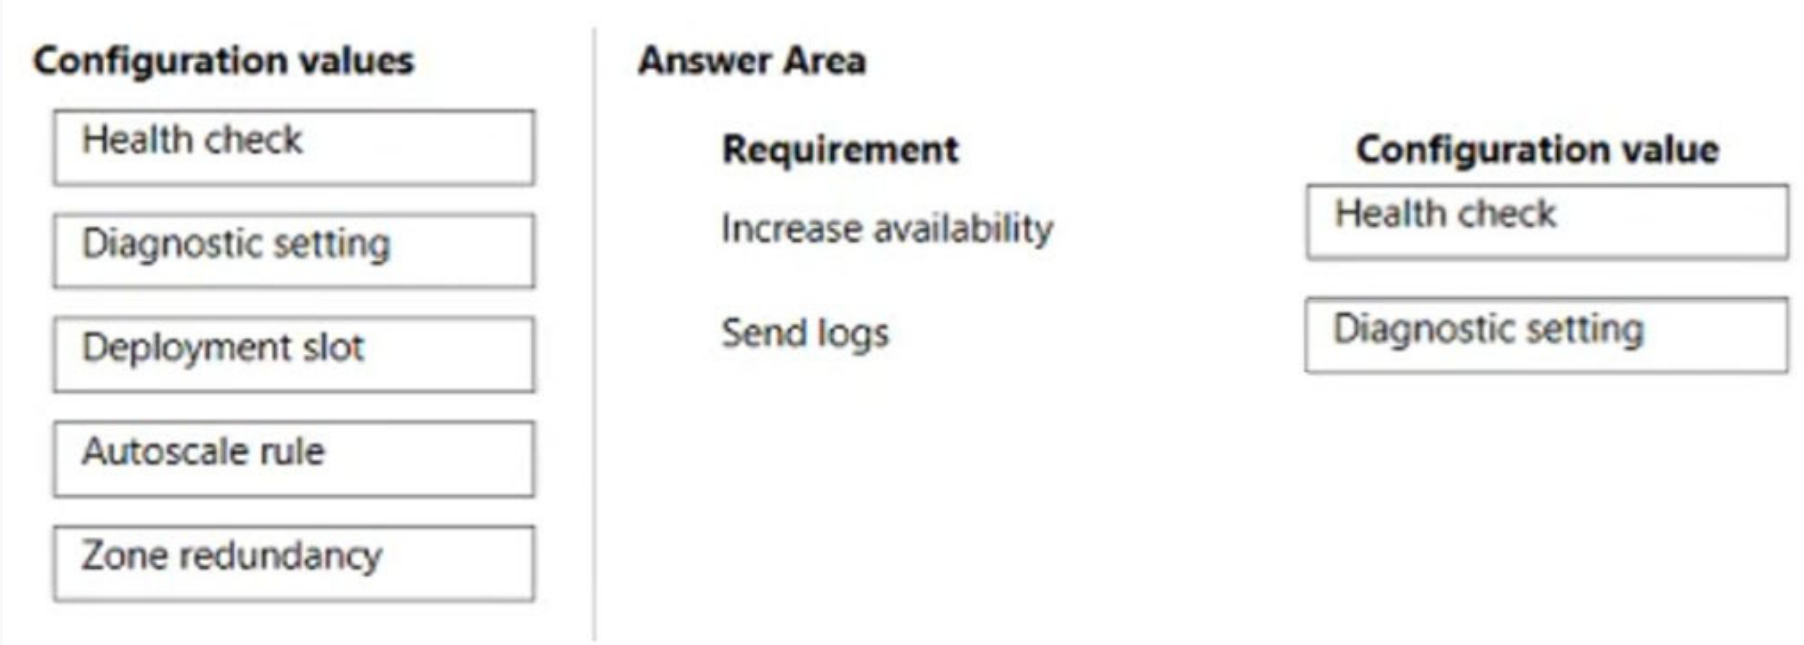

Box 1: Health check -

Health check increases your application's availability by re-routing requests away from unhealthy instances, and replacing instances if they remain unhealthy. Your

App Service plan should be scaled to two or more instances to fully utilize Health check.

Box 2: Diagnostic setting -

Azure provides built-in diagnostics to assist with debugging an App Service app.

With the new Azure Monitor integration, you can create Diagnostic Settings to send logs to Storage Accounts, Event Hubs and Log Analytics.

Reference:

https://docs.microsoft.com/en-us/azure/app-service/monitor-instances-health-check https://docs.microsoft.com/en-us/azure/app-service/troubleshoot-diagnostic-logs

OPT_001122 Highly Voted 11 months, 1 week ago

Box 1: Health check

Box 2: Diagnostic setting

upvoted 30 times

le129 Highly Voted 1 year, 1 month ago

disregard my answer please

upvoted 22 times

Ranzzzan Most Recent 9 months ago

https://learn.microsoft.com/en-us/azure/app-service/monitor-instances-health-check?tabs=dotnet#what-happens-if-my-app-is-running-on-a-single-instance

upvoted 2 times

Billabongs 1 year ago

The answer looks correct,
https://learn.microsoft.com/en-us/azure/app-service/monitor-instances-health-check?tabs=dotnet#what-happens-if-my-app-is-running-on-a-single-instance

upvoted 7 times

Dani_ac7 1 year, 1 month ago

looks ok

upvoted 4 times

le129 1 year, 1 month ago

1. Autoscale rule
2. Diagnostic setting

upvoted 21 times

surprise0011 6 months, 3 weeks ago

There is specified that you have configured four instances, it seems like it is an requirement , so autoscale doesn't fit

upvoted 5 times

Question #32 *Topic 6*

You develop a solution that uses Azure Virtual Machines (VMs).
The VMs contain code that must access resources in an Azure resource group. You grant the VM access to the resource group in Resource Manager.
You need to obtain an access token that uses the VM's system-assigned managed identity.
Which two actions should you perform? Each correct answer presents part of the solution.

A. From the code on the VM, call Azure Resource Manager using an access token.

B. Use PowerShell on a remote machine to make a request to the local managed identity for Azure resources endpoint.

C. Use PowerShell on the VM to make a request to the local managed identity for Azure resources endpoint.

D. From the code on the VM, call Azure Resource Manager using a SAS token.

E. From the code on the VM, generate a user delegation SAS token.

Correct Answer: *BD*

Community vote distribution
AC (100%)

le129 Highly Voted 1 year, 1 month ago

Selected Answer: AC

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm

upvoted 22 times

dddddd111 Most Recent 1 week, 6 days ago

Selected Answer: AC

A and C are correct answers.

upvoted 1 times

aragones 5 months, 2 weeks ago

Got this 2023-05-12.

my cases:
case: You need to configure authorization.
case: You need to ensure the app does not time out and processes the blob data
case study: VanArsdel Inc Canada

upvoted 2 times

adilkhan 6 months, 2 weeks ago

A, C chatGPT

upvoted 1 times

red0400 7 months, 2 weeks ago

Selected Answer: AC

Two parts of the question - Code, Powershell, for code version, A is definitely the choice, and for PowerShell, since it's for virtual machine, we will do it for virtual machine, so it's C.

upvoted 2 times

adilkhan 7 months, 4 weeks ago

AC chatGPT

upvoted 1 times

Bear_Polar 11 months ago

Selected Answer: AC

a and c

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: AC

a and c

upvoted 2 times

POOOJAAAAAAAAAA 10 months, 1 week ago

IMO .. option B
Now that you've created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.

Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.
Reference: https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm

upvoted 1 times

POOOJAAAAAAAAAA 10 months, 1 week ago

My bad its A and C

upvoted 1 times

lombster 1 year ago

Selected Answer: AC

From le129 link
1. Call local managed identity: C
$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}
2. Extract token and call Azure Resource Manager using the access token: A
(Invoke-WebRequest -Uri https://management.azure.com/subscriptions/<SUBSCRIPTION ID>/resourceGroups/<RESOURCE GROUP>?api-version=2016-06-01 -Method GET -ContentType "application/json" -Headers @{ Authorization ="Bearer $ArmToken"}).content

upvoted 4 times

VACAS 1 year, 1 month ago

Selected Answer: AC

https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm

upvoted 4 times

Question #33 *Topic 6*

You are developing a road tollway tracking application that sends tracking events by using Azure Event Hubs using premium tier.
Each road must have a throttling policy uniquely assigned.
You need to configure the event hub to allow for per-road throttling.
What should you do?

A. Use a unique consumer group for each road.

B. Ensure each road stores events in a different partition.

C. Ensure each road has a unique connection string.

D. Use a unique application group for each road.

Correct Answer: *B*

Community vote distribution

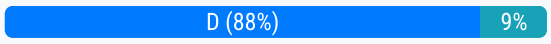

le129 Highly Voted 1 year, 1 month ago

Selected Answer: D

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features.

Application groups
An application group is a collection of client applications that connect to an Event Hubs namespace sharing a unique identifying condition such as the security context - shared access policy or Azure Active Directory (Azure AD) application ID.

Azure Event Hubs enables you to define resource access policies such as throttling policies for a given application group and controls event streaming (publishing or consuming) between client applications and Event Hubs.

For more information, see Resource governance for client applications with application groups.

upvoted 18 times

gmishra88 Highly Voted 1 year ago

Selected Answer: D

Definitely D. Never heard of before and wondering how many things to remember

upvoted 6 times

dddddd111 Most Recent 1 week, 6 days ago

Selected Answer: D

D?????

upvoted 1 times

TusharChandak 1 month, 3 weeks ago

Selected Answer: A

A is the correct option

upvoted 1 times

Dats1987 6 months, 2 weeks ago

Selected Answer: D

Correct ans: D

upvoted 1 times

Sulzirsha 8 months, 3 weeks ago

Looks like option D

Refer: https://learn.microsoft.com/en-us/azure/event-hubs/resource-governance-overview

Azure Event Hubs enables you to govern event streaming workloads of client applications that connect to Event Hubs. You can create logical groups known as application groups where each group is a collection of client applications, and then apply quota and access management policies for an application group (group of client applications).
Application groups are available only in premium and dedicated tiers
Each application group can contain zero or more policies that control the data plane access of the client applications that are part of the application group. Application groups currently support throttling policies.

upvoted 1 times

Osip21 9 months ago

Selected Answer: A

ChatGPT: A. Use a unique consumer group for each road.

In order to configure the Event Hub to allow for per-road throttling, you can use a unique consumer group for each road. This way, you can assign a unique throttling policy to each consumer group, which in turn controls the number of events that can be sent and received by the corresponding road.

Option B and C are not correct solutions to this problem, because partitioning and unique connection strings are not related to throttling policies.

Option D is not correct solution too, application groups are not related to throttling policies and they are used to group related Event Hubs together.

upvoted 2 times

MoroMind 6 months, 1 week ago

Seriously? Do we, software developers, want to use chatGPT as the oracle? Just familiarize yourself with the basic working principles of chatGPT to know that this is not a good idea in case of this kind of test.

I recommend asking chatGPT why D is the correct answer, and you'll get a "credible" explanation too ;).

upvoted 2 times

adilkhan 7 months, 4 weeks ago

In many questions ChatGPT is wrong. For this one as well D is a correct answer!

upvoted 1 times

adilkhan 7 months, 4 weeks ago

Application groups
An application group is a collection of client applications that connect to an Event Hubs namespace sharing a unique identifying condition such as the security context - shared access policy or Azure Active Directory (Azure AD) application ID.

Azure Event Hubs enables you to define resource access policies such as throttling policies for a given application group and controls event streaming (publishing or consuming) between client applications and Event Hubs.

upvoted 2 times

Saurabh0074321 10 months ago

Selected Answer: B

B is the correct answer

upvoted 1 times

David_123 10 months ago

Can you share the resource you used to claim this answer ?

upvoted 2 times

Nhiendo 10 months ago

Selected Answer: B

upvoted 2 times

finnishr 1 year, 1 month ago

Selected Answer: D

D is correct.

upvoted 4 times

Junius 1 year, 1 month ago

Answer is correct

upvoted 2 times

Question #34 *Topic 6*

You develop and deploy an ASP.NET Core application that connects to an Azure Database for MySQL instance.
Connections to the database appear to drop intermittently and the application code does not handle the connection failure.
You need to handle the transient connection errors in code by implementing retries.
What are three possible ways to achieve this goal? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Close the database connection and immediately report an error.

B. Disable connection pooling and configure a second Azure Database for MySQL instance.

C. Wait five seconds before repeating the connection attempt to the database.

D. Set a maximum number of connection attempts to 10 and report an error on subsequent connections.

E. Increase connection repeat attempts exponentially up to 120 seconds.

Correct Answer: *ACD*

Community vote distribution

CDE (100%)

Amy_ant Highly Voted 1 year, 1 month ago

C,D,E

The first and second case are fairly straight forward to handle. Try to open the connection again. When you succeed, the transient error has been mitigated by the system. You can use your Azure Database for MySQL again. We recommend having waits before retrying the connection. Back off if the initial retries fail. This way the system can use all resources available to overcome the error situation. A good pattern to follow is:

Wait for 5 seconds before your first retry.
For each following retry, the increase the wait exponentially, up to 60 seconds.
Set a max number of retries at which point your application considers the operation failed.

https://docs.microsoft.com/en-us/azure/mysql/single-server/concepts-connectivity

upvoted 12 times

deborahQ 1 year ago

but in E, it says increase the attempts exponentially to 120 seconds (but it might be typo). Agree if that is wait exponentially.

upvoted 1 times

RochaG2 Highly Voted 1 year, 1 month ago

Selected Answer: CDE

CDE for sure

upvoted 7 times

dddddd111 Most Recent 1 week, 6 days ago

Selected Answer: CDE

Option A ("Close the database connection and immediately report an error") is not a suitable retry strategy as it does not involve retrying the connection after a transient failure.

upvoted 1 times

JH81 4 months, 1 week ago

Selected Answer: CDE

C, D, E
It says use retries so A and B Are not correct.

upvoted 2 times

AQMA 7 months, 2 weeks ago

Selected Answer: CDE

A transient error is an error that will resolve itself. Transient errors can occur for example when hardware or network failure happens.
A best practice for designing and developing applications in the cloud is to expect transient errors. Assume they can happen in any component at any time and have the appropriate logic to handle these situations.
Transient errors should be handled using retry logic.
A good pattern to follow is:

Wait for 5 seconds before your first retry.
For each following retry, the increase the wait exponentially, up to 60 seconds.
Set a max number of retries at which point your application considers the operation failed.

Reference: https://learn.microsoft.com/en-us/azure/mysql/single-server/concepts-connectivity#handling-transient-errors

upvoted 1 times

NombreFalso 8 months, 1 week ago

Selected Answer: CDE

C, D, E because A and B are just "give up"

upvoted 2 times

adilkhan 9 months ago

C,D,E

upvoted 3 times

OPT_001122 10 months, 3 weeks ago

Selected Answer: CDE

C, D, E

upvoted 4 times

at_sayali 11 months, 3 weeks ago

Got this in 10/2022
Went with C, D, E

upvoted 5 times

Nokaido 1 year, 1 month ago

Selected Answer: CDE

C, D, E are the only ones implementing retries

upvoted 7 times

finnishr 1 year, 1 month ago

Correct answers are C, D, E

upvoted 6 times

Junius 1 year, 1 month ago

Given answers are correct

upvoted 4 times

Question #35 *Topic 6*

You are building a B2B web application that uses Azure B2B collaboration for authentication. Paying customers authenticate to Azure B2B using federation.
The application allows users to sign up for trial accounts using any email address.
When a user converts to a paying customer, the data associated with the trial should be kept, but the user must authenticate using federation.
You need to update the user in Azure Active Directory (Azure AD) when they convert to a paying customer.
Which Graph API parameter is used to change authentication from one-time passcodes to federation?

A. resetRedemption

B. Status

C. userFlowType

D. invitedUser

Correct Answer: *B*

Community vote distribution

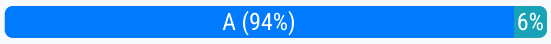

finnishr Highly Voted 1 year, 1 month ago

Selected Answer: A

"When a user redeems a one-time passcode and later obtains an MSA, Azure AD account, or other federated account, they'll continue to be authenticated using a one-time passcode. If you want to update the user's authentication method, you can reset their redemption status."

https://docs.microsoft.com/en-us/azure/active-directory/external-identities/one-time-passcode#user-experience-for-one-time-passcode-guest-users

upvoted 11 times

warchoon 7 months, 3 weeks ago

invitedUser is also used :)
https://learn.microsoft.com/en-us/azure/active-directory/external-identities/reset-redemption-status#:~:text=reset%3E%3E%22%20%20%0A%7D%2C%20%0A%22-,resetRedemption,-%22%3A%20true
Who asks such questions?

upvoted 1 times

damianadalid 5 months, 1 week ago

It is involved in the process, but the fact that invitedUser is an object instead of a boolean, makes me stick with resetRedemption regardless.

upvoted 1 times

Dhanishetty Highly Voted 11 months, 2 weeks ago

Selected Answer: A

https://learn.microsoft.com/en-us/azure/active-directory/external-identities/reset-redemption-status

upvoted 6 times

dddddd111 Most Recent 1 week, 6 days ago

Selected Answer: A

A and B are applicable, but A is more relevant to the question.
Resetting redemption effectively allows a user to be re-invited, and this process can be used to change the authentication method for the user, such as from one-time passcodes to federation. The resetRedemption property is used to reset the redemption status for a user, allowing them to go through the invitation process again, potentially changing the authentication method in the process.

upvoted 1 times

applepie 3 months ago

got this question today, answer A - 7/30/2023, score 895/1000

upvoted 3 times

red0400 8 months ago

The correct answer is C.

The Graph API parameter used to change authentication from one-time passcodes to federation is userFlowType. This parameter is used to update the user flow type for a user in Azure AD B2B collaboration.

Option A, resetRedemption, is a parameter used to reset the invitation redemption status of a user in Azure AD B2B collaboration.

Option B, Status, is a parameter used to update the status of a user in Azure AD, such as activating or deactivating a user account.

Option D, invitedUser, is a parameter used to invite a user to an Azure AD tenant.

Therefore, the correct answer is C, userFlowType. This parameter should be used to update the user's authentication method from one-time passcodes to federation when they convert to a paying customer in the B2B web application.

upvoted 2 times

adilkhan 9 months ago

A is correct . When a user redeems a one-time passcode and later obtains an MSA, Azure AD account, or other federated account, they'll continue to be authenticated using a one-time passcode. If you want to update the user's authentication method, you can reset their redemption status.

https://learn.microsoft.com/en-us/azure/active-directory/external-identities/one-time-passcode

upvoted 2 times

Osip21 9 months ago

Selected Answer: C

C. userFlowType

upvoted 2 times

OPT_001122 10 months, 3 weeks ago

Selected Answer: A

A. resetRedemption Most Voted

upvoted 3 times

LGWJ12 11 months, 2 weeks ago

Selected Answer: A

https://learn.microsoft.com/en-us/azure/active-directory/external-identities/reset-redemption-status

upvoted 4 times

coffecold 1 year ago

Selected Answer: A

https://learn.microsoft.com/en-us/azure/active-directory/external-identities/reset-redemption-status

upvoted 3 times

willchenxa 1 year, 1 month ago

Selected Answer: A

I would select A
https://docs.microsoft.com/en-us/azure/active-directory/external-identities/one-time-passcode#user-experience-for-one-time-passcode-guest-users

upvoted 3 times

le129 1 year, 1 month ago

Selected Answer: A

https://docs.microsoft.com/en-us/azure/active-directory/external-identities/redemption-experience

upvoted 3 times

Question #36 *Topic 6*

HOTSPOT
-

You develop an image upload service that is exposed using Azure API Management. Images are analyzed after upload for automatic tagging.

Images over 500 KB are processed by a different backend that offers a lower tier of service that costs less money. The lower tier of service is denoted by a header named x-large-request. Images over 500 KB must never be processed by backends for smaller images and must always be charged the lower price.

You need to implement API Management policies to ensure that images are processed correctly.
How should you complete the API Management inbound policy? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
<inbound>
  <base/>
  <set-variable name="imageSize" value="@(context.Request.Headers["Content-Length="][0])"/>
  <choose>
    <when condition="@(int.Parse(context.Variables.GetValueOrDefault<string>("imageSize"))<512000)">
      <set-headername="x-large-request"exists-action="[ ]">
        <value>true</value>
      </set-header>
    </when>
    <otherwise>
      <[ ] [ ] ="{{large-image-host}}"/>

    </otherwise>
  </choose>
</inbound>
```

exists-action: skip, append, delete, override

<[ ] [ ]: set-body, forward-request, set-backend-service, set-query-parameter | base-url, dimension, vary-by-header, publish-to-dapr

Correct Answer:

Answer Area

```
<inbound>
  <base/>
  <set-variable name="imageSize" value="@(context.Request.Headers["Content-Length="][0])"/>
  <choose>
    <when condition="@(int.Parse(context.Variables.GetValueOrDefault<string>("imageSize"))<512000)">
      <set-headername="x-large-request"exists-action="[delete]">
        <value>true</value>
      </set-header>
    </when>
    <otherwise>
      <[set-backend-service] [base-url] ="{{large-image-host}}"/>

    </otherwise>
  </choose>
</inbound>
```

exists-action: skip, append, **delete**, override

<[ ] [ ]: **set-backend-service**, set-query-parameter | **base-url**, dimension, vary-by-header, publish-to-dapr

kvtborad Highly Voted 2 months, 3 weeks ago

I got this question on 6th August 2023. chose given answers. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.

upvoted 8 times

Firo Most Recent 3 months ago

The given answer looks correct
the condition here is for "Images < 500KB" so based on the below reference we need to delete the header.

https://learn.microsoft.com/en-us/azure/api-management/set-header-policy#attributes
This is also the reference for back end service
https://learn.microsoft.com/en-us/azure/api-management/set-backend-service-policy

upvoted 2 times

VarunDashora26 2 months, 1 week ago

In my opinion, why would there be a need to Delete the header when " The lower tier of service is denoted by a header named x-large-request" .... instead it should be override, set-backend, baseURI

upvoted 1 times

Heighte 6 months, 3 weeks ago

question seem broken, it doesn't make any sense to me to have delete and value true below.

upvoted 1 times

damianadalid 5 months, 1 week ago

Indeed. And if the answer were override (to set it to true), we would need a complementary set-header for x-large-request to false in the otherwise block...

upvoted 1 times

adilkhan 7 months, 4 weeks ago

<when condition="@((int)context.Request.Headers.GetValueOrDefault("Content-Length", "0") > 500000)">
<set-header name="x-large-request" exists-action="override">
<value>true</value>
</set-header>
<set-backend-service base-url="https://backend-large-service.com" />
</when>
<otherwise>
<set-header name="x-large-request" exists-action="override">
<value>false</value>
</set-header>
<set-backend-service base-url="https://backend-small-service.com" />
</otherwise> chat gpt

upvoted 3 times

JustHereToLearn 9 months ago

The given answer looks correct
https://learn.microsoft.com/en-us/azure/api-management/set-header-policy#attributes
https://learn.microsoft.com/en-us/azure/api-management/set-backend-service-policy

upvoted 2 times

IvanIco 4 months ago

it's definitely not delete

upvoted 1 times

daothanhyen 8 months, 2 weeks ago

Size < 512 should process.
Why is it delete?

upvoted 3 times

NK203 8 months, 2 weeks ago

x-large-request should be set on Size>=512.if Size<512kb,remove this header.

upvoted 7 times

Question #37 *Topic 6*

HOTSPOT
-

You develop several Azure Functions app functions to process JSON documents from a third-party system. The third-party system publishes events to Azure Event Grid to include hundreds of event types, such as billing, inventory, and shipping updates.

Events must be sent to a single endpoint for the Azure Functions app to process. The events must be filtered by event type before processing. You must have authorization and authentication control to partition your tenants to receive the event data.

You need to configure Azure Event Grid.

Which configuration should you use? To answer, select the appropriate values in the answer area.

NOTE: Each correct selection is worth one point.

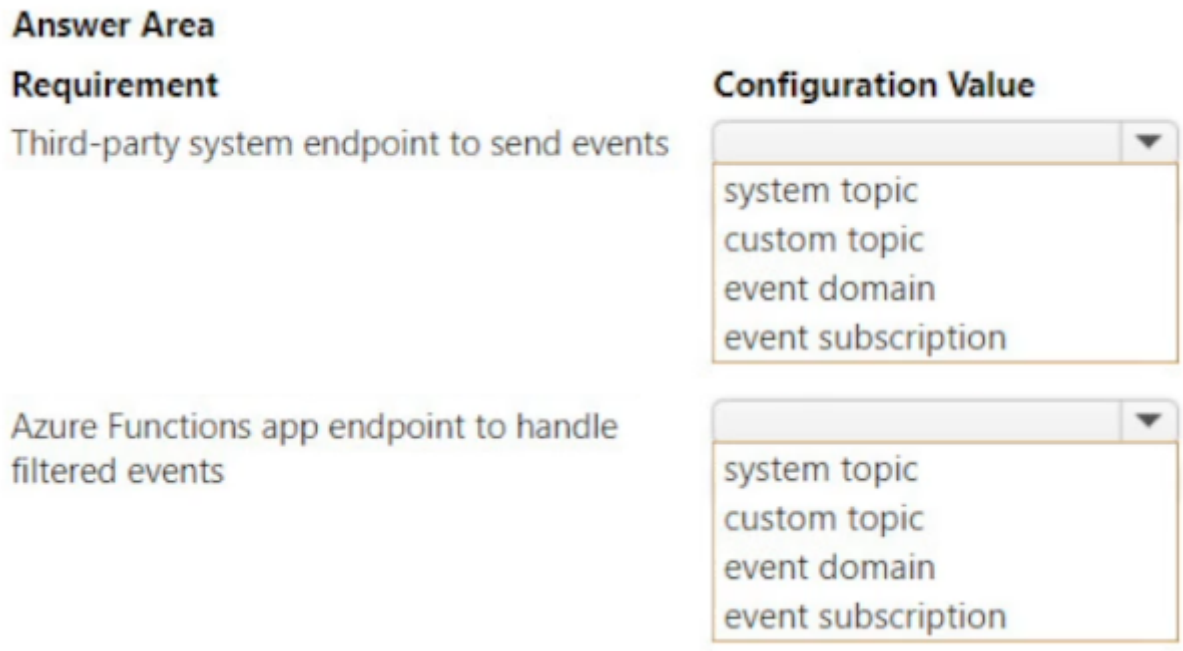

Correct Answer:

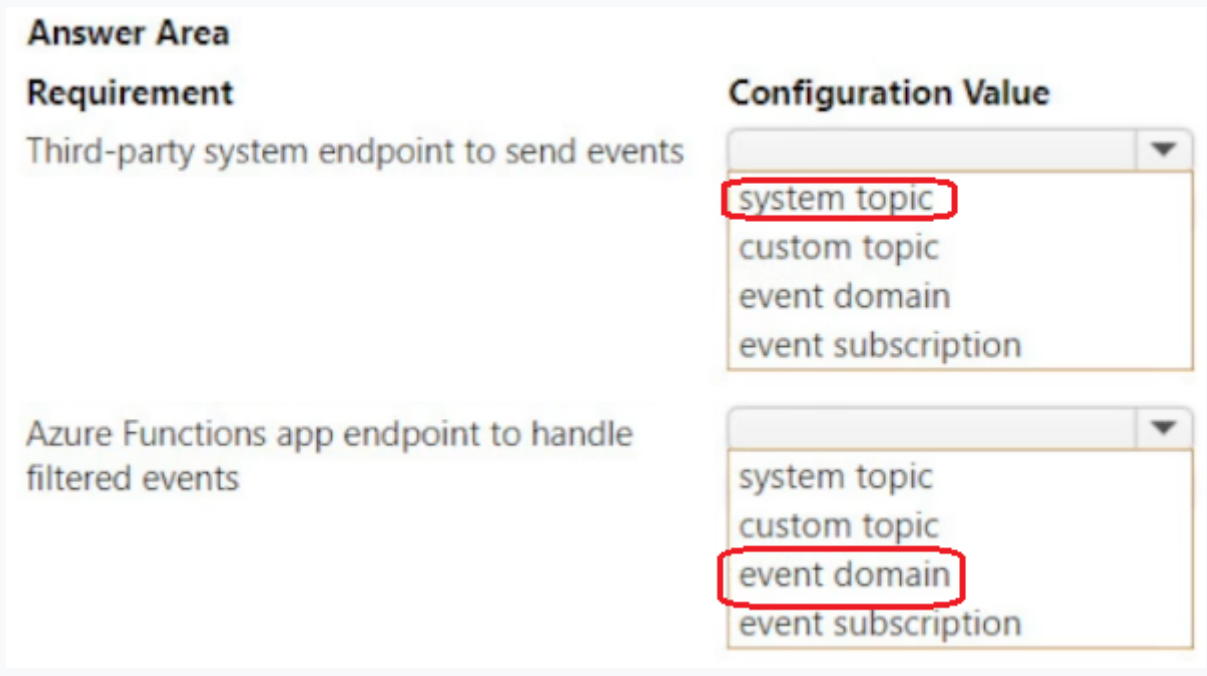

MatejB Highly Voted 7 months, 1 week ago

Check this from official trainig material:
https://learn.microsoft.com/en-us/training/modules/azure-event-grid/2-event-grid-overview

1. custom topic
Because according to link:
"Topics. The event grid topic provides an endpoint where the source sends events(...)Custom topics are application and third-party topics"

2. event subscription
Because according to link:
"Event subscriptions. A subscription tells Event Grid which events on a topic you're interested in receiving. When creating the subscription, you provide an endpoint for handling the event"

upvoted 21 times

Chris2349 4 months, 2 weeks ago

Received the same question today on the exam, 15th of June 2023. Scored 887.
Answered the same.

upvoted 2 times

mmdex Highly Voted 8 months, 1 week ago

I don't know, but from the description of the event domain, I'd say the correct answer is:

1. event domain
2. event subscription

https://learn.microsoft.com/en-us/azure/event-grid/event-domains

1. event domain: Using event domain, the third party system does not send data directly to topics. Event domain provides a single endpoint for the system to publish events.

2. event subscription: Azure functions are event handlers. They use subscriptions. It is true that the event domain also provides a single subscription endpoint, but the handler would then receive all events, not just a filtered set as required:
"Event domains also allow for domain-scope subscriptions. An event subscription on an event domain will receive all events sent to the domain regardless of the topic the events are sent to."
https://learn.microsoft.com/en-us/azure/event-grid/event-domains#domain-scope-subscriptions

upvoted 7 times

warchoon 7 months, 3 weeks ago

So. It must be
1. event domain
2. event domain
You can see the endpoint is the same
https://learn.microsoft.com/en-us/azure/event-grid/includes/media/event-grid-domain-example-use-case/contoso-construction-example.png

upvoted 1 times

warchoon 7 months, 3 weeks ago

https://learn.microsoft.com/en-us/azure/event-grid/how-to-event-domains?tabs=azurecli#create-topics-and-subscriptions:~:text=topics/demotopic1%22%20%5C-,%2D%2Dendpoint,-https%3A//contoso.azurewebsites

upvoted 1 times

AndySmith Most Recent 4 days, 13 hours ago

According to below link / Figure:
1) 3rd party publishes to Domain EP of Event Grid service
2) Azure Functions and others subscribe to Event Subscription.
https://learn.microsoft.com/en-us/azure/event-grid/event-domains?tabs=event-grid-event-schema#example-use-case

The excerpt from that page, which supports the Q's requirements (for event domain):
"Use event domains to:
1) Manage multitenant eventing architectures at scale.
2) Manage your authentication and authorization.
3) Partition your topics without managing each individually.
4) Avoid individually publishing to each of your topic endpoints. "

So, I would go with: 1) event domain 2) event subscription

upvoted 1 times

devex 1 month, 3 weeks ago

1. event domain
2. event subscription

Regarding 1....
Requirement: "You must have authorization and authentication control to partition your tenants to receive the event data."
From: https://learn.microsoft.com/en-us/azure/event-grid/event-domains
"Domains also give you authentication and authorization control over each topic so you can partition your tenants"

upvoted 1 times

kvtborad 2 months, 3 weeks ago

I got this question on 6th August 2023. I was confused so chose a custom topic for the both boxs passed with 904. I got a Case study: city and Lights. All questions are from ExamTopics.

upvoted 3 times

dy0917 5 months, 2 weeks ago

1.Custom topic
2.Event subscriptions

upvoted 1 times

nekkilodeon 6 months, 1 week ago

1.Custom topic
2.Event domain

upvoted 2 times

surprise0011 6 months, 2 weeks ago

i received this one on 2023-04-17, I did not know what to do I selected custom topic twice. I passed with 926. This definitely need more reasarch because it is still there

upvoted 5 times

adilkhan 6 months, 1 week ago

were all the questions from exam topics? please reply

upvoted 2 times

adilkhan 7 months, 3 weeks ago

Custom topic for both ChatGPT.
You should use a custom topic to configure the third-party system to send events to Azure Event Grid. A custom topic allows you to specify the endpoint that the events should be sent to, and it provides you with control over authentication and authorization for the topic.
A custom topic allows you to define your own schema and can be used to filter events based on event type, subject, and other criteria. You can also control access to the custom topic by configuring authorization and authentication.

upvoted 1 times

warchoon 7 months, 3 weeks ago

1. event domain
2. event domain => "You must have authorization and authentication control ..."

upvoted 1 times

red0400 8 months ago

Third party system endpoint to send events= custom topic
Azure functions app endpoint to handle filtered events= event subscription

upvoted 2 times

Asten 8 months ago

custom topic for third-party topics

upvoted 1 times

Asten 8 months ago

the first one should be custom topic

upvoted 1 times

Luza93 9 months, 2 weeks ago

1. I'd say that the first one is custom topic because:
"Custom topics are application and third-party topics" - and we know from the question requirements that there is a third-party app that is sending the events
reference: https://learn.microsoft.com/en-us/azure/event-grid/custom-topics

Additionally :
"Only Azure services can publish events to system topics. Therefore, you don't get an endpoint or access keys that you can use to publish events like you do for custom topics or event domains." - but we need an endpoint for the third-party app
reference: https://learn.microsoft.com/en-us/azure/event-grid/system-topics

2. As ManuKuns wrote the event domain is right because of requirement for "authorization and authentication control ..."
reference: https://learn.microsoft.com/en-us/azure/event-grid/event-domains

upvoted 6 times

ManuKuns 9 months, 3 weeks ago

1. Why system topic? "A system topic in Event Grid represents one or more events published by Azure services..." Is is not said how the third party system publishes the events. https://learn.microsoft.com/en-us/azure/event-grid/system-topics
2. event domain is right: "You must have authorization and authentication control ..." https://learn.microsoft.com/en-us/azure/event-grid/event-domains

upvoted 3 times

Question #38 *Topic 6*

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?

A.
```
az group create
  --name fridge-rg
  --location fridge-loc
```

B.
```
New-AzureRmServiceBusQueue
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Name fridge-q
  -EnablePartitioning $False
```

C.
```
connectionString=$(az servicebus namespace authorization-rule keys list
  --resource-group fridge-rg
  --fridge-ns fridge-ns
  --name RootManageSharedAccessKey
  --query primaryConnectionString --output tsv)
```

D.
```
az servicebus namespace create
  --resource-group fridge-rg
  --name fridge-ns
  --location fridge-loc
```

Correct Answer: *B*

Community vote distribution

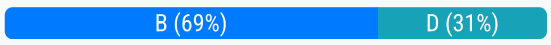

Net_IT Highly Voted 8 months, 2 weeks ago

Selected Answer: B

The next step in the process should be to create a queue IMO (as discussed in similar questions earlier). How the commans looks depends on what type of command prompt you are using (Azure CLI, PowerShell, Bash, ...).

upvoted 8 times

Shayman204 Most Recent 2 weeks, 4 days ago

there is a keyword to note in the question ..'Central location' I will go with B since it has enablePartitioning = false

upvoted 1 times

davidkerr7 4 months, 2 weeks ago

Selected Answer: B

We have been over this before, its B
Bus already exist and look at the location for D makes no sense

upvoted 3 times

jga_private 5 months ago

Selected Answer: D

NewAzureRmServiceBusQueue looks obsolete and replaced by New-AzServiceBusQueue, it needs first a Namespace

upvoted 2 times

JustHereToLearn 9 months, 1 week ago

Selected Answer: D

The answer should be D based on solution provided here https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-cli#provision-resources

Also looks like there is no command like New-AzureRmServiceBusQueue

upvoted 3 times

mabdo 8 months, 2 weeks ago

Agreed, B method was deprecated!

upvoted 1 times

ajayasa 9 months ago

there were similar question earlier it is posted as answer B. as service bus already exists we need to complete the configuration which is queue. the command is available :https://learn.microsoft.com/zh-cn/powershell/module/azurerm.servicebus/new-azurermservicebusqueue?view=azurermps-6.13.0&viewFallbackFrom=azurermps-6.11.0

upvoted 8 times

Michael2023 9 months ago

and also the --location value in D looks wrong... so B. would be the only right answer

upvoted 5 times

Question #39 *Topic 6*

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?

A.
```
New-AzureRmServiceBusNamespace
  -ResourceGroupName fridge-rg
  -NamespaceName 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.
```
connectionString=$(az servicebus namespace authorization-rule keys list
  --resource-group fridge-rg
  --fridge-ns fridge-ns
  --name RootManageSharedAccessKey
  --query primaryConnectionString --output tsv)
```

Correct Answer: *C*

Community vote distribution

C (100%)

halfway Highly Voted 6 months, 2 weeks ago

Selected Answer: C

Service Bus Queue

upvoted 5 times

macobuzi Most Recent 2 months ago

Lol, there are so many duplications for this question!

upvoted 2 times

Chris2349 4 months, 3 weeks ago

The answer seems correct although uses an outdated command:
https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-powershell

upvoted 1 times

Question #40 *Topic 6*

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?

A.
```
az group create
  --name fridge-rg
  --location fridge-loc
```

B.
```
New-AzureRmServiceBusNamespace
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Location fridge-loc
```

C.
```
New-AzureRmServiceBusQueue
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Name fridge-q
  -EnablePartitioning $False
```

D.
```
New-AzureRmResourceGroup
-Name fridge-rg
-Location fridge-loc
```

Correct Answer: *C*

damianadalid Highly Voted 5 months, 1 week ago

I found this question three times already, please admins, avoid duplicated (triplicated) questions, I did not pay my contributor access for this.

upvoted 19 times

ENGs Most Recent 3 weeks, 4 days ago

One of these Frige Questions where my exam 2023-10 before the Update of the Exam

upvoted 1 times

paunski7 6 months, 1 week ago

C is the correct as answer is similar to previous same questions

upvoted 3 times

Question #41 *Topic 6*

DRAG DROP

-

You develop and deploy several APIs to Azure API Management.

You create the following policy fragment named APICounts:

```
<fragment>
    <emit-metric value="1" namespace="custom-metrics">
        <dimension name="User ID" />
        <dimension name="Operation ID" />
        <dimension name="API ID" />
        <dimension name="Client IP" value="@(context.Request.IpAddress)" />
    </emit-metric>
</fragment>
```

The policy fragment must be reused across various scopes and APIs. The policy fragment must be applied to all APIs and run when a calling system invokes any API.

You need to implement the policy fragment.

XML elements

- name
- inbound
- outbound
- set-variable
- fragment-id
- include-fragment

Answer Area

```
<policies>
    < XML >
        < XML | XML ="APICounts" />
        <base />
    </ XML >
    . . .
</policies>
```

Correct Answer:

Answer Area

```
<policies>
    < inbound >
        < include-fragment | fragment-id ="APICounts" />
        <base />
    </ inbound >
    . . .
</policies>
```

juanckar Highly Voted 3 months, 3 weeks ago

This was on the exam (July 2023). Went with proposed answer. Scored 917

upvoted 7 times

phantom31 3 months, 3 weeks ago

how many percentage of questions come from examtopics?

upvoted 1 times

rishabh12307 3 weeks, 5 days ago

around 60-70%

upvoted 1 times

calmriver44 Most Recent 5 months ago

Answers are correct. Here is an explanation of Policy fragments

https://learn.microsoft.com/en-us/azure/api-management/policy-fragments

For example, insert the policy fragment named ForwardContext in the inbound policy section:

<policies>
<inbound>
<include-fragment fragment-id="ForwardContext" />
<base />
</inbound>
[...]

upvoted 1 times

e1odie 6 months, 2 weeks ago

provided answer seems correct according to https://learn.microsoft.com/en-us/azure/api-management/include-fragment-policy

upvoted 4 times

Question #42 *Topic 6*

HOTSPOT
-

You are developing a solution that uses several Azure Service Bus queues. You create an Azure Event Grid subscription for the Azure Service Bus namespace. You use Azure Functions as subscribers to process the messages.

You need to emit events to Azure Event Grid from the queues. You must use the principal of least privilege and minimize costs.

Which Azure Service Bus values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

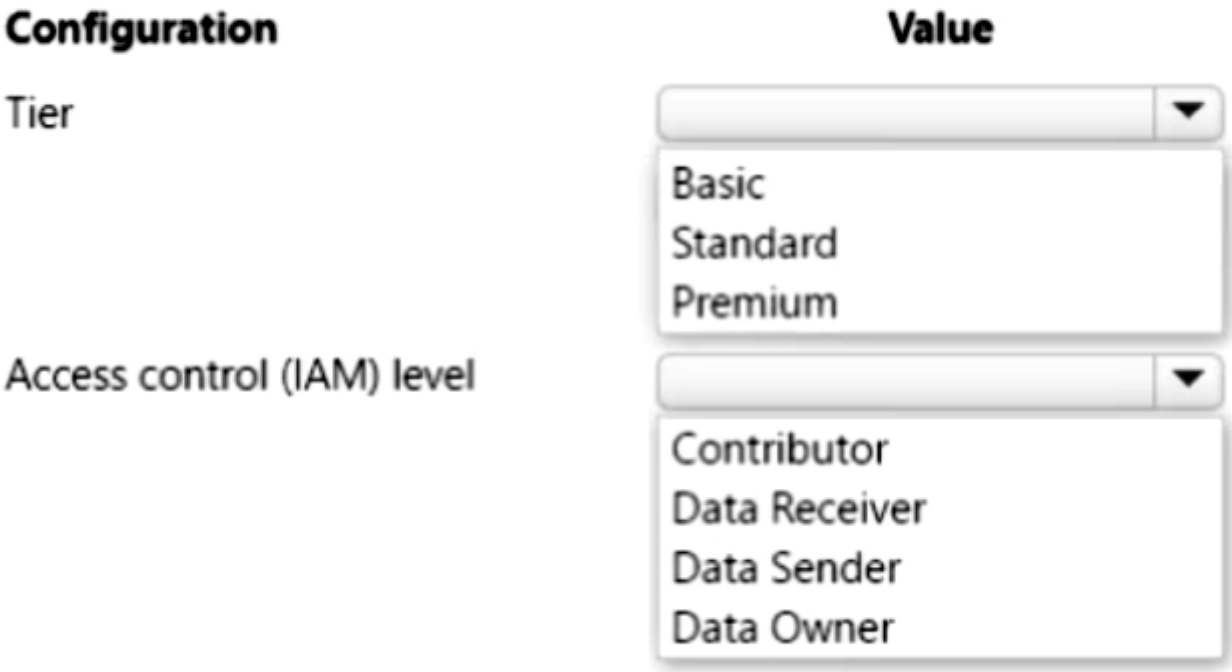

Answer Area

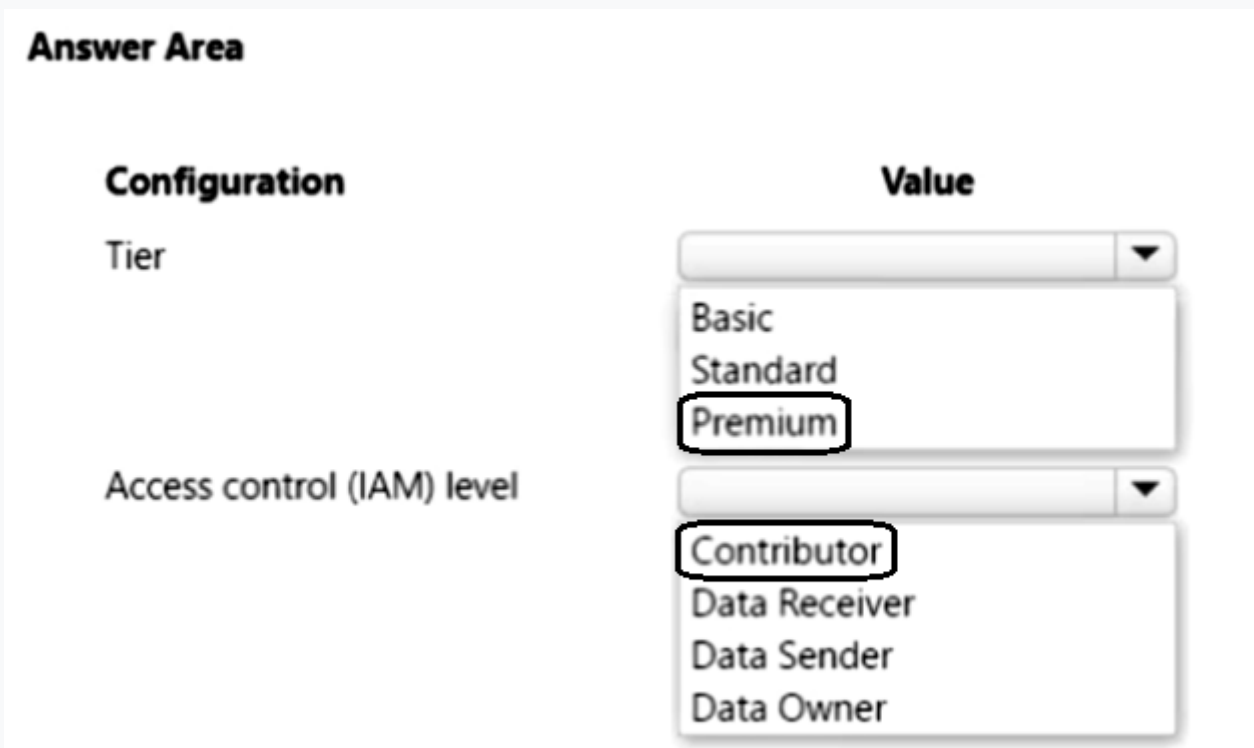

Correct Answer:

adilkhan Highly Voted 6 months, 1 week ago

Remember with Service bus + event grid its always premimum

upvoted 12 times

surprise0011 5 months, 1 week ago

Exactly, this requires Service Bus Premium tier: https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept?tabs=event-grid-event-schema

upvoted 1 times

kvtborad Highly Voted 2 months, 3 weeks ago

I got this question on 6th August 2023. chose given answers. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.

upvoted 5 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 2 times

Chris2349 4 months, 2 weeks ago

Received the question, 15th of June 2023. Scored 887.
Went with the given answers.

upvoted 3 times

SharpZx 6 months, 2 weeks ago

Looks correct:

https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept?tabs=event-grid-event-schema

upvoted 3 times

halfway 6 months, 2 weeks ago

Correct.

Ref: https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept?tabs=event-grid-event-schema

upvoted 2 times

Question #43 *Topic 6*

You are developing several Azure API Management (APIM) hosted APIs.

The APIs have the following requirements:

- Require a subscription key to access all APIs.
- Include terms of use that subscribers must accept to use the APIs.
- Administrators must review and accept or reject subscription attempts.
- Limit the count of multiple simultaneous subscriptions.

You need to implement the APIs.

What should you do?

A. Configure and apply header-based versioning.

B. Create and publish a product.

C. Configure and apply query string-based versioning.

D. Add a new revision to all APIs. Make the revisions current and add a change log entry.

Correct Answer: *C*

Community vote distribution

B (100%)

HoneyMax Highly Voted 3 months, 3 weeks ago

Selected Answer: B

Only product allows to set the terms of use

upvoted 5 times

dddddd111 Most Recent 1 week, 5 days ago

Selected Answer: B

B is the right answer.

upvoted 1 times

kvtborad 2 months, 3 weeks ago

I got this question on 6th August 2023. chose highly voted B. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.

upvoted 4 times

EliteAllen 3 months ago

Selected Answer: B

B. Create and publish a product.

To meet the requirements of needing a subscription key, including terms of use, having administrators review subscriptions, and limiting simultaneous subscriptions for Azure API Management hosted APIs, you should create and publish a product

upvoted 3 times

lyggwtwtczxnhznebw 3 months ago

Selected Answer: B

Voting B

upvoted 2 times

75daltonic 3 months, 2 weeks ago

Selected Answer: B

I think B

upvoted 2 times

AbidooKing 3 months, 3 weeks ago

The given answer is wrong. The correct solution is B "Create and publish a product".

https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-add-products?tabs=azure-portal#create-and-publish-a-product

upvoted 4 times

junix_ 3 months, 3 weeks ago

I believe it's B

refs: https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-add-products?tabs=azure-portal

upvoted 2 times

Question #44 *Topic 6*

HOTSPOT
-

You are developing a solution by using the Azure Event Hubs SDK. You create a standard Azure Event Hub with 16 partitions. You implement eight event processor clients.

You must balance the load dynamically when an event processor client fails. When an event processor client fails, another event processor must continue processing from the exact point at which the failure occurred. All events must be aggregate and upload to an Azure Blob storage account.

You need to implement event processing recovery for the solution.

Which SDK features should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Feature
Ensure that event process clients mark the position within an event sequence.	Offset / Checkpoint / Namespace / Capture
Mark the event processor client position within a partition event sequence.	Offset / Checkpoint / Namespace / Capture

Correct Answer:

Answer Area

Requirement	Feature
Ensure that event process clients mark the position within an event sequence.	**Offset** / Checkpoint / Namespace / Capture
Mark the event processor client position within a partition event sequence.	Offset / **Checkpoint** / Namespace / Capture

AbidooKing Highly Voted 3 months, 3 weeks ago

This is a really stupid and confusing question but I think it's checkpoint to ensure marking and offset to mark...

This ref says: "Checkpointing is a process by which readers mark or commit their position within a partition event sequence."
https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features#checkpointing

But this ref here says: "Consumers also engage in checkpointing. Through this process, subscribers use offsets to mark their position within a partition event sequence. An offset is a placeholder that works like a bookmark to identify the last event that the consumer read."
https://learn.microsoft.com/en-us/azure/architecture/reference-architectures/event-hubs/partitioning-in-event-hubs-and-kafka#dataflow

upvoted 6 times

AbidooKing 3 months, 2 weeks ago

I actually got this on my exam 2023July13!! I chickened out and selected checkpoint on both just to make sure i at least get one them correct. Scored 917.

upvoted 6 times

EliteAllen Most Recent 3 months ago

Checkpoint for both.

The features needed to implement event processing recovery for the solution, and specifically to handle the marking of the event processor client position within a partition event sequence, are related to checkpointing. Checkpointing is the process of tracking the processing of events within a partition and is essential for ensuring that processing can continue from the exact point at which a failure occurred.

1. Ensure that event process clients mark the position within an event sequence: Again, use Checkpointing. By regularly recording checkpoints, the event processor clients can ensure that they mark their progress through the event stream, allowing for recovery from failures.

2. Mark the event processor client position within a partition event sequence: Use Checkpointing. This allows an event processor to record its position in the event stream and resume from that position if it's restarted.

upvoted 1 times

lyggwtwtczxnhznebw 3 months ago

An offset is the position of an event within a partition.
Checkpointing is a process by which readers mark or commit their position within a partition event sequence.
Can't make my mind, I'll say select checkpoint on both and call it a day.

upvoted 2 times

junix_ 3 months, 3 weeks ago

I believe the given answer is correct.

refs: https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features

upvoted 2 times

RenatoJr 3 months, 3 weeks ago

Should it be both checkpoint?

Checkpoints allow an event processor client to keep track of its progress within the event stream by storing the offset or sequence number of the last event it successfully processed.

upvoted 2 times

Question #45 *Topic 6*

HOTSPOT
-

You are developing a new API to be hosted by Azure API Management (APIM). The backend service that implements the API has not been completed. You are creating a test API and operation.

You must enable developers to continue with the implementation and testing of the APIM instance integrations while you complete the backend API development.

You need to configure a test API response.

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

Answer Area

APIM Configuration Setting	APIM Configuration Value
Policy	proxy set-status mock-response forward-request
Policy section	inbound backend on-error outbound
HTTP response code	200 400 500 501

Correct Answer:

Answer Area

APIM Configuration Setting	APIM Configuration Value
Policy	proxy set-status **mock-response** forward-request
Policy section	inbound **backend** on-error outbound
HTTP response code	**200** 400 500 501

AbidooKing Highly Voted 3 months, 3 weeks ago

Answer is wrong. It should be inbound not backend.
Correct is: Mock-response, inbound and 200

https://learn.microsoft.com/en-us/azure/api-management/mock-api-responses?tabs=azure-portal#enable-response-mocking

upvoted 12 times

junix_ Highly Voted 3 months, 3 weeks ago

Answer should be:
1. mock-response
2. inbound
3. 200

refs: https://learn.microsoft.com/en-us/azure/api-management/mock-api-responses?tabs=azure-cli

upvoted 7 times

applepie 3 months ago

got this question today, went with this answer - 7/30/2023, score 895/1000

upvoted 3 times

Question #46 *Topic 6*

You are developing several Azure API Management (APIM) hosted APIs.

You must inspect request processing of the APIs in APIM. Requests to APIM by using a REST client must also be included. The request inspection must include the following information:

- requests APIM sent to the API backend and the response it received
- policies applied to the response before sending back to the caller
- errors that occurred during the processing of the request and the policies applied to the errors
- original request APIM received from the caller and the policies applied to the request

You need to inspect the APIs.

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

NOTE: Each correct selection is worth one point.

A. Enable the Allow tracing setting for the subscription used to inspect the API.

B. Add the Ocp-Apim-Trace header value to the API call whit a value set to true.

C. Add the Ocp-Apim-Subscription-Key header value to the key for a subscription that allows access to the API.

D. Create and configure a custom policy. Apply the policy to the inbound policy section with a global scope.

E. Create and configure a custom policy. Apply the policy to the outbound policy section with an API scope.

Correct Answer: *ACE*

Community vote distribution

ABC (85%) Other

75daltonic Highly Voted 3 months, 1 week ago

Selected Answer: ABC

From: https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-api-inspector
-To trace request processing, you must enable the Allow tracing setting for the subscription used to debug your API.
-When making requests to API Management using curl, a REST client such as Postman, or a client app, enable tracing by adding the following request headers:
Ocp-Apim-Trace - set value to true
Ocp-Apim-Subscription-Key - set value to the key for a tracing-enabled subscription that allows access to the API

upvoted 5 times

uncledana 3 months, 1 week ago

Agree - ABC.

https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-api-inspector

upvoted 1 times

NPE_ Most Recent 4 days, 16 hours ago

Selected Answer: ACE

How about E? E is definitely the only correct answer to fulfill the outbound response requirement in the question description or? So the answer should be ABCE. In this case, the provided answer can be correct

upvoted 1 times

EliteAllen 3 months ago

Selected Answer: ABC

In Azure API Management (APIM), you can inspect the request processing of APIs, including the details mentioned in the question, by performing certain actions.

Here's how you can achieve the requirements:

Enable the Allow tracing setting for the subscription used to inspect the API:
This will allow you to trace the requests and responses, including details of the policies applied and any errors that occurred.

Add the Ocp-Apim-Trace header value to the API call with a value set to true:
This header tells APIM to include trace information in the response. It will include details about the request and response, policies applied, and errors that occurred.

Add the Ocp-Apim-Subscription-Key header value to the key for a subscription that allows access to the API:
This header is used to pass the subscription key, which is required to call an API if the subscription key is mandated.

upvoted 2 times

MikeAWS 1 week, 3 days ago

what is this, a ChatGPT answer?

upvoted 1 times

RenatoJr 3 months, 3 weeks ago

Selected Answer: BC

I believe the answer should be BCD

B. This header enables request and response tracing for the API call.
C. This header ensures that the API call is associated with a valid subscription.
D. Apply the policy to the inbound policy section with a global scope. By applying the custom policy to the inbound section with a global scope, you can inspect the original request received by APIM, apply policies to the request, inspect requests sent to the API backend, and capture responses received from the backend. This allows you to gather the requested information about the request processing, applied policies, and errors.

https://learn.microsoft.com/en-us/azure/api-management/api-management-policies
https://learn.microsoft.com/en-us/azure/api-management/api-management-subscriptions
https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-api-inspector

upvoted 1 times

AbidooKing 3 months, 3 weeks ago

I think it should be ABC.
Ocp-Apim-Trace - set value to true and set value to the Ocp-Apim-Subscription-Key when making requests to API Management using curl or a REST client such as Postman,

And to trace request processing, you must enable the Allow tracing setting for the subscription used to debug your API.

No need to configure a custom policy when tracing Outbound shows the policies applied to the response before sending back to the caller. Right?

https://learn.microsoft.com/en-us/azure/governance/policy/tutorials/create-custom-policy-definition

upvoted 3 times

junix_ 3 months, 3 weeks ago

Selected Answer: ABC

I believe the answer should be A,B,C

refs: https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-api-inspector

upvoted 4 times

75daltonic 3 months, 2 weeks ago

From: https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-api-inspector
-To trace request processing, you must enable the Allow tracing setting for the subscription used to debug your API.
-When making requests to API Management using curl, a REST client such as Postman, or a client app, enable tracing by adding the following request headers:
Ocp-Apim-Trace - set value to true
Ocp-Apim-Subscription-Key - set value to the key for a tracing-enabled subscription that allows access to the API

upvoted 1 times

Question #47 *Topic 6*

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?

A.
```
Get-AzureRmServiceBusKey
  -ResourceGroupName fridge-rg
  -Namespace fridge-ns
  -Name RootManageSharedAccessKey
```

B.
```
New-AzureRmResourceGroup
  -Name fridge-rg
  -Location fridge-loc
```

C.
```
New-AzureRmServiceBusNamespace
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Location fridge-loc
```

D.
```
New-AzureRmServiceBusQueue
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Name fridge-q
  -EnablePartitioning $False
```

Correct Answer: *D*

Community vote distribution

D (100%)

Rodikito Highly Voted 2 months, 1 week ago

this question is repeated several times in this site

upvoted 8 times

nardk Most Recent 1 month, 2 weeks ago

Selected Answer: D

Must be D

upvoted 2 times

Question #48 *Topic 6*

HOTSPOT
-

You plan to implement an Azure Functions app.

The Azure Functions app has the following requirements:

- Must be triggered by a message placed in an Azure Storage queue.
- Must use the queue name set by an app setting named input_queue.
- Must create an Azure Blob Storage named the same as the content of the message.

You need to identify how to reference the queue and blob name in the function.json file of the Azure Functions app.

How should you reference the names? To answer, select the appropriate values in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Reference type	Value
Queue name	input_queue {input_queue} %input_queue%
Blob name	{queueTrigger} {input_queue}/{id} %input_queue%/{filename}

Correct Answer:

Answer Area

Reference type	Value
Queue name	input_queue **{input_queue}** %input_queue%
Blob name	**{queueTrigger}** {input_queue}/{id} %input_queue%/{filename}

Firo Highly Voted 3 months ago

Queue name: %input_queue%
Reference: https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-expressions-patterns#binding-expressions---app-settings

Blob name: {queue Trigger}
Reference: https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-expressions-patterns#trigger-metadata

Blobname

upvoted 15 times

Ayaz_b Most Recent 1 month ago

```
{
"bindings": [
{
"type": "queueTrigger",
"name": "myQueueItem",
"direction": "in",
"queueName": "%input_queue%" // Reference the input_queue app setting
},
{
"type": "blob",
"name": "outputBlob",
"path": "{QueueTrigger}", // Use the value from the queue message as the blob name
"connection": "AzureWebJobsStorage",
"direction": "out",
"containerName": "%output_blobcontainer%" // Reference the output_blobcontainer app setting
}
]
}
```

upvoted 1 times

Question #49 *Topic 6*

HOTSPOT
-

You have an Azure API Management instance named API1 that uses a managed gateway.

You plan to implement a policy that will apply at a product scope and will set the header of inbound requests to include information about the region hosting the gateway of API1. The policy definition contains the following content:

```
<policies>
    <inbound>
        TARGET1
        <set-header name="x-request-context-data" exists-action="override">
            <value>@(TARGET2.Deployment.Region)</value>
      </set-header>
    </inbound>
</policies>
```

You have the following requirements for the policy definition:

• Ensure that the header contains the information about the region hosting the gateway of API1.
• Ensure the policy applies only after any global level policies are processed first.

You need to complete the policy definition.

Which values should you choose? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

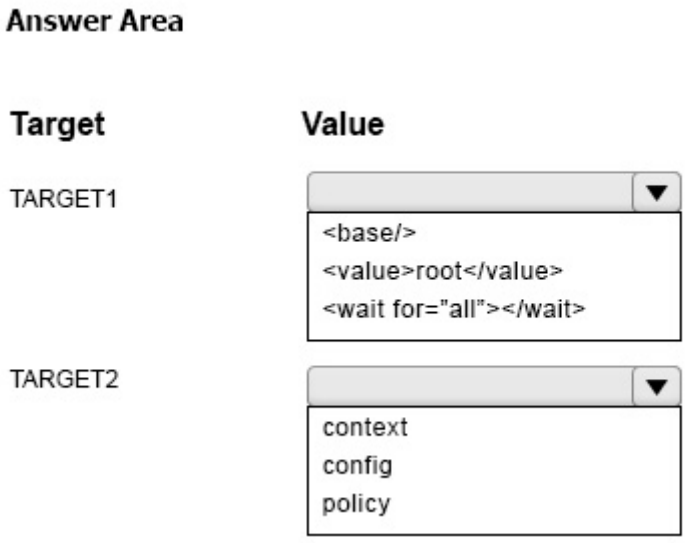

Correct Answer:

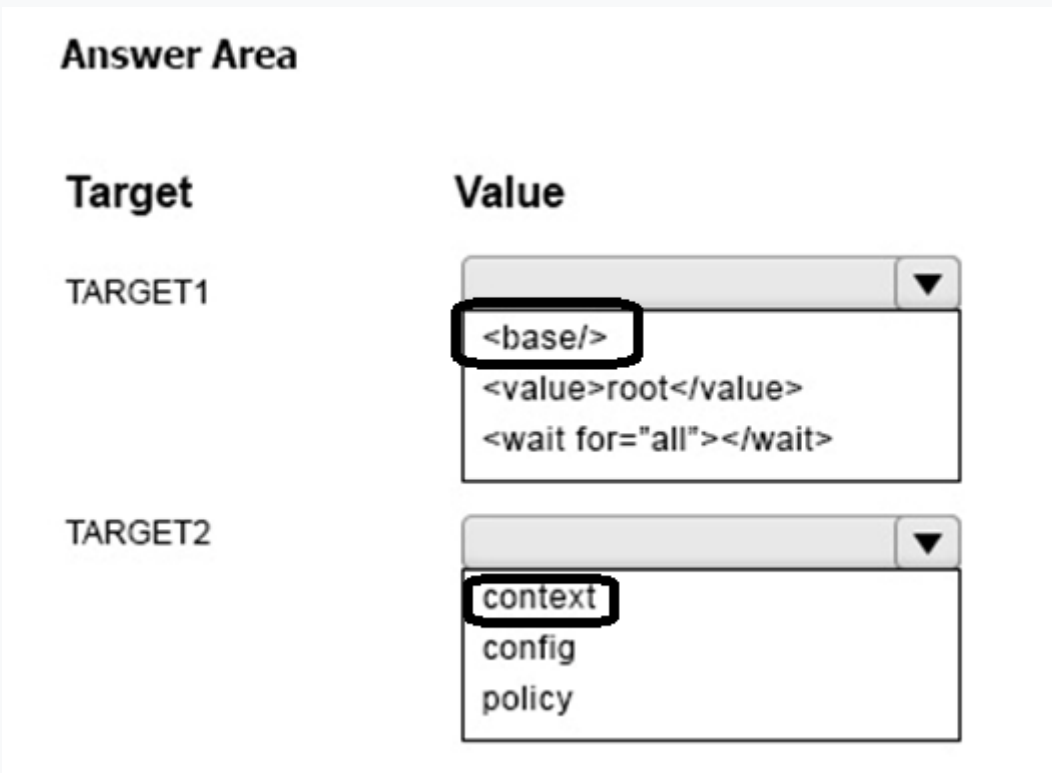

Capital 1 month, 1 week ago

Given answer is correct:
https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-policies#use-policy-expressions-to-modify-requests

upvoted 3 times

oCaio 2 months, 2 weeks ago

Given answer correct.

Ref: https://safe.menlosecurity.com/https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-policies
Article Section: Use policy expressions to modify requests

upvoted 4 times

Firo 3 months ago

<base /> is correct: https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-expressions-patterns#trigger-metadata

upvoted 2 times

Question #50 *Topic 6*

You are developing several Azure API Management (APIM) hosted APIs.

You must make several minor and non-breaking changes to one of the APIs. The API changes include the following requirements:

- Must not disrupt callers of the API.
- Enable roll back if you find issues.
- Documented to enable developers to understand what is new.
- Tested before publishing.

You need to update the API.

What should you do?

A. Configure and apply header-based versioning.

B. Create and publish a product.

C. Configure and apply a custom policy.

D. Add a new revision to the API.

E. Configure and apply query string-based versioning.

Correct Answer: *E*

Community vote distribution

D (100%)

ddddd111 1 week, 5 days ago

Selected Answer: D

D is more relevant.

upvoted 1 times

Ciupaz 1 month, 1 week ago

The best approach for making non-breaking changes to an APIM-hosted API is to create a new revision of the API.

upvoted 1 times

lyggwtwtczxnhznebw 2 months, 1 week ago

Selected Answer: D

Voting D

upvoted 2 times

kabbas 2 months, 3 weeks ago

Should be revision and not versioning. Anew version will disrupt current use

upvoted 2 times

Firo 3 months ago

Selected Answer: D

D is correct:
In Azure API Management, use revisions to make non-breaking API changes so you can model and test changes safely. When ready, you can make a revision current and replace your current API.
https://learn.microsoft.com/en-us/azure/api-management/api-management-get-started-revise-api?tabs=azure-portal

upvoted 4 times

Question #51 *Topic 6*

HOTSPOT
-

You are developing an application to store millions of images in Azure blob storage.

The application has the following requirements:

• Store the Exif (exchangeable image file format) data from the image as blob metadata when the application uploads the image.
• Retrieve the Exif data from the image while minimizing bandwidth and processing time.
• Utilizes the REST API.

You need to use the image Exif data as blob metadata in the application.

Which HTTP verbs should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Application Metadata Action	HTTP Verb
Store Exif data.	GET / PUT / POST / HEAD
Retrieve Exif data.	PUT / POST / HEAD / CONNECT

Correct Answer:

Answer Area

Application Metadata Action	HTTP Verb
Store Exif data.	GET / **PUT** / POST / HEAD
Retrieve Exif data.	PUT / POST / **HEAD** / CONNECT

danisancho 4 days, 2 hours ago

According to ChatGPT, first answer is correct: "When storing Exif data from an image as blob metadata in Azure Blob Storage, you can use the HTTP `PUT` method. This method is commonly used for uploading a resource to the specified URI, which in this case would be the URL of your Azure Blob Storage container. When making the `PUT` request, you can include the Exif data as metadata headers in the request to associate the metadata with the uploaded blob."

upvoted 1 times

Ciupaz 1 month, 1 week ago

Strange question and answers. The response should be POST and GET.

upvoted 1 times

Fabiann 2 months, 2 weeks ago

we want to store as Metadata :

so PUT for set and GET/HEAD to retrieve.

https://learn.microsoft.com/en-us/training/modules/work-azure-blob-storage/6-set-retrieve-properties-metadata-rest?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wwl.develop-solutions-that-use-blob-storage

upvoted 3 times

the_odd_one 5 days, 14 hours ago

I see where you come from on the link, however generally POST is for new data and PUT is for overriding existing data.
When the question is phrased "Store the Exif (exchangeable image file format) data from the image as blob metadata when the application uploads the image." I lean towards using POST and not PUT

upvoted 1 times

oCaio 2 months, 2 weeks ago

Retrieve via HEAD because

"For example, if your application needs the Exif (exchangeable image format) data from a photo, it can retrieve the photo and extract it. To save bandwidth and improve performance, your application can store the Exif data in the blob's metadata when the application uploads the photo. You can then retrieve the Exif data in metadata using only a HEAD request."

Ref: https://learn.microsoft.com/en-us/azure/storage/blobs/storage-performance-checklist#use-metadata

upvoted 3 times

kvtborad 2 months, 3 weeks ago

I got this question on 6th August 2023. chose highly voted first box: PUT, Second Box: POST. passed with 904. I got a Case study: city and Lights. All questions are from ExamTopics.

upvoted 3 times

NguyenDo 3 months ago

Store can be PUT or POST, make no sense, I think Head for store and retrieve.

upvoted 3 times

Firo 3 months ago

Store: Head (not sure)
Retrieve: Head
https://learn.microsoft.com/en-us/azure/storage/blobs/storage-performance-checklist#use-metadata

upvoted 1 times

Question #52 *Topic 6*

You are developing several microservices to run on Azure Container Apps for a company. External TCP ingress traffic from the internet has been enabled for the microservices.

The company requires that the microservices must scale based on an Azure Event Hub trigger.

You need to scale the microservices by using a custom scaling rule.

Which two Kubernetes Event-driven Autoscaling (KEDA) trigger fields should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. metadata

B. type

C. authenticationRef

D. name

E. metricType

Correct Answer: *AB*

Community vote distribution

AB (100%)

applepie Highly Voted 3 months ago

https://keda.sh/docs/2.11/scalers/azure-event-hub/

upvoted 7 times

nardk Most Recent 1 month, 2 weeks ago

Selected Answer: AB

A. metadata
B. type

upvoted 2 times

Topic 7 - Testlet 1

Question #1 *Topic 7*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment -

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk Server 2016. The VM runs the following workflows:

Ocean Transport `" This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.

Inland Transport `" This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

Container API `" This API provides container information including weight, contents, and other attributes.

Location API `" This API provides location information regarding shipping ports of call and trucking stops.

Shipping REST API `" This API provides shipping information for use and display on the shipping website.

Shipping Data -

The application uses MongoDB JSON document storage database for all container and transport information.

Shipping Web Site -

The site displays shipping container tracking information and container contents. The site is located at http://shipping.wideworldimporters.com/

Proposed solution -

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:

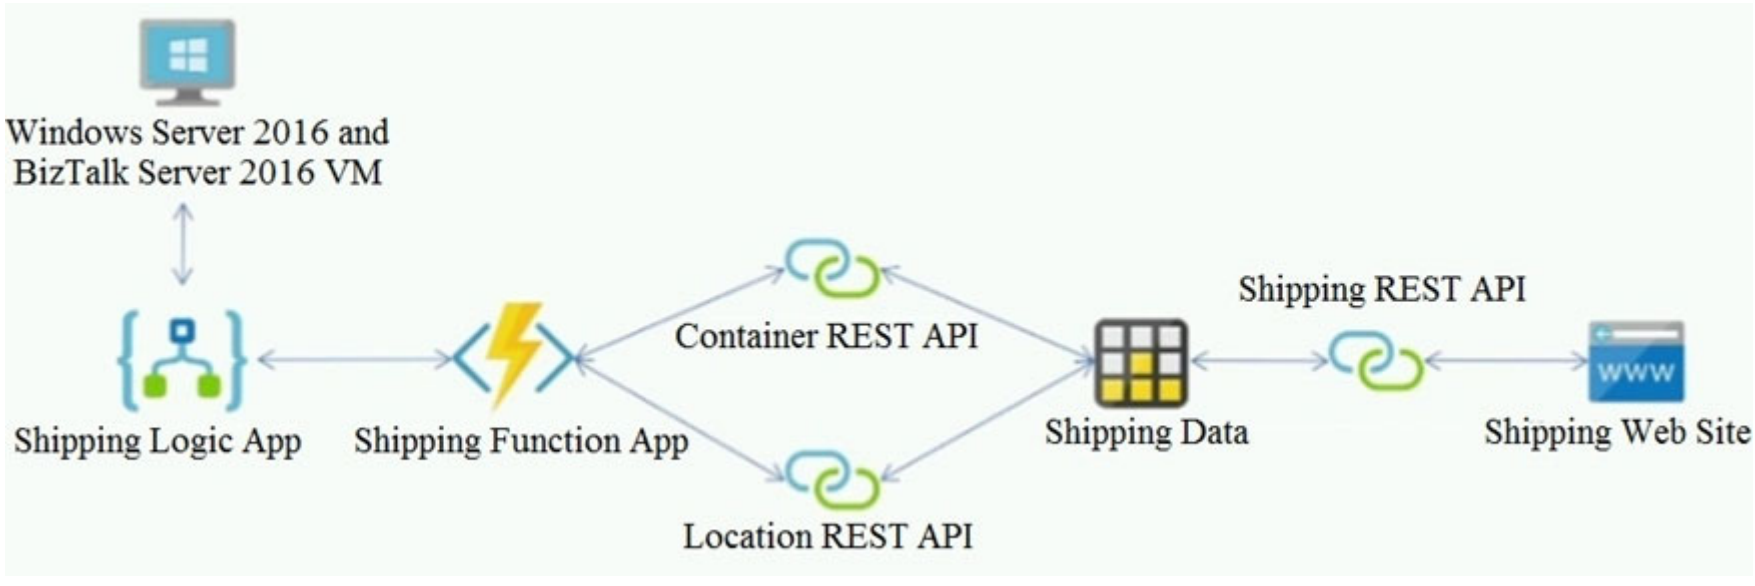

Requirements -

Shipping Logic app -
The Shipping Logic app must meet the following requirements:
Support the ocean transport and inland transport workflows by using a Logic App.
Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app -
Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

REST APIs -
The REST API's that support the solution must meet the following requirements:
Secure resources to the corporate VNet.
Allow deployment to a testing location within Azure while not incurring additional costs.
Automatically scale to double capacity during peak shipping times while not causing application downtime.
Minimize costs when selecting an Azure payment model.

Shipping data -
Data migration from on-premises to Azure must minimize costs and downtime.

Shipping website -
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Issues -

Windows Server 2016 VM -
The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

Shipping website and REST APIs -
The following error message displays while you are testing the website:
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource.
Origin 'http://test.wideworldimporters.com/' is therefore not allowed access.

Question
HOTSPOT -
You need to configure Azure CDN for the Shipping web site.
Which configuration options should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

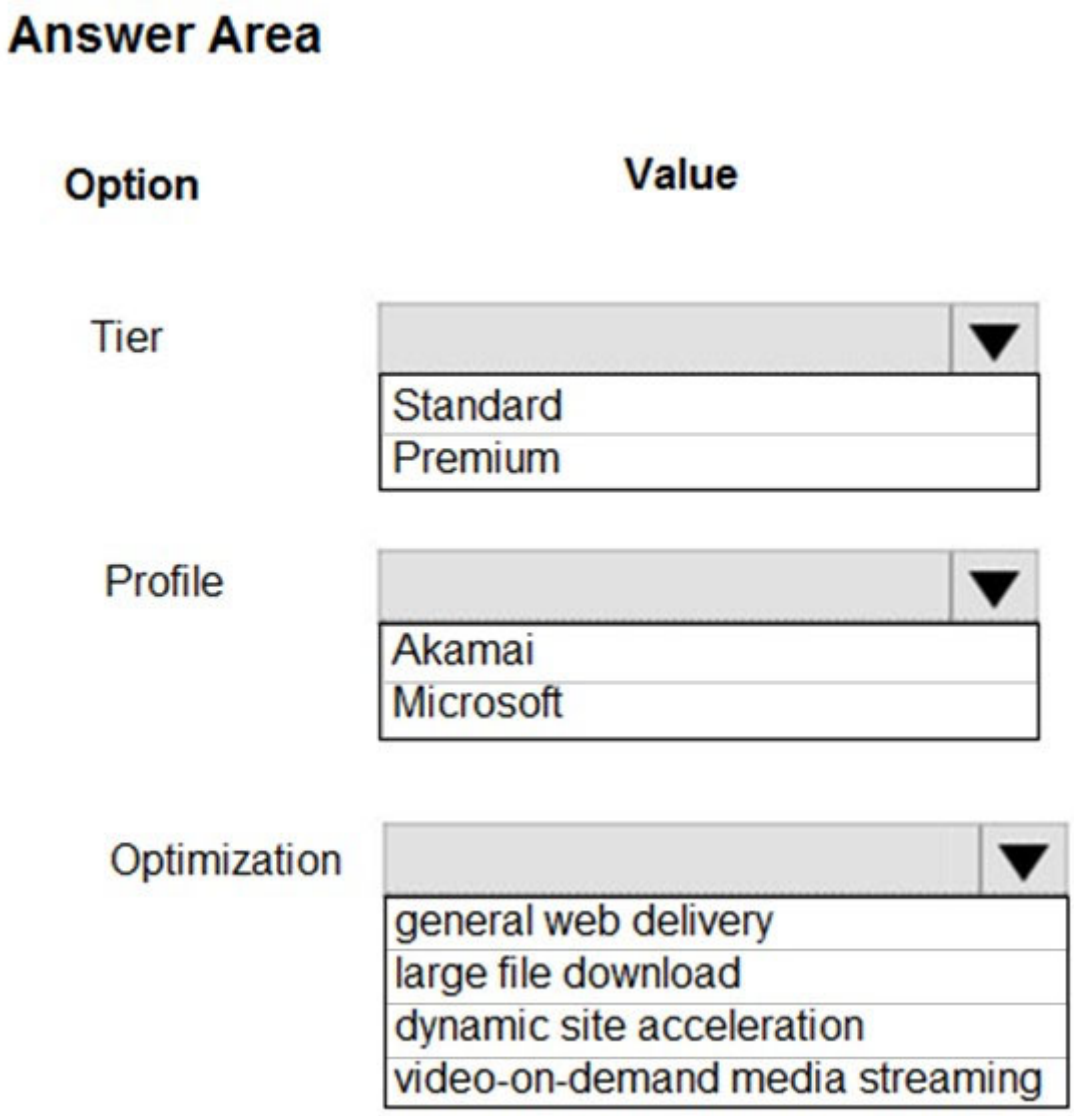

Correct Answer:

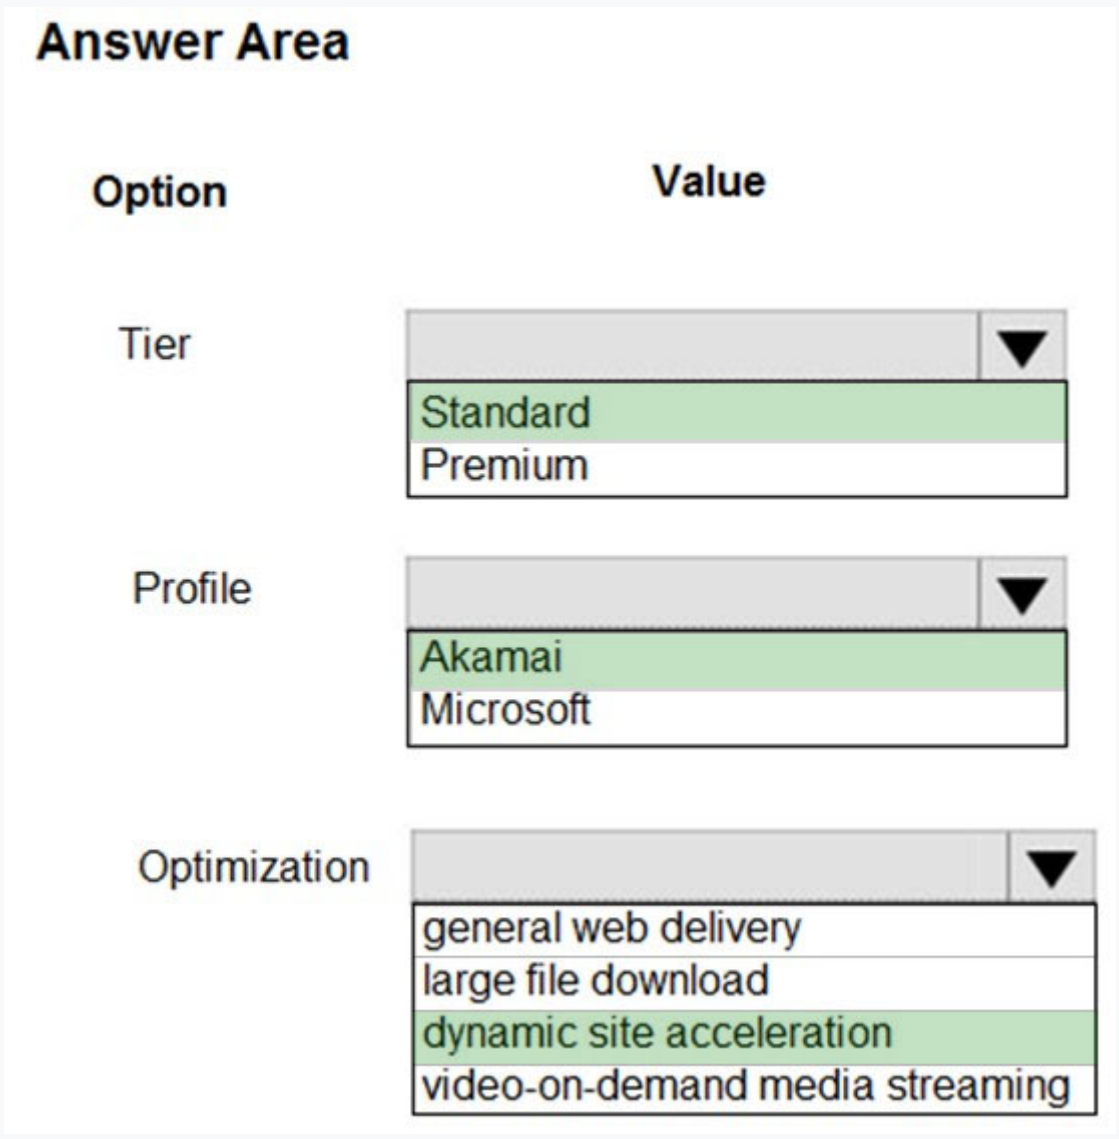

Scenario: Shipping website -

Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Tier: Standard -

Profile: Akamai -

Optimization: Dynamic site acceleration

Dynamic site acceleration (DSA) is available for Azure CDN Standard from Akamai, Azure CDN Standard from Verizon, and Azure CDN Premium from Verizon profiles.

DSA includes various techniques that benefit the latency and performance of dynamic content. Techniques include route and network optimization, TCP optimization, and more.

You can use this optimization to accelerate a web app that includes numerous responses that aren't cacheable. Examples are search results, checkout transactions, or real-time data. You can continue to use core Azure CDN caching capabilities for static data.
Reference:
https://docs.microsoft.com/en-us/azure/cdn/cdn-optimization-overview

coffecold Highly Voted 1 year ago
To prevent reading the cases multiple times:
Please see the spots below where you can find the questions (page/topic/question/subject)
Wide World Importers
46 7 1 Azure CDN tiers
46 7 2 azure tools on/with VM
49 12 1 secure Function app : auth + token + trigger type
49 12 2 secure Logic App : tool used
55 23 1 setup workflow in logic app
55 23 2 setup network to trigger logic app from external server
upvoted 26 times

OPT_001122 10 months, 3 weeks ago
Great help
upvoted 4 times

SoftSol Highly Voted 1 year, 6 months ago
Correct. Standard Akamai does support dynamic site acceleration. Microsoft requires front door, meaning additional costs. Verizon should also work I suppose, but isn't an options, so let's ignore that.
upvoted 10 times

Tyrel 3 weeks, 5 days ago
Akamai is retiring. Standard - Microsoft - General Web Delivery should be used instead.

General web delivery Note:
"If you are using an Azure CDN Standard from Akamai profile, select this optimization type if your average file size is smaller than 10 MB. Otherwise, if your average file size is larger than 10 MB, select Large file download from the Optimized for drop-down list."

Reference: https://learn.microsoft.com/en-us/azure/cdn/cdn-optimization-overview#general-web-delivery
upvoted 1 times

EliteAllen Most Recent 3 months ago
Azure CDN from Akamai is scheduled to be retired on October 31, 2023. You can no longer create new Azure CDN from Akamai after June 1, 2023. For more information, see Migrate CDN provider for guidance on migrating to another Azure CDN provider.

https://learn.microsoft.com/en-us/azure/cdn/cdn-features
upvoted 3 times

uncledana 3 months, 1 week ago
Standard
Microsoft
General Web Delivery

https://learn.microsoft.com/en-us/azure/cdn/cdn-features
upvoted 4 times

ExamDev 4 months, 1 week ago
About Akami... here some note:
Azure CDN from Akamai is scheduled to be retired on October 31, 2023. You can no longer create new Azure CDN from Akamai after June 1, 2023. For more information, see Migrate CDN provider for guidance on migrating to another Azure CDN provider.
upvoted 4 times

phantom31 3 months, 3 weeks ago
how many case study have you got totally?
upvoted 1 times

ExamDev 4 months, 1 week ago
Sorry Akamai made some typo :)
upvoted 1 times

yosri_c_sharp 4 months, 2 weeks ago
Azure CDN will be retiring Azure CDN Standard from Akamai on 31 October 2023
I thing the question will be removed
source: https://azure.microsoft.com/en-us/updates/azure-cdn-will-be-retiring-azure-cdn-standard-from-akamai-on-31-october-2023-transition-to-another-cdn-profile-to-avoid-servi/
upvoted 1 times

Question #2 *Topic 7*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment -

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk Server 2016. The VM runs the following workflows:

Ocean Transport `" This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.

Inland Transport `" This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

Container API `" This API provides container information including weight, contents, and other attributes.

Location API `" This API provides location information regarding shipping ports of call and trucking stops.

Shipping REST API `" This API provides shipping information for use and display on the shipping website.

Shipping Data -

The application uses MongoDB JSON document storage database for all container and transport information.

Shipping Web Site -

The site displays shipping container tracking information and container contents. The site is located at http://shipping.wideworldimporters.com/

Proposed solution -

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:

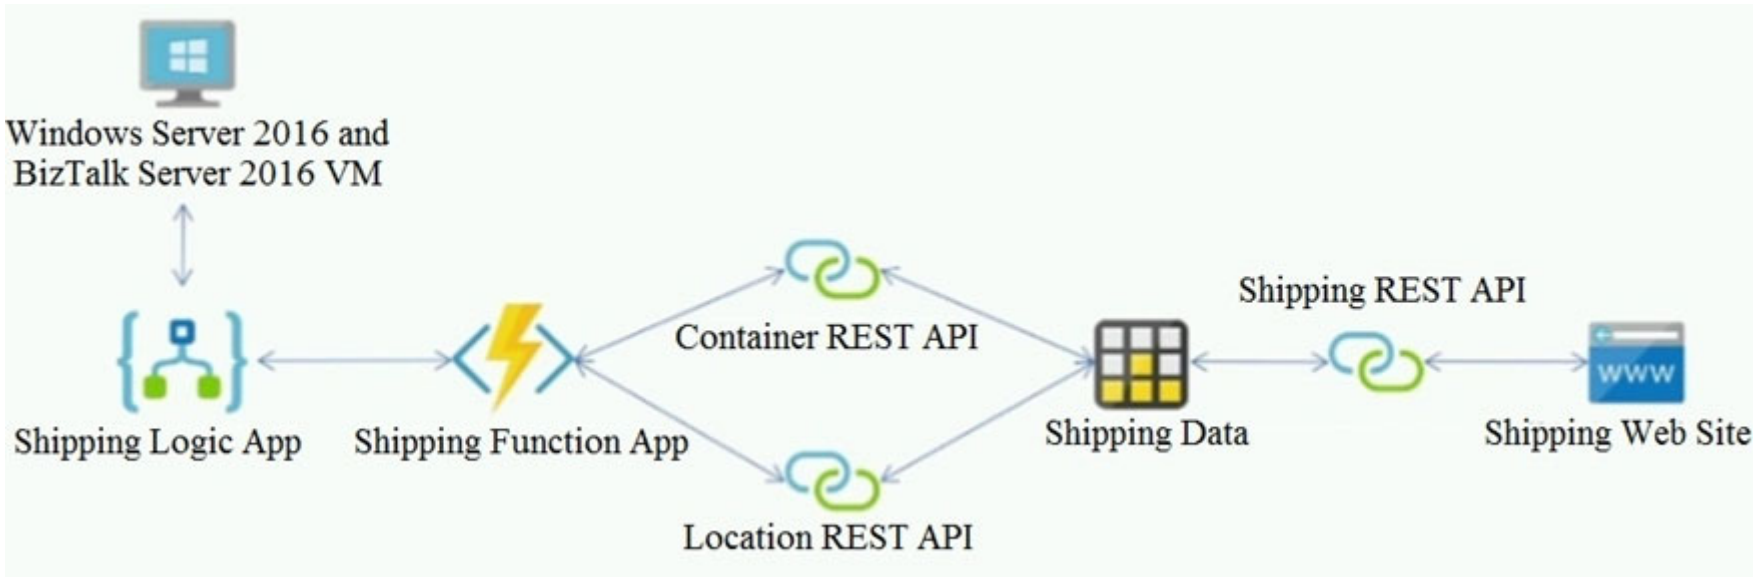

Requirements -

Shipping Logic app -
The Shipping Logic app must meet the following requirements:
Support the ocean transport and inland transport workflows by using a Logic App.
Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app -
Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

REST APIs -
The REST API's that support the solution must meet the following requirements:
Secure resources to the corporate VNet.
Allow deployment to a testing location within Azure while not incurring additional costs.
Automatically scale to double capacity during peak shipping times while not causing application downtime.
Minimize costs when selecting an Azure payment model.

Shipping data -
Data migration from on-premises to Azure must minimize costs and downtime.

Shipping website -
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Issues -

Windows Server 2016 VM -
The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

Shipping website and REST APIs -
The following error message displays while you are testing the website:
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource.
Origin 'http://test.wideworldimporters.com/' is therefore not allowed access.

Question
HOTSPOT -
You need to correct the VM issues.
Which tools should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Issue	Tool
Backup and Restore	Azure Site Recovery Azure Backup Azure Data Box Azure Migrate
Performance	Azure Network Watcher Azure Traffic Manager ExpressRoute Accelerated Networking

Correct Answer:

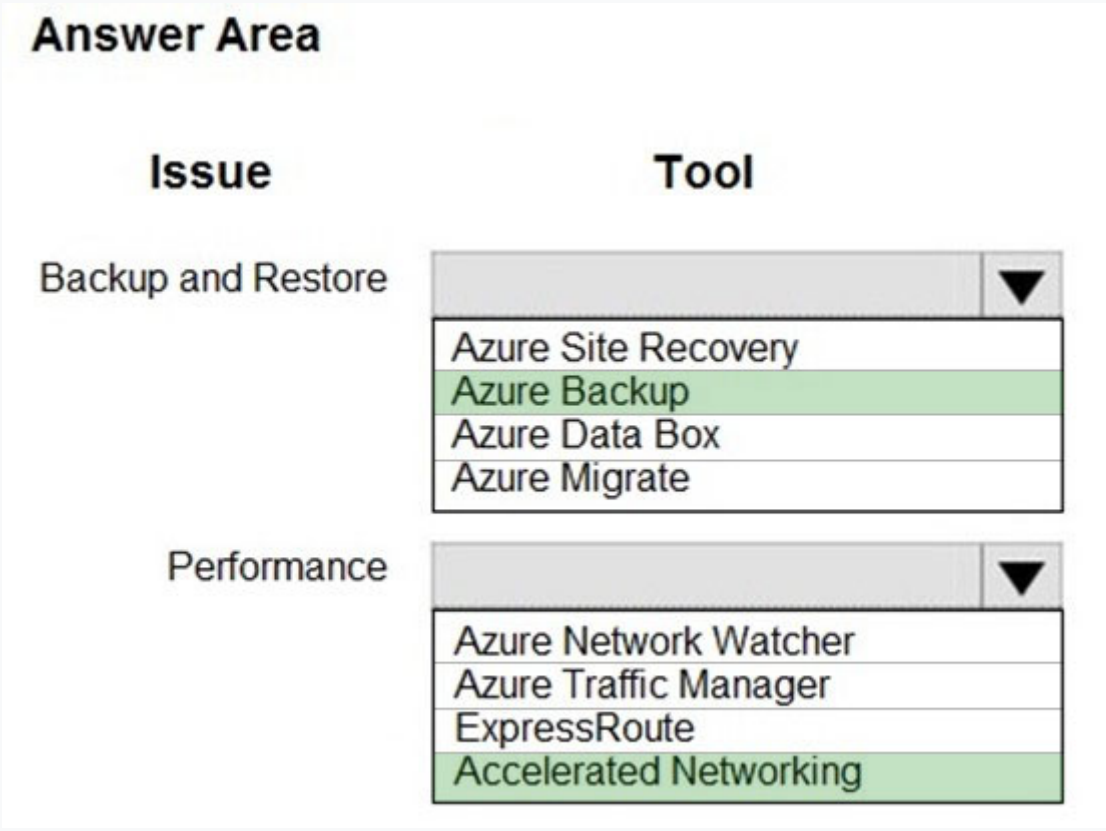

Box 1: Azure Backup -

The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

In-Place restore of disks in IaaS VMs is a feature of Azure Backup.

Performance: Accelerated Networking

Scenario: The VM shows high network latency, jitter, and high CPU utilization.

Box 2: Accelerated networking -

The VM shows high network latency, jitter, and high CPU utilization.

Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, greatly improving its networking performance. This high-performance path bypasses the host from the datapath, reducing latency, jitter, and CPU utilization, for use with the most demanding network workloads on supported VM types.

Reference:

https://azure.microsoft.com/en-us/blog/an-easy-way-to-bring-back-your-azure-vm-with-in-place-restore/

JOSEEVILLASMIL Highly Voted 7 months, 1 week ago

Hoow the f could you memorize all this.

upvoted 23 times

coffecold Highly Voted 1 year ago

Accelerated Networking only works in conjunction with a Azure Virtual Network (VNet). That is mentioned in the case and is a clue as well.

upvoted 8 times

yosri_c_sharp Most Recent 4 months, 2 weeks ago

Virtual Machines no longer exists on az 204

upvoted 3 times

nvtienanh 1 year ago

The given answer is correct.
https://learn.microsoft.com/en-us/azure/virtual-network/accelerated-networking-overview#benefits

upvoted 4 times

Topic 8 - Testlet 10

Question #1 *Topic 8*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web

API project that is deployed as an Azure Web App.

Overall architecture -

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing -

Employees may upload receipts in two ways:

Uploading using an Azure Files mounted folder

Uploading using the web application

Data Storage -

Receipt and employee information is stored in an Azure SQL database.

Documentation -

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for

Azure File upload, and instructions on how to configure the mounted folder.

Solution details -

Users table -

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application -

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing -

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored

in Azure Blob
Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging -
Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled.
Application Insights must always contain all log messages.

Requirements -

Receipt processing -
Concurrent processing of a receipt must be prevented.

Disaster recovery -
Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security -
User's SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
Receipt data must always be encrypted at rest.
All data must be protected in transit.
User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues -

Upload format issue -
Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File
Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal
Server error page.

Capacity issue -
During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

Log capacity issue -
Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Application code -

Processing.cs -

```
PC01 public static class Processing
PC02 {
PC03   public static class Function
PC04   {
PC05     [FunctionName("IssueWork")]
PC06     public static async Task Run([TimerTrigger("0 */5 * * * *")] TimerInfo timer, ILogger
log)
PC07     {
PC08       var container = await GetCloudBlobContainer();
PC09       foreach (var fileItem in await ListFiles())
PC10       {
PC11         var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12         var ms = new MemoryStream();
PC13         await file.DownloadToStreamAsync(ms);
PC14         var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15         await blob.UploadFromStreamAsync(ms);
PC16
PC17       }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob(CloudBlockBlob sourceBlob)
PC20     {
PC21        . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25       var cloudBlobClient = new CloudBlobClient(new Uri(". . ."), await GetCredentials());
PC26
PC27       await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28       return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32        . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36        . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(". . .");
PC38   }

PC39 }
```

Database.cs -

```
DB01 public class Database
DB02 {
DB03   private string ConnectionString =
DB04
DB05   public async Task<object> LoadUserDetails(string userId)
DB06   {
DB07
DB08     return await policy.ExecuteAsync(async () =>
DB09     {
DB10       using (var connection = new SqlConnection(ConnectionString))
DB11       {
DB12         await connection.OpenAsync();
DB13         using (var command = new SqlCommand("…", connection))
DB14         using (var reader = command.ExecuteReader())
DB15         {
DB16           …
DB17         }
DB18       }
DB19     });
DB20   }

DB21 }
```

ReceiptUploader.cs -

```
RU01 public class ReceiptUploader
RU02 {
RU03   public async Task UploadFile(string file, byte[] binary)
RU04   {
RU05     var httpClient = new HttpClient();
RU06     var response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU07     while (ShouldRetry(response))
RU08     {
RU09       response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU10     }
RU11   }
RU12   private bool ShouldRetry(HttpResponseMessage response)
RU13   {
RU14
RU15   }
RU16 }
```

ConfigureSSE.ps1 -

```
CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "..." -AccountName "..."
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "..."
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "..."
CS04 Set-AzureRmKeyVaultAccessPolicy `
CS05   -VaultName $keyVault.VaultName `
CS06   -ObjectId $storageAccount.Identity.PrincipalId `
CS07
CS08
CS09 Set-AzureRmStorageAccount `
CS10   -ResourceGroupName $storageAccount.ResourceGroupName `
CS11   -AccountName $storageAccount.StorageAccountName `
CS12   -EnableEncryptionService File `
CS13   -KeyvaultEncryption `
CS14   -KeyName $key.Name
CS15   -KeyVersion $key.Version `
CS16   -KeyVaultUri $keyVault.VaultUri
```

Question

DRAG DROP -

You need to add code at line PC32 in Processing.cs to implement the GetCredentials method in the Processing class.

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:

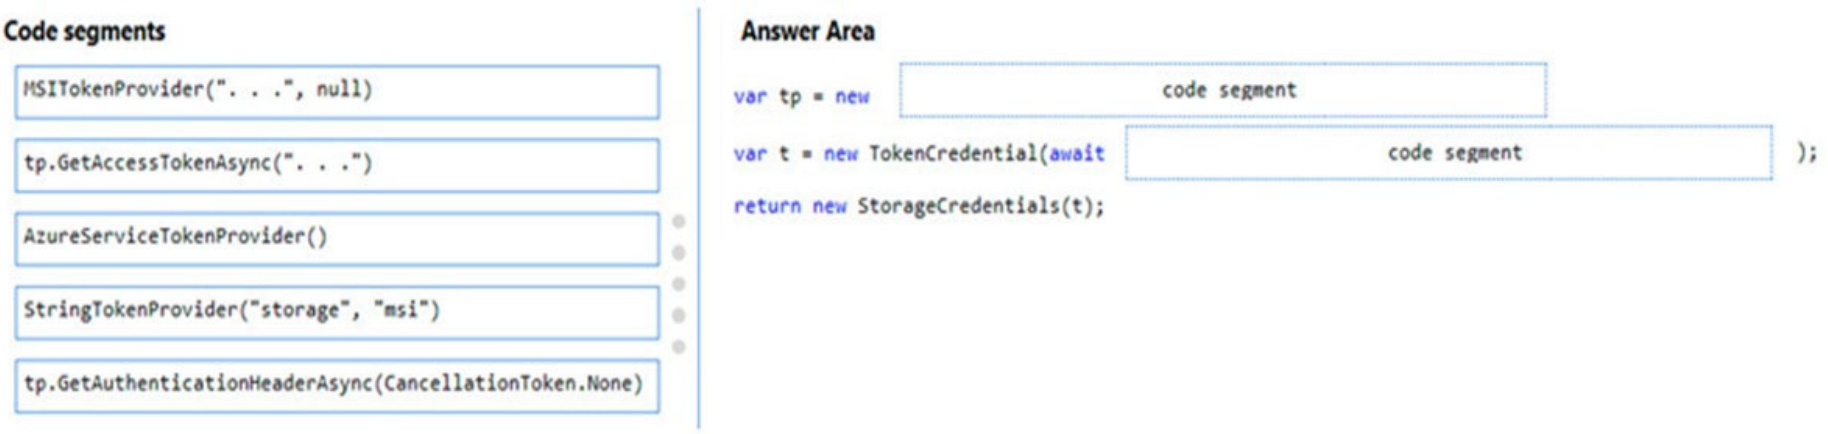

Correct Answer:

Box 1: AzureServiceTokenProvider()

Box 2: tp.GetAccessTokenAsync("..")

Acquiring an access token is then quite easy. Example code:

private async Task<string> GetAccessTokenAsync()

{

var tokenProvider = new AzureServiceTokenProvider();

return await tokenProvider.GetAccessTokenAsync("https://storage.azure.com/");

}

Reference:

https://joonasw.net/view/azure-ad-authentication-with-azure-storage-and-managed-service-identity

coffecold Highly Voted 1 year ago

To prevent reading the cases multiple times:
Please see the spots below where you can find the questions (page/topic/question/subject)
Litware Inc
47 8 1 code : GetCredentials method
47 8 2 code : disaster recovery requirements
51 14 1 code : set encryption with request to blob storage
51 14 2 ARM set keyVault access policy
53 20 1 prevent blob writing concurrency problems
54 20 2 resolve capacity issue at function processing
54 20 3 resolve capacity issue at extensive logging

upvoted 29 times

warchoon 6 months, 4 weeks ago

Wide World Importers 7/1 7/2 12/1 12/2 23/1 23/2
Litware Inc. 8/1 8/2 14/1 14/2 20/1 20/2 20/3
Coho Winery produces 9/1 9/2 22/1 22/2 28/1 28/2
VanArsdel, Ltd. 10/1 10/2 10/3 10/4 15/1 15/2 21/1 29/1 29/2 30/1
Contoso, Ltd. 11/1 11/2 11/3 11/4 11/5 16/1 16/2 27/1 27/2 31/1 31/2 31/3
City Power & Light 13/1 13/2 13/3 13/4 13/5 13/6 17/1 17/2 18/1 24/1 24/2 32/1
Proseware, Inc. 19/1 19/2 19/3 19/4 25/1 25/2 25/3 26/1

upvoted 12 times

alejary 6 months, 1 week ago

I know that 7/1 means topic 7, question 1, but where we can find use case description?

upvoted 1 times

sadish003 6 months, 2 weeks ago

Sorry, couldn't understand what you mean.
You mean this code (47 8 1) should be placed in the URL to access this content?
Could you please elaborate more on this?

upvoted 1 times

karthikwarrior 8 months, 1 week ago

did not follow or understand what you meant

upvoted 1 times

OPT_001122 10 months, 3 weeks ago

Great Help

upvoted 1 times

rqb11 Highly Voted 2 years, 6 months ago

Answer is correct: https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity?context=azure%2Factive-directory%2Fmanaged-identities-azure-resources%2Fcontext%2Fmsi-context&tabs=python#asal

upvoted 26 times

Faizs 2 years, 5 months ago

Absolutely

upvoted 3 times

devex Most Recent 2 months ago

Answer is correct.

TokenCredential class receives a string in the constructor:
- https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.auth.tokencredential?view=azure-dotnet#constructors

So the answer cannot be MSITokenProvider.GetAuthenticationHeaderAsync which returns an AuthenticationHeaderValue:
- https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.management.resourcemanager.fluent.authentication.msitokenprovider.getauthenticationheaderasync

AzureServiceTokenProvider.GetAccessTokenAsync on the other hand returns a string:
- https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.services.appauthentication.azureservicetokenprovider.getaccesstokenasync

upvoted 2 times

adilkhan 9 months ago

https://learn.microsoft.com/en-us/dotnet/api/overview/azure/service-to-service-authentication?view=azure-dotnet Answers are correct

upvoted 2 times

OPT_001122 10 months, 3 weeks ago

AzureServiceTokenProvider
GetAccessTokenAsync(...)

upvoted 2 times

coffecold 1 year ago

This is Microsoft.Azure.Services.AppAuthentication and is since 2019 legacy.
Instead use the Azure Identity client library for .NET Azure.Identity
using Azure.Identity;
using Azure.Security.KeyVault.Secrets;

// Create a secret client using the DefaultAzureCredential
var client = new SecretClient(new Uri("https://myvault.azure.vaults.net/"), new DefaultAzureCredential());
etc.

upvoted 6 times

edmond 10 months, 2 weeks ago

Or use ManagedIdentityCredential() if you want to make sure it's using managed identity exclusively

upvoted 3 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer:
- AzureServiceTokenProvider
- GetAccessTokenAsync(...)

upvoted 3 times

UnknowMan 2 years, 5 months ago

Yep :

- AzureServiceTokenProvider
- GetAccessTokenAsync

upvoted 5 times

Faizs 2 years, 5 months ago

Very correct

upvoted 2 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web

API project that is deployed as an Azure Web App.

Overall architecture -

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing -

Employees may upload receipts in two ways:

Uploading using an Azure Files mounted folder

Uploading using the web application

Data Storage -

Receipt and employee information is stored in an Azure SQL database.

Documentation -

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for

Azure File upload, and instructions on how to configure the mounted folder.

Solution details -

Users table -

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application -

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing -

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob

Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must

remain valid if the email is forwarded to another user.

Logging -
Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled.
Application Insights must always contain all log messages.

Requirements -

Receipt processing -
Concurrent processing of a receipt must be prevented.

Disaster recovery -
Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security -
User's SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
Receipt data must always be encrypted at rest.
All data must be protected in transit.
User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues -

Upload format issue -
Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File
Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal
Server error page.

Capacity issue -
During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

Log capacity issue -
Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Application code -

Processing.cs -

```
PC01 public static class Processing
PC02 {
PC03   public static class Function
PC04   {
PC05     [FunctionName("IssueWork")]
PC06     public static async Task Run([TimerTrigger("0 */5 * * * *")] TimerInfo timer, ILogger
log)
PC07     {
PC08       var container = await GetCloudBlobContainer();
PC09       foreach (var fileItem in await ListFiles())
PC10       {
PC11         var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12         var ms = new MemoryStream();
PC13         await file.DownloadToStreamAsync(ms);
PC14         var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15         await blob.UploadFromStreamAsync(ms);
PC16
PC17       }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob(CloudBlockBlob sourceBlob)
PC20     {
PC21        . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25       var cloudBlobClient = new CloudBlobClient(new Uri(". . ."), await GetCredentials());
PC26
PC27       await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28       return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32        . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36        . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(". . .");
PC38   }

PC39 }
```

Database.cs -

```
DB01 public class Database
DB02 {
DB03   private string ConnectionString =
DB04
DB05   public async Task<object> LoadUserDetails(string userId)
DB06   {
DB07
DB08     return await policy.ExecuteAsync(async () =>
DB09     {
DB10       using (var connection = new SqlConnection(ConnectionString))
DB11       {
DB12         await connection.OpenAsync();
DB13         using (var command = new SqlCommand("…", connection))
DB14         using (var reader = command.ExecuteReader())
DB15         {
DB16           …
DB17         }
DB18       }
DB19     });
DB20   }

DB21 }
```

ReceiptUploader.cs -

```
RU01 public class ReceiptUploader
RU02 {
RU03   public async Task UploadFile(string file, byte[] binary)
RU04   {
RU05     var httpClient = new HttpClient();
RU06     var response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU07     while (ShouldRetry(response))
RU08     {
RU09       response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU10     }
RU11   }
RU12   private bool ShouldRetry(HttpResponseMessage response)
RU13   {
RU14
RU15   }
RU16 }
```

ConfigureSSE.ps1 -

```
CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "..." -AccountName "..."
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "..."
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "..."
CS04 Set-AzureRmKeyVaultAccessPolicy `
CS05  -VaultName $keyVault.VaultName `
CS06  -ObjectId $storageAccount.Identity.PrincipalId `
CS07
CS08
CS09 Set-AzureRmStorageAccount `
CS10  -ResourceGroupName $storageAccount.ResourceGroupName `
CS11  -AccountName $storageAccount.StorageAccountName `
CS12  -EnableEncryptionService File `
CS13  -KeyvaultEncryption `
CS14  -KeyName $key.Name
CS15  -KeyVersion $key.Version `
CS16  -KeyVaultUri $keyVault.VaultUri
```

Question

DRAG DROP -

You need to ensure disaster recovery requirements are met.

What code should you add at line PC16?

To answer, drag the appropriate code fragments to the correct locations. Each code fragment 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:

Values

- true
- SingleTransferContext
- ShouldTransferCallbackAsync
- false
- DirectoryTransferContext
- ShouldOverwriteCallbackAsync

Answer Area

```
var copyOptions = new CopyOptions { };
var context = new      [Value]  = (source, destination) => Task.FromResult(true);
context.       [Value]  = (source, destination) => Task.FromResult(true);
await TransferManager.CopyAsync(blob, GetDRBlob(blob), isServiceCopy:  [Value]
  , context: context, options:copyOptions);
```

Correct Answer:

Scenario: Disaster recovery. Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Box 1: DirectoryTransferContext -
We transfer all files in the directory.
Note: The TransferContext object comes in two forms: SingleTransferContext and DirectoryTransferContext. The former is for transferring a single file and the latter is for transferring a directory of files.
Box 2: ShouldTransferCallbackAsync
The DirectoryTransferContext.ShouldTransferCallbackAsync delegate callback is invoked to tell whether a transfer should be done.

Box 3: False -
If you want to use the retry policy in Copy, and want the copy can be resume if break in the middle, you can use SyncCopy (isServiceCopy = false).
Note that if you choose to use service side copy ('isServiceCopy' set to true), Azure (currently) doesn't provide SLA for that. Setting 'isServiceCopy' to false will download the source blob loca
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-use-data-movement-library https://docs.microsoft.com/en-us/dotnet/api/microsoft.windowsazure.storage.datamovement.directorytransfercontext.shouldtransfercallbackasync?view=azure- dotnet

rdemontis Highly Voted 2 years, 7 months ago
Answers in box 1 and 2 arent' correct. They should be SingleTranferConext and ShouldOverwriteCallbackAsync because we are copying a single blob (see CopyAsync method). We are inside a foreach loop that scan each file in the share.
upvoted 54 times

ray01 2 years, 6 months ago
Correct
1. https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.datamovement.singletransfercontext?view=azure-dotnet
2. According to a link above, there is only "ShouldOverwriteCallbackAsync" available for "SingleTranferConext "
3. https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.datamovement.transfermanager.copyasync?view=azure-dotnet#Microsoft_Azure_Storage_DataMovement_TransferManager_CopyAsync_Microsoft_Azure_Storage_Blob_CloudBlob_Microsoft_Azure_Storage_Blob_CloudBlob_System_Boolean_Microsoft_Azure_Storage_DataMovement_CopyOptions_Microsoft_Azure_Storage_DataMovement_SingleTransferContext_
The question is should it be server side copy or not? I would say "true".
upvoted 8 times

d0bermannn 2 years, 2 months ago
yes, foreach construction applied
upvoted 2 times

clarionprogrammer Highly Voted 2 years, 6 months ago
SingleTranferConext
ShouldOverwriteCallbackAsync
true
upvoted 27 times

ZodiaC 2 years, 3 months ago
ITS FALSE 10000000000%
upvoted 3 times

ZodiaC 2 years, 3 months ago
Srry im didnt read well
upvoted 7 times

NombreFalso 8 months, 1 week ago

It's all good buddy

upvoted 2 times

anastakasim 2 years, 6 months ago

The last box is "true"? really?

upvoted 2 times

ExamDev Most Recent 4 months, 1 week ago

About isServiceCopy:
A flag indicating whether the copy is service-side asynchronous copy or not. If this flag is set to true, service-side asychronous copy will be used; if this flag is set to false, file is downloaded from source first, then uploaded to destination.

upvoted 1 times

Kluk_Kluk 7 months, 2 weeks ago

https://learn.microsoft.com/en-us/azure/storage/common/storage-use-data-movement-library#track-transfer-progress

upvoted 1 times

Kluk_Kluk 7 months, 2 weeks ago

"To achieve this scenario, we need to create a TransferContext object. The TransferContext object comes in two forms: SingleTransferContext and DirectoryTransferContext. The former is for transferring a single file and the latter is for transferring a directory of files."

upvoted 1 times

Kluk_Kluk 7 months, 2 weeks ago

the foreach in the lines above PC16 indicates single file transfer

upvoted 2 times

OPT_001122 10 months, 3 weeks ago

1.SingleTranferConext
2.ShouldOverwriteCallbackAsync
3.true (If this flag is set to true, service-side asychronous copy will be used)

upvoted 3 times

coffecold 1 year ago

Old version 11 code Microsoft.Azure.Storage.DataMovement
Version 12 : Azure.Storage.Blobs
You would probably use :
StartCopyFromUriAsync(Uri, BlobCopyFromUriOptions, CancellationToken)

upvoted 2 times

gmishra88 1 year ago

For the non Microsoft stable tied horses, Task.FromResult(true) creates a task that returns true.

upvoted 4 times

coffecold 1 year ago

usage : See https://github.com/Azure/azure-storage-net-data-movement/blob/master/samples/DataMovementSamples/DataMovementSamples/Samples.cs

upvoted 1 times

gmishra88 1 year ago

Also Microsoft, you do not need to know .Net to be able to write AZ-204.

upvoted 2 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer:
1.SingleTranferConext (because it is inside loop)
2.ShouldOverwriteCallbackAsync (for singleTransferContext this is only option)
3.true (If this flag is set to true, service-side asychronous copy will be used)

upvoted 5 times

leonidn 1 year, 9 months ago

SingleTransferContext because we transfer blob by blob because line PC16 is in the ForEach.
IsServiceCopy true to allow service side asynchronous processing instead of downloading/uploading blobs.

ShouldOverwriteCallbackAsync because ShouldTransferCallbackAsync is not a property of SingleTransferContext.

upvoted 4 times

asdasdasg2 1 year, 9 months ago

No one is talking about how the line on the first answer box is invalid no matter that answer you put.

I believe it's a copy paste error and the line should be:
var context = new [ ];

with answer SingleTransferContext

upvoted 2 times

asdasdasg2 1 year, 9 months ago

i realize that is confusing, i meant that the line in the hot area should be :
var context = new ____ ();

upvoted 2 times

BrettusMaximus 2 years, 6 months ago

isServiceCopy: True
Boolean
A flag indicating whether the copy is service-side asynchronous copy or not. If this flag is set to true, service-side asychronous copy will be used; if this flag is set to false, file is downloaded from source first, then uploaded to destination

upvoted 2 times

Topic 9 - Testlet 11

Question #1 *Topic 9*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

LabelMaker app -

Coho Winery produces, bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends data to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the LabelMaker application to include artwork and text for custom label designs.

Requirements. Data -

You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using SQL.

Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Requirements. Security -

You have the following security requirements:

Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.

▪

External partners must use their own credentials and authenticate with their organization's identity management solution.

External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.

Storage of e-commerce application settings must be maintained in Azure Key Vault.

E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).

Conditional access policies must be applied at the application level to protect company content.

The LabelMaker application must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

Requirements. LabelMaker app -

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure

Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture -

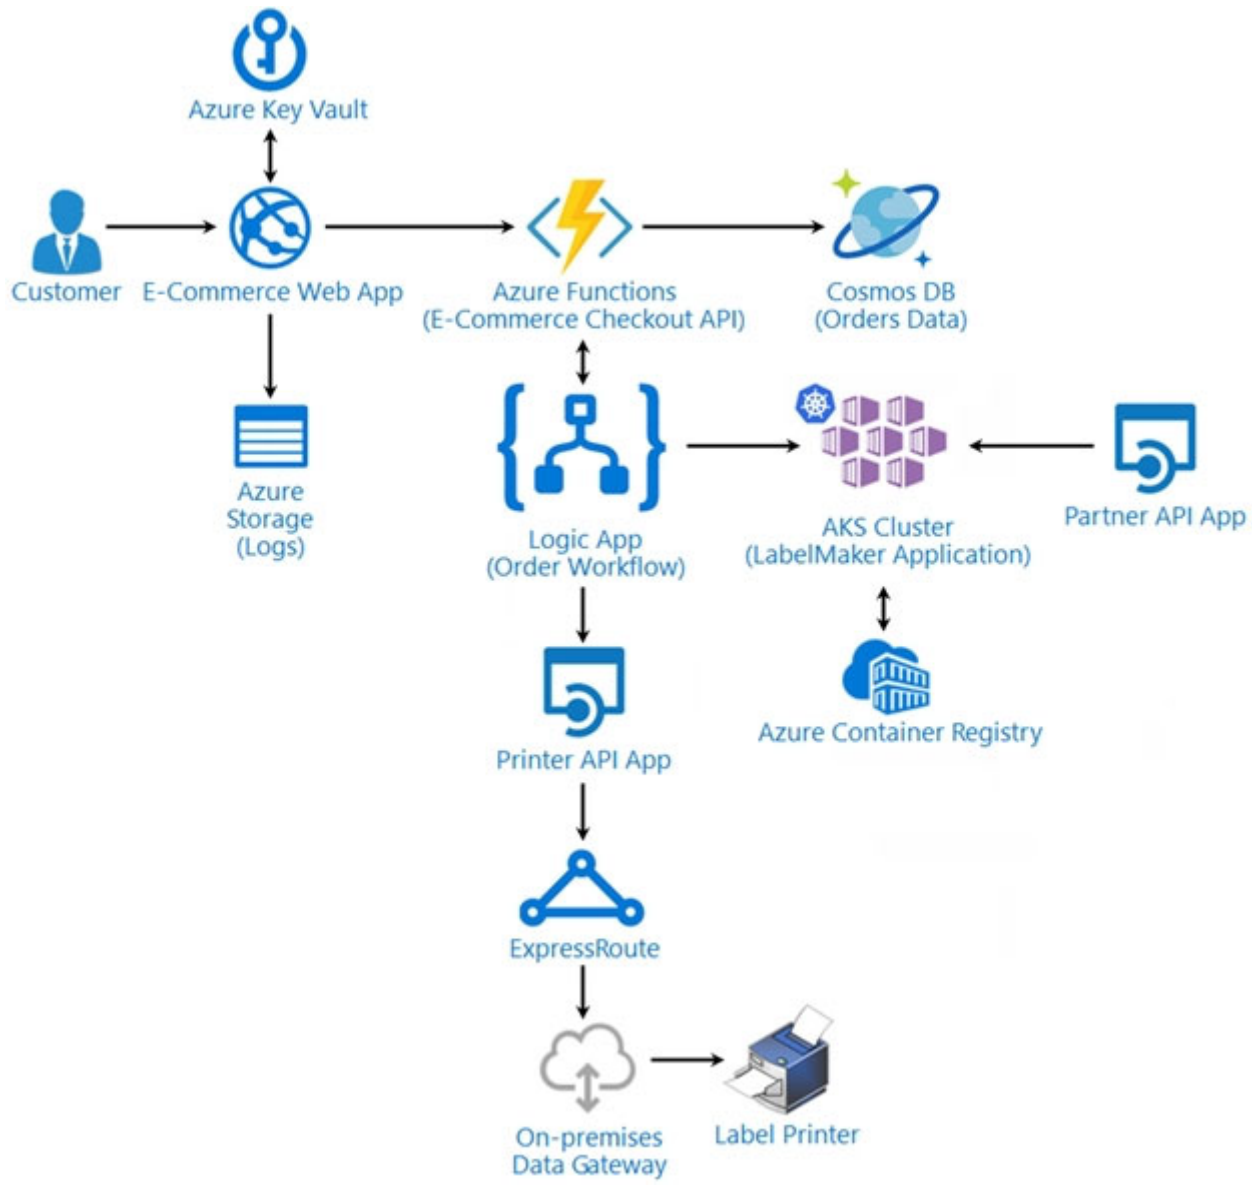

Issues -

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order.json -

Relevant portions of the app files are shown below. Line numbers are included for reference only.

This JSON file contains a representation of the data for an order that includes a single item.

Order.json -

Order.json

```
{
  "id" : 1,
  "customers" : [
  {
     "familyName" : "Doe",
     "givenName" : "John",
     "customerid" : 5
   }
  ],
  "line_items" : [
  {
    "fulfillable_quantity" : 1,
    "id": 6,
    "price" : "199.99" ,
    "product_id" : 7513594,
    "quantity": 1,
    "requires_shipping" : true ,
    "sku": "SFC-342-N" ,
    "title" : "Surface Go" ,
    "vendor" : "Microsoft" ,
    "name" : "Surface Go - 8GB" ,
    "taxable" : true ,
    "tax_lines" : [
    {
      "title" : "State Tax" ,
      "price" : "3.98" ,
      "rate" : 0.06
     }
    ],
    "total_discount" : "5.00" ,
    "discount_allocations" : [
     {
       "amount" : "5.00" ,
       "discount_application_index" : 2
       }
    ]
```

Question

HOTSPOT -

You need to configure Azure Cosmos DB.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Setting	Value
Consistency Level	Strong Bounded-staleness Session Eventual
API	SQL MongoDB Graph Table

Correct Answer:

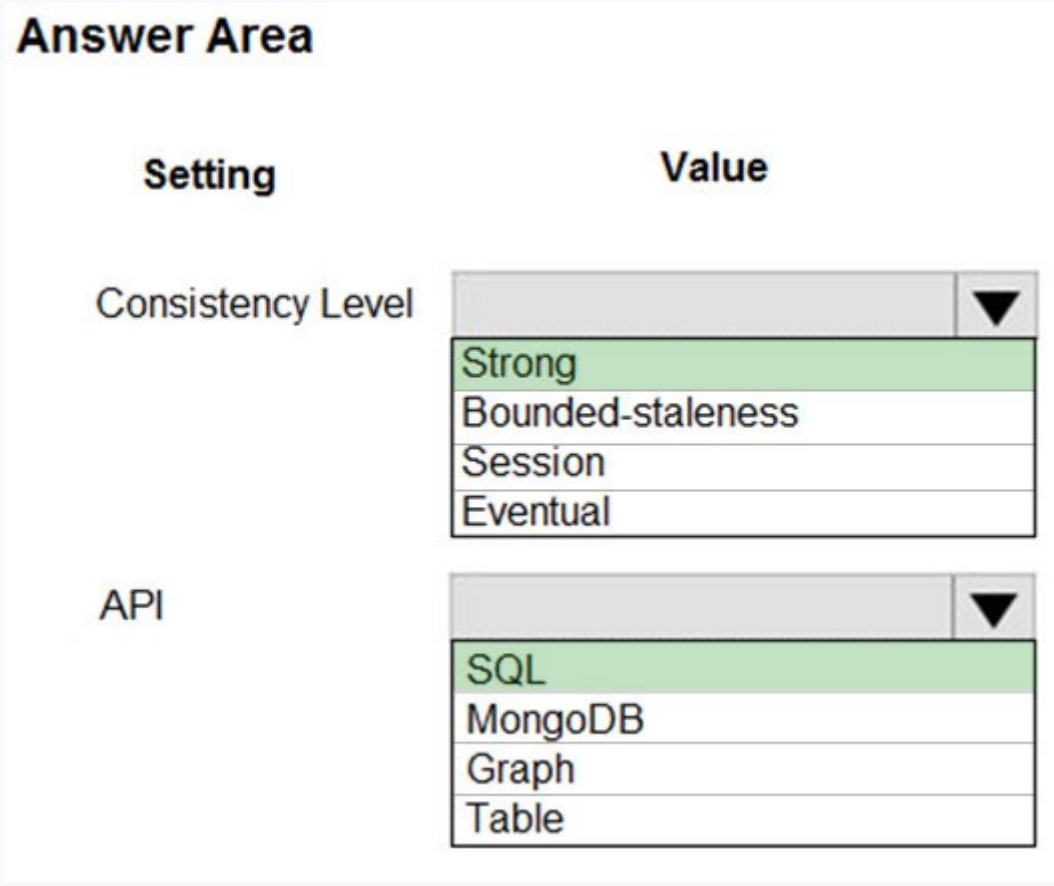

Box 1: Strong -

When the consistency level is set to strong, the staleness window is equivalent to zero, and the clients are guaranteed to read the latest committed value of the write operation.

Scenario: Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Note: You can choose from five well-defined models on the consistency spectrum. From strongest to weakest, the models are: Strong, Bounded staleness,

Session, Consistent prefix, Eventual

Box 2: SQL -

Scenario: You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using Structured Query Language (SQL).

coffecold Highly Voted 1 year ago

To prevent reading the cases multiple times:
Please see the spots below where you can find the questions (page/topic/question/subject)
Coho Winery
47 9 1 configure Azure Cosmos DB
47 9 2 SQL
54 22 1 Review logs/history to troubleshoot workflow
54 22 2 ARM function configuration : retries etc
57 28 1 workflow for Azure Container Registry
57 28 2 access data from the user claim object

upvoted 17 times

OPT_001122 10 months, 3 weeks ago

Great Help!!

upvoted 2 times

SenseiJC 11 months, 1 week ago

Thanks so much for this! Big help!

upvoted 3 times

noip Highly Voted 2 years, 2 months ago

correct 100%

upvoted 17 times

SivajiTheBoss Most Recent 1 year, 7 months ago

correct Answer:
1. Strong consistency
2. Sql Api

upvoted 4 times

asdasdasg2 1 year, 9 months ago

Answer is correct - Strong consistency matches requirements and MS recommends that if you have no preference or specific reason to use any of the other cosmos APIs, you should use the SQL api

upvoted 5 times

Sukon_Desknot 2 years, 2 months ago

(/◕ヮ◕)/*:･ﾟ✧
upvoted 6 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

LabelMaker app -

Coho Winery produces, bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends data to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the LabelMaker application to include artwork and text for custom label designs.

Requirements. Data -

You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using SQL.

Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Requirements. Security -

You have the following security requirements:

Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.

▪

External partners must use their own credentials and authenticate with their organization's identity management solution.

External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.

Storage of e-commerce application settings must be maintained in Azure Key Vault.

E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).

Conditional access policies must be applied at the application level to protect company content.

The LabelMaker application must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

Requirements. LabelMaker app -

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure

Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture -

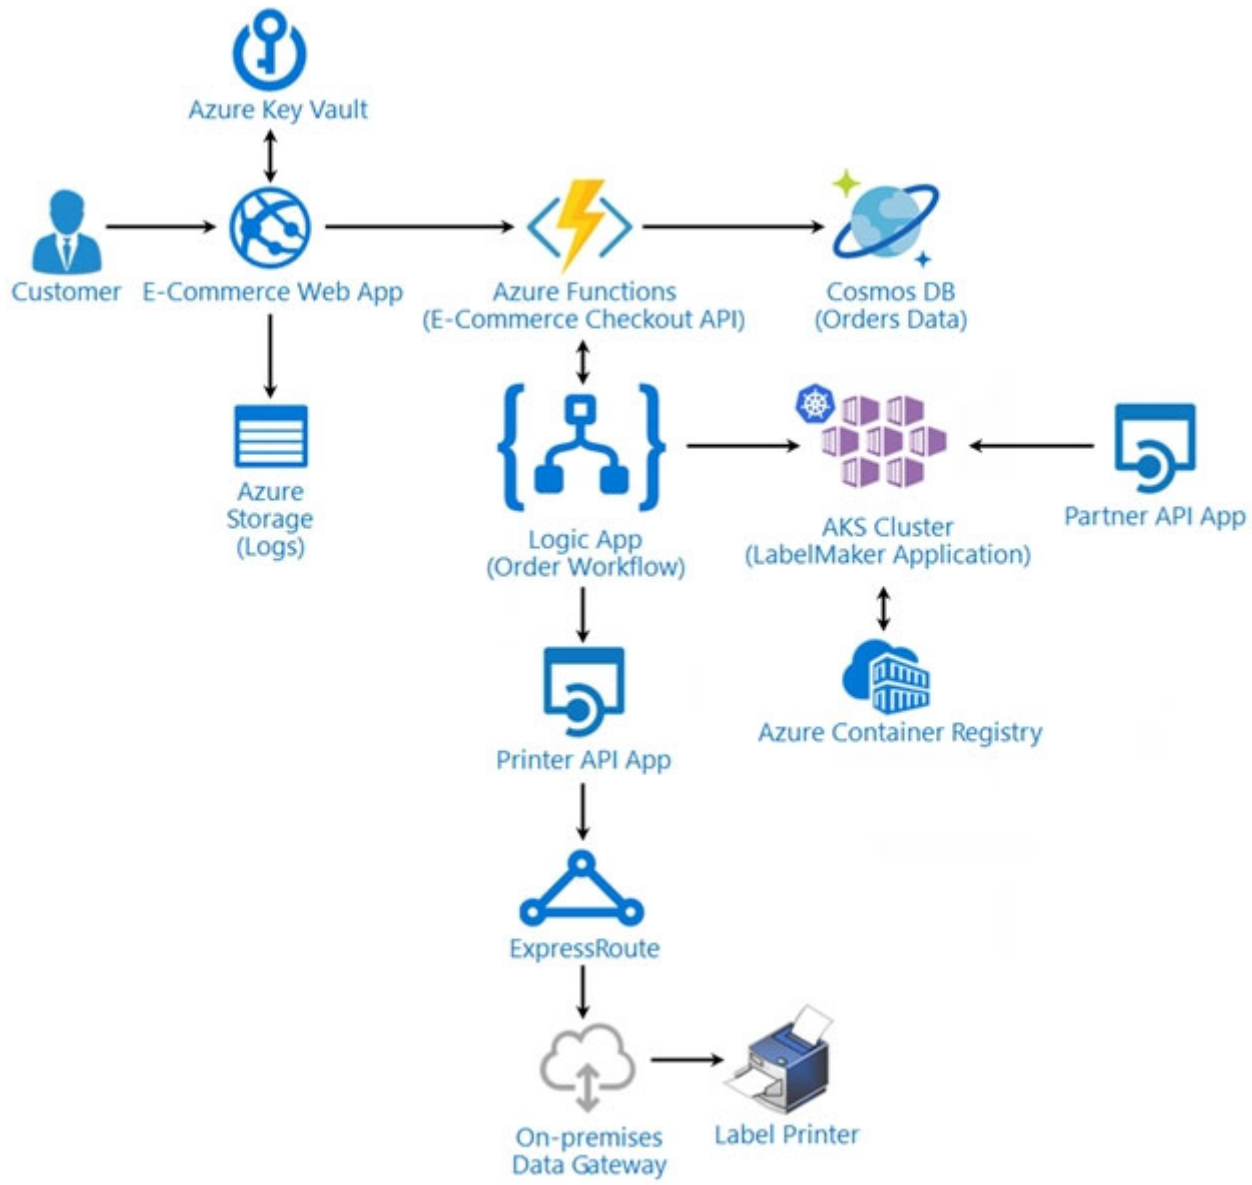

Issues -

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order.json -

Relevant portions of the app files are shown below. Line numbers are included for reference only.

This JSON file contains a representation of the data for an order that includes a single item.

Order.json -

Order.json

```
{
  "id" : 1,
  "customers" : [
  {
     "familyName" : "Doe",
     "givenName" : "John",
     "customerid" : 5
   }
  ],
  "line_items" : [
  {
    "fulfillable_quantity" : 1,
    "id": 6,
    "price" : "199.99" ,
    "product_id" : 7513594,
    "quantity": 1,
    "requires_shipping" : true ,
    "sku": "SFC-342-N" ,
    "title" : "Surface Go" ,
    "vendor" : "Microsoft" ,
    "name" : "Surface Go - 8GB" ,
    "taxable" : true ,
    "tax_lines" : [
    {
      "title" : "State Tax" ,
      "price" : "3.98" ,
      "rate" : 0.06
     }
    ],
    "total_discount" : "5.00" ,
    "discount_allocations" : [
     {
       "amount" : "5.00" ,
       "discount_application_index" : 2
       }
    ]
```

Question

HOTSPOT -

You need to retrieve all order line items from Order.json and sort the data alphabetically by the city.

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:

Answer Area

SELECT li.id AS lineitemid, li.price

FROM [Orders o / Lineltems li]

JOIN [li / o] IN [o.line_items / li.line_items / o.address]

ORDER BY [o.address.city / li.address.city / o.city / li.city] ASC

Correct Answer:

Answer Area

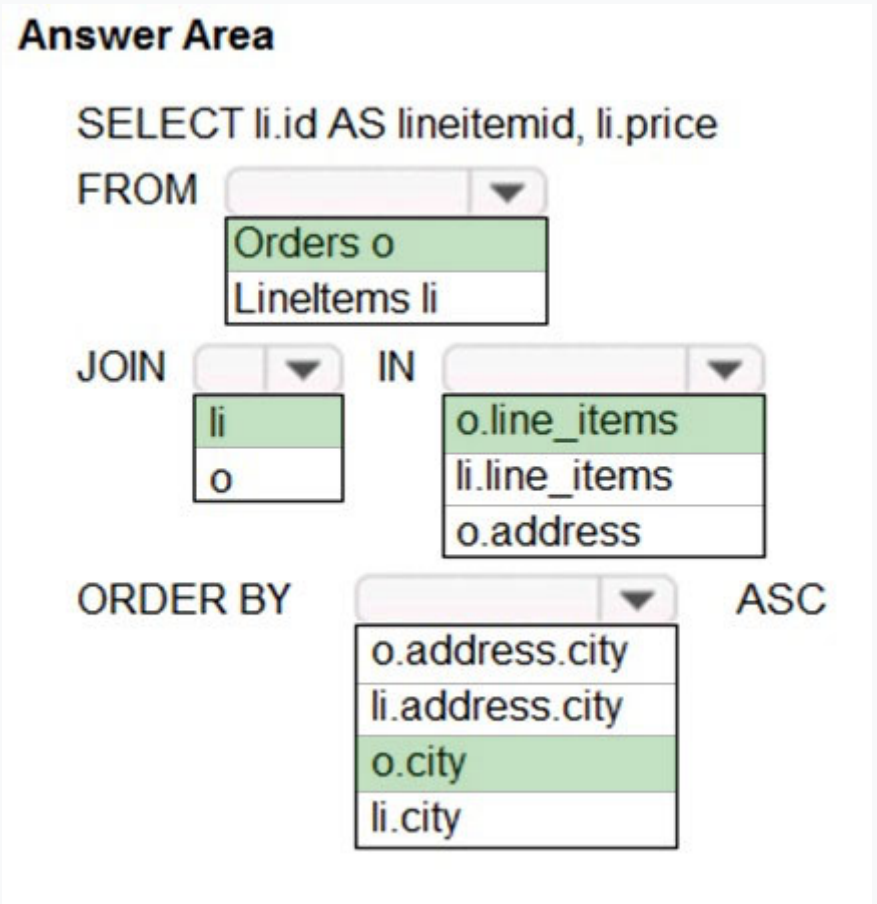

Box 1: orders o -

Scenario: Order data is stored as nonrelational JSON and must be queried using SQL.

Box 2:li -

Box 3: o.line_items -

Box 4: o.city -

The city field is in Order, not in the 2s.

Droplex Highly Voted 2 years, 1 month ago

goodluck!

upvoted 71 times

cool_tool Highly Voted 2 years, 2 months ago

The last one should be o.address.city

upvoted 23 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000.

upvoted 8 times

Alex42 2 years, 1 month ago

address is not joined, so it has to be o.city

upvoted 3 times

SaintBahamut 1 year, 6 months ago

joining is needed because we want to list line items from all orders as separate rows (it will give us order x item line of rows), adress is is just a simple object from order, that will be included for all line items due join we made

upvoted 2 times

mariodarken 1 year, 9 months ago

In CosmosDB SQL, if the address is not an array, it seems that you can skip the join https://docs.microsoft.com/es-es/azure/cosmos-db/sql/sql-query-getting-started#query-the-json-items

upvoted 3 times

d0bermannn 2 years, 2 months ago

exactly

upvoted 2 times

70PineApple Most Recent 8 months, 1 week ago

Got this in exam today..20/02/23

upvoted 1 times

jjjita 9 months, 1 week ago

This question is probably part of Case study Coho Winery

upvoted 1 times

tunaparker 1 year, 7 months ago

Now let's review all questions again because there are dozens of them lol

upvoted 5 times

SD5713 1 year, 8 months ago

I passed with only these questions, so don't stress guys. 80% of these questions is exactly the same. Good luck!

upvoted 19 times

GunjanLiebwein 1 year, 8 months ago

I have my exam on 25.02.2022. So let's see.

upvoted 5 times

GunjanLiebwein 1 year, 8 months ago

I Cleared it with 88%. And I got this question too.:-)

upvoted 11 times

altafpatel1984 1 year, 10 months ago

I got this question and json had address node inside main json. So answer would be o. address.city

upvoted 8 times

warchoon 6 months, 3 weeks ago

Thank you. I was sure the question is incomplete

upvoted 1 times

mkqwerty 1 year, 11 months ago

HOPE NOT TO SEE YOU AGAIN

upvoted 20 times

Younes364 1 year, 1 month ago

Me too

upvoted 2 times

tunaparker 1 year, 7 months ago

Hahahahhah lol :D

upvoted 1 times

phvogel 2 years ago

I suspect that the JSON document is incomplete. But if the city were going to be anywhere it would be in the Order, not the line. So the last answer is either o.city or o.address.city. For that we have to guess because there's way to tell if the address information is stored inline in the order or in a class called address. It could be either of:
"street": "828 Broadway",
"city": "New York"
or
"address": {
"street": "828 Broadway",
"city": "New York"
}

upvoted 16 times

altafpatel1984 1 year, 10 months ago

I got this question and json had address node inside main json. So answer would be o. address.city

upvoted 11 times

ks321 3 months, 1 week ago

Thanks Bro!

upvoted 1 times

ning 2 years, 2 months ago

Where is the city field in json???

upvoted 4 times

altafpatel1984 1 year, 10 months ago

I got this question and json had address node inside main json. So answer would be o. address.city

upvoted 7 times

GreenPanda 2 years, 2 months ago

https://docs.microsoft.com/ja-jp/azure/cosmos-db/sql-query-getting-started#query-the-json-items

upvoted 2 times

noip 2 years, 2 months ago

where is the city field in order?

upvoted 7 times

vokep77043 2 years, 2 months ago

Whole orders json is not showed in case study - it seems Orders.json is a single document of another Orders table/json.

upvoted 2 times

altafpatel1984 1 year, 10 months ago

I got this question and json had address node inside main json. So answer would be o. address.city

upvoted 3 times

gesl 1 year, 9 months ago

Are these questions enough to pass the exam?

upvoted 1 times

mariodarken 1 year, 9 months ago

I hope so.. good luck mate

upvoted 2 times

gesl 1 year, 8 months ago

Thank you mariodarken. For all wondering ... Feb./2022... Aprox. 40 out of 54 questions are covered with course. Enough to pass the exam. Good luck

upvoted 1 times

Topic 10 - Testlet 12

Question #1 *Topic 10*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -
The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
HOTSPOT -
You need to implement the Azure Function for delivery driver profile information.
Which configurations should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

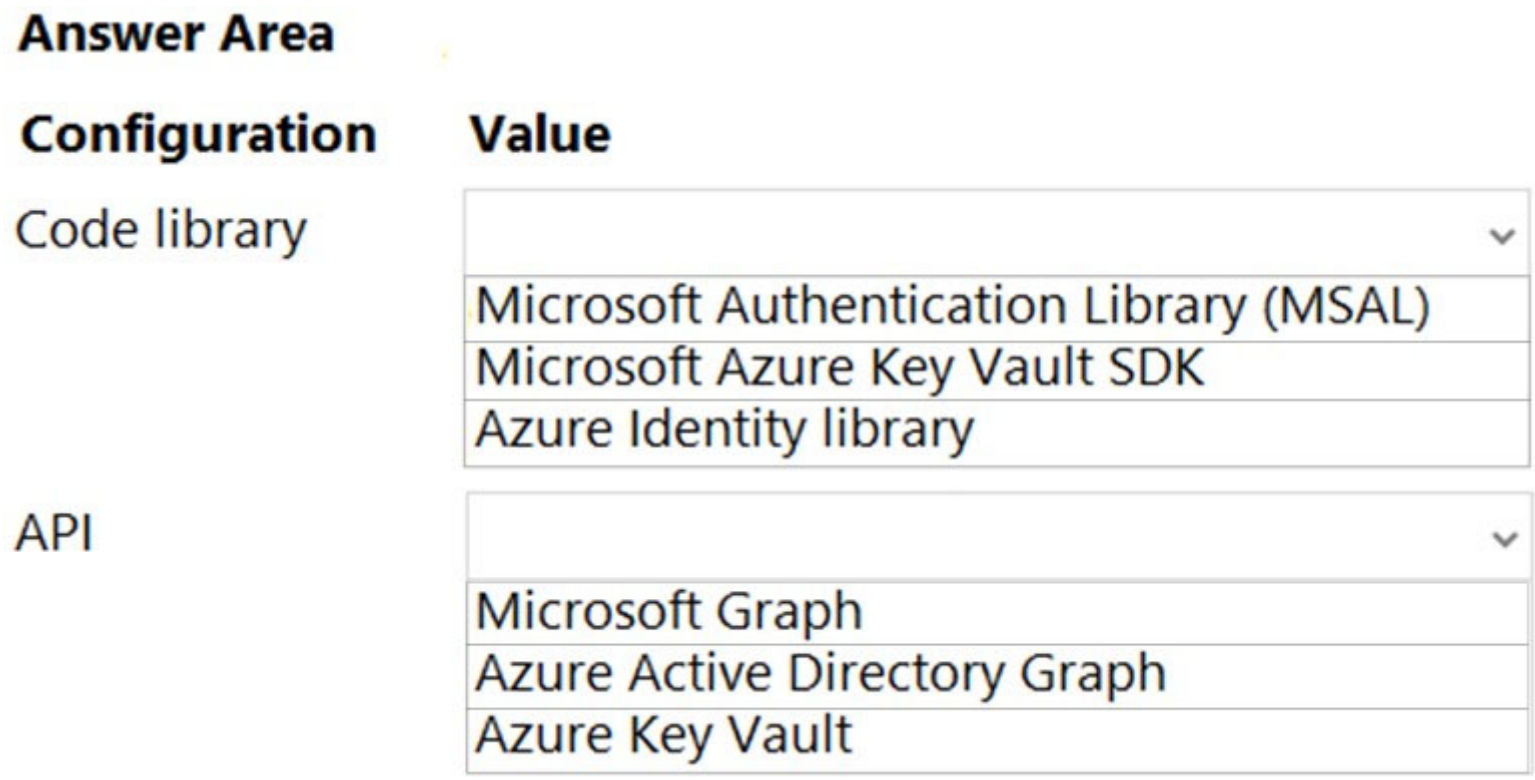

Correct Answer:

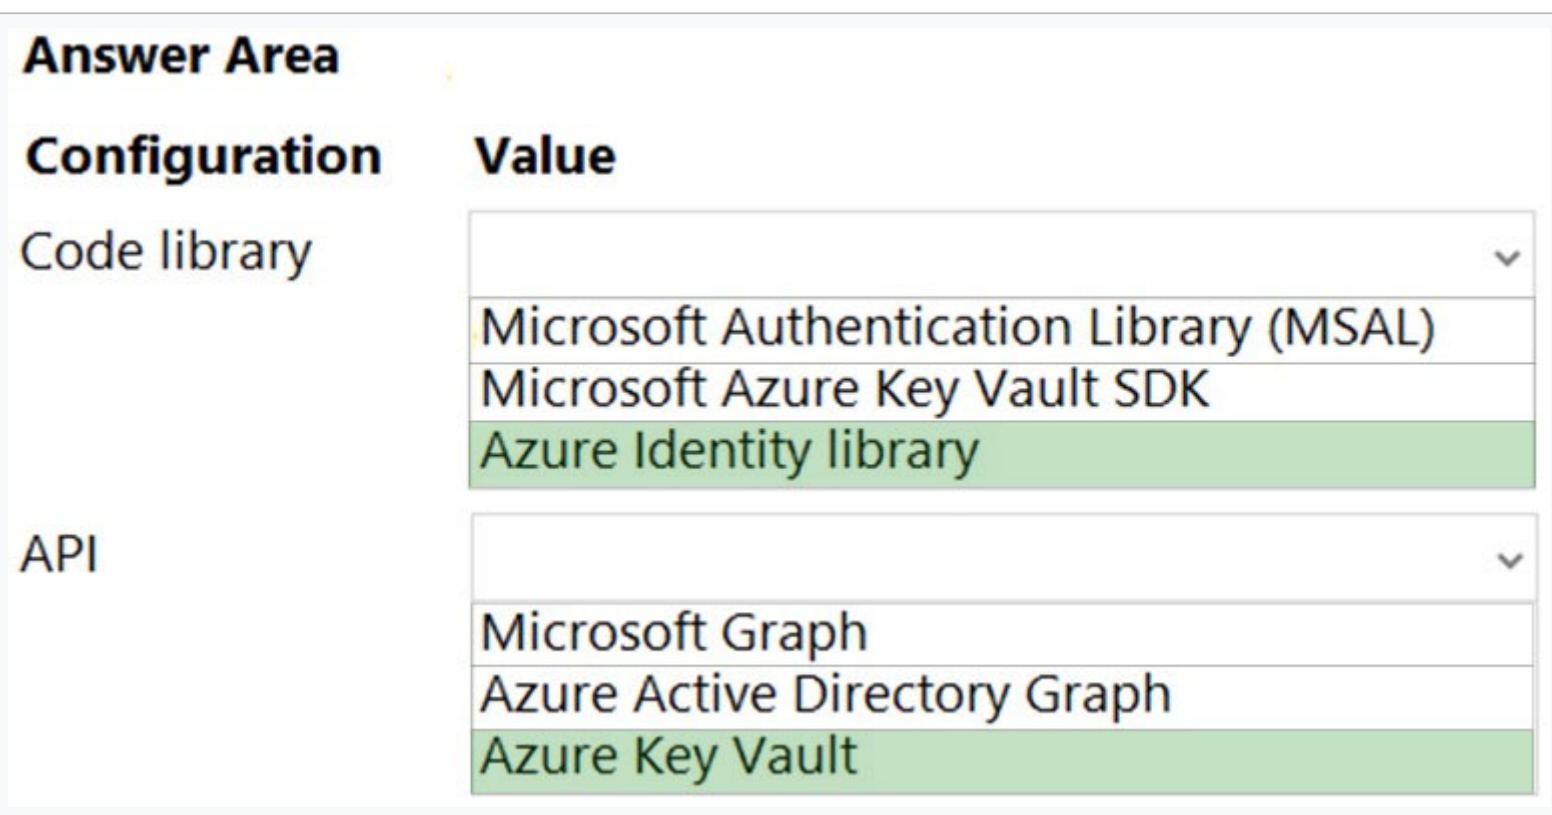

Box 1: Azure Identity library -
Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.
We recommend that you use a managed identity for applications deployed to Azure.
The preceding authentication scenarios are supported by the Azure Identity client library and integrated with Key Vault SDKs.
Note: What is Managed Service Identity?

Azure Key Vault avoids the need to store keys and secrets in application code or source control. However, in order to retrieve keys and secrets from Azure Key
Vault, you need to authorize a user or application with Azure Key Vault, which in its turn needs another credential. Managed Service Identity avoids the need of storing credentials for Azure Key Vault in application or environment settings by creating a Service Principal for each application or cloud service on which
Managed Service Identity is enabled. This Service Principal enables you to call a local MSI endpoint to get an access token from Azure AD using the credentials of the Service Principal. This token is then used to authenticate to an Azure Service, for example Azure Key Vault.

Box 2: Azure Key Vault -
Azure Key Vault allows you to securely access sensitive information from within your applications:
* Keys, secrets, and certificates are protected without your having to write the code yourself, and you can easily use them from your applications.
Use Azure Key Vault to store only secrets for your application. Examples of secrets that should be stored in Key Vault include:

Client application secrets -

Connection strings -

Passwords -

Shared access keys -

SSH keys -
Reference:
https://docs.microsoft.com/en-us/azure/key-vault/general/developers-guide https://integration.team/blog/retrieve-azure-key-vault-secrets-using-azure-functions-and-managed-service-identity

willchenxa Highly Voted 1 year, 1 month ago
Store delivery driver profile information in Azure Active Directory (Azure AD), I would use:
Code Library: MSAL
API: Microsoft Graph
https://docs.microsoft.com/en-us/azure/active-directory/develop/msal-overview
upvoted 52 times

davidkerr7 4 months, 2 weeks ago
correct, profile is stored in AD not keyvault
"Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website."
upvoted 1 times

Dani_ac7 1 year, 1 month ago
i think your answer is correct
upvoted 3 times

coffecold Highly Voted 1 year ago
To prevent reading the cases multiple times:
Please see the spots below where you can find the questions (page/topic/question/subject)
VanArsdel, Ltd.
47 10 1 delivery profile information in code
48 10 2 grant access SAS or AD?
48 10 3 JWT +claims
48 10 4 place to store the RSA-HSM key
51 15 1 reduce read latency Cosmos DB
51 15 2 audit transactions in blob
54 21 1 systems for source-receiver-handler
57 29 1 funcion configuration : binding , trigger, direction
57 29 2 tier of SPA web app
57 30 1 Azure Blob features for point-in-time restore and accidental deletion of data
upvoted 16 times

dancing__donkey 8 months, 1 week ago
Where are these pages? link pls
upvoted 3 times

OPT_001122 10 months, 3 weeks ago
great help!
upvoted 2 times

jaf19f Most Recent 2 months, 2 weeks ago

I got this question (12-Aug-2023) and I chose highly voted answer - 932 passed

upvoted 3 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with highly voted. Scored 917

upvoted 5 times

JH81 4 months ago

Got this question and case study 6/29/2023 and passed with 850. MASL/Graph is correct.

upvoted 1 times

OPT_001122 10 months, 3 weeks ago

MSAL
MS Graph

upvoted 6 times

coffecold 1 year ago

Why...o...why are so many of the revealed answers incorrect? Without the topic discussions you would have hard times...

upvoted 10 times

Ciupaz 2 weeks, 1 day ago

You're right.

upvoted 1 times

Ishan_63 1 year ago

Microsoft Graph API is used in order to get user provide information from Azure AD. Key valult is used to manage Secrets, Certificates and Keys.

upvoted 2 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
You need to grant access to the retail store location data for the inventory service development effort.
What should you use?

A. Azure AD access token

B. Azure RBAC role

C. Shared access signature (SAS) token

D. Azure AD ID token

E. Azure AD refresh token

Correct Answer: *C*
A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:
What resources the client may access.
What permissions they have to those resources.
How long the SAS is valid.
Note: Inventory services:
The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

kampatra Highly Voted 1 year, 1 month ago
Seems Correct as access is required to access for 3 months
upvoted 12 times

davidkerr7 4 months, 2 weeks ago
where ar u getting the 3 months from
its SAS because they only need access to upload data
'Store locations send data every hour to an Azure Blob storage account'
upvoted 1 times

gmishra88 Highly Voted 1 year ago
All that reading to check if the person knows SAS exist with an expiry date?
upvoted 8 times

NombreFalso 8 months, 1 week ago
This is why I hate case studies
upvoted 5 times

jaf19f Most Recent 2 months, 2 weeks ago
I got this question (12-Aug-2023) and I chose C as my answer - 932 passed
upvoted 2 times

juanckar 3 months, 3 weeks ago
This was on the exam (July 2023). Went with highly voted. Scored 917
upvoted 3 times

OPT_001122 10 months, 3 weeks ago

https://www.examtopics.com/exams/microsoft/az-204/view/47/
Topic 10 - Testlet 12
Question #1

It has all the questions listed for VanArsdel casestudy
upvoted 2 times

idolle 4 months, 3 weeks ago
Page not found
upvoted 2 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
HOTSPOT -
You need to reliably identify the delivery driver profile information.
How should you configure the system? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Configuration	Value
JSON web token (JWT) type	ID / Refresh / Access
Payload claim value	oid / aud / idp

Correct Answer:

Answer Area

Configuration	Value
JSON web token (JWT) type	**ID** / Refresh / Access
Payload claim value	**oid** / aud / idp

Box 1: ID -
Scenario: Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate

website.
ID token - A JWT that contains claims that you can use to identify users in your application. This token is securely sent in HTTP requests for communication between two components of the same application or service. You can use the claims in an ID token as you see fit. They're commonly used to display account information or to make access control decisions in an application. ID tokens are signed, but the're not encrypted. When your application or API receives an ID token, it must validate the signature to prove that the token is authentic. Your application or API must also validate a few claims in the token to prove that it's valid.
Depending on the scenario requirements, the claims validated by an application can vary, but your application must perform some common claim validations in every scenario.

Box 2: Oid -
Oid - The immutable identifier for the "principal" of the request - the user or service principal whose identity has been verified. In ID tokens and app+user tokens, this is the object ID of the user. In app-only tokens, this is the object ID of the calling service principal. It can also be used to perform authorization checks safely and as a key in database tables. This ID uniquely identifies the principal across applications - two different applications signing in the same user will receive the same value in the oid claim.
Incorrect:
Aud - Identifies the intended recipient of the token. For Azure AD B2C, the audience is the application ID. Your application should validate this value and reject the token if it doesn't match. Audience is synonymous with resource.
Idp - Records the identity provider that authenticated the subject of the token. This value is identical to the value of the Issuer claim unless the user account not in the same tenant as the issuer - guests, for instance. If the claim isn't present, it means that the value of iss can be used instead. For personal accounts being used in an organizational context (for instance, a personal account invited to an Azure AD tenant), the idp claim may be 'live.com' or an STS URI containing the
Microsoft account tenant.
Reference:
https://docs.microsoft.com/en-us/azure/active-directory-b2c/tokens-overview https://docs.microsoft.com/en-us/azure/active-directory/develop/access-tokens

AbdulMannan Highly Voted 1 year, 1 month ago
Got this question on 30-Sep-2022 exam.
Answer is correct. Passed with 870 score.
upvoted 11 times

ArturKon Highly Voted 1 year, 1 month ago
Looks correct.
Ref.: https://learn.microsoft.com/en-us/azure/active-directory/develop/id-tokens
upvoted 6 times

jaf19f Most Recent 2 months, 2 weeks ago
I got this question (12-Aug-2023) and I chose highly voted answer - 932 passed
upvoted 3 times

BaoNguyen2411 3 months ago
got this question on 29/02/2023
upvoted 1 times

juanckar 3 months, 3 weeks ago
This was on the exam (July 2023). Went with highly voted. Scored 917
upvoted 3 times

winterthor4 7 months, 1 week ago
Look correct. Got this question on 26-Mar-2023 exam. Go with suggested answer got 890 score.
upvoted 5 times

adilkhan 6 months, 1 week ago
all questions were from exam topics?
upvoted 2 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
You need to secure the Azure Functions to meet the security requirements.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Store the RSA-HSM key in Azure Key Vault with soft-delete and purge-protection features enabled.

B. Store the RSA-HSM key in Azure Blob storage with an immutability policy applied to the container.

C. Create a free tier Azure App Configuration instance with a new Azure AD service principal.

D. Create a standard tier Azure App Configuration instance with an assigned Azure AD managed identity.

E. Store the RSA-HSM key in Azure Cosmos DB. Apply the built-in policies for customer-managed keys and allowed locations.

Correct Answer: *AD*
Scenario: All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Microsoft Azure Key Vault is a cloud-hosted management service that allows users to encrypt keys and small secrets by using keys that are protected by hardware security modules (HSMs).
You need to create a managed identity for your application.
Reference:
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

Community vote distribution
AD (100%)

gmishra88 Highly Voted 1 year ago
There is no need to read all that text to identify the answer.
upvoted 16 times

sutozoli 12 months ago
Which sentence helped you to identify the answer?
upvoted 2 times

hubekpeter 11 months, 1 week ago
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.
upvoted 5 times

AbdulMannan Highly Voted 1 year, 1 month ago
Got this question on 30-Sep-2022 exam.
Answer is correct. Passed with 870 score.
upvoted 9 times

Frazerdon Most Recent 1 month, 2 weeks ago
Selected Answer: AD
Correct!!!
upvoted 1 times

jaf19f 2 months, 2 weeks ago
I got this question (12-Aug-2023) and I chose A and D as my answer - 932 passed

upvoted 3 times

ReyPirata 2 months, 1 week ago

Correct. This was on the exam (08/20/2023). Scored 925

upvoted 1 times

BaoNguyen2411 3 months ago

got this question on 29/06/2023

upvoted 2 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with highly voted. Scored 917

upvoted 3 times

beebert69 9 months ago

When we have started answering questions on the Case Study can we go back to the Case Study documents?

upvoted 1 times

mmdex 8 months, 1 week ago

All the Case Study information is available to you as you are answering the questions. You can go back and forth between questions and the Case Study information as you like. But after you answer all questions and finish the Case Study, you cannot go back. It's like a mini-exam within the main exam.

upvoted 8 times

kampatra 1 year, 1 month ago

Correct!

upvoted 3 times

sghaha 1 year, 6 months ago

in Korean
https://docs.microsoft.com/ko-kr/azure/app-service/app-service-key-vault-references

upvoted 5 times

Topic 11 - Testlet 13

Question #1 *Topic 11*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -
You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14         "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

DRAG DROP -

You need to add markup at line AM04 to implement the ContentReview role.

How should you complete the markup? To answer, drag the appropriate json segments to the correct locations. Each json segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

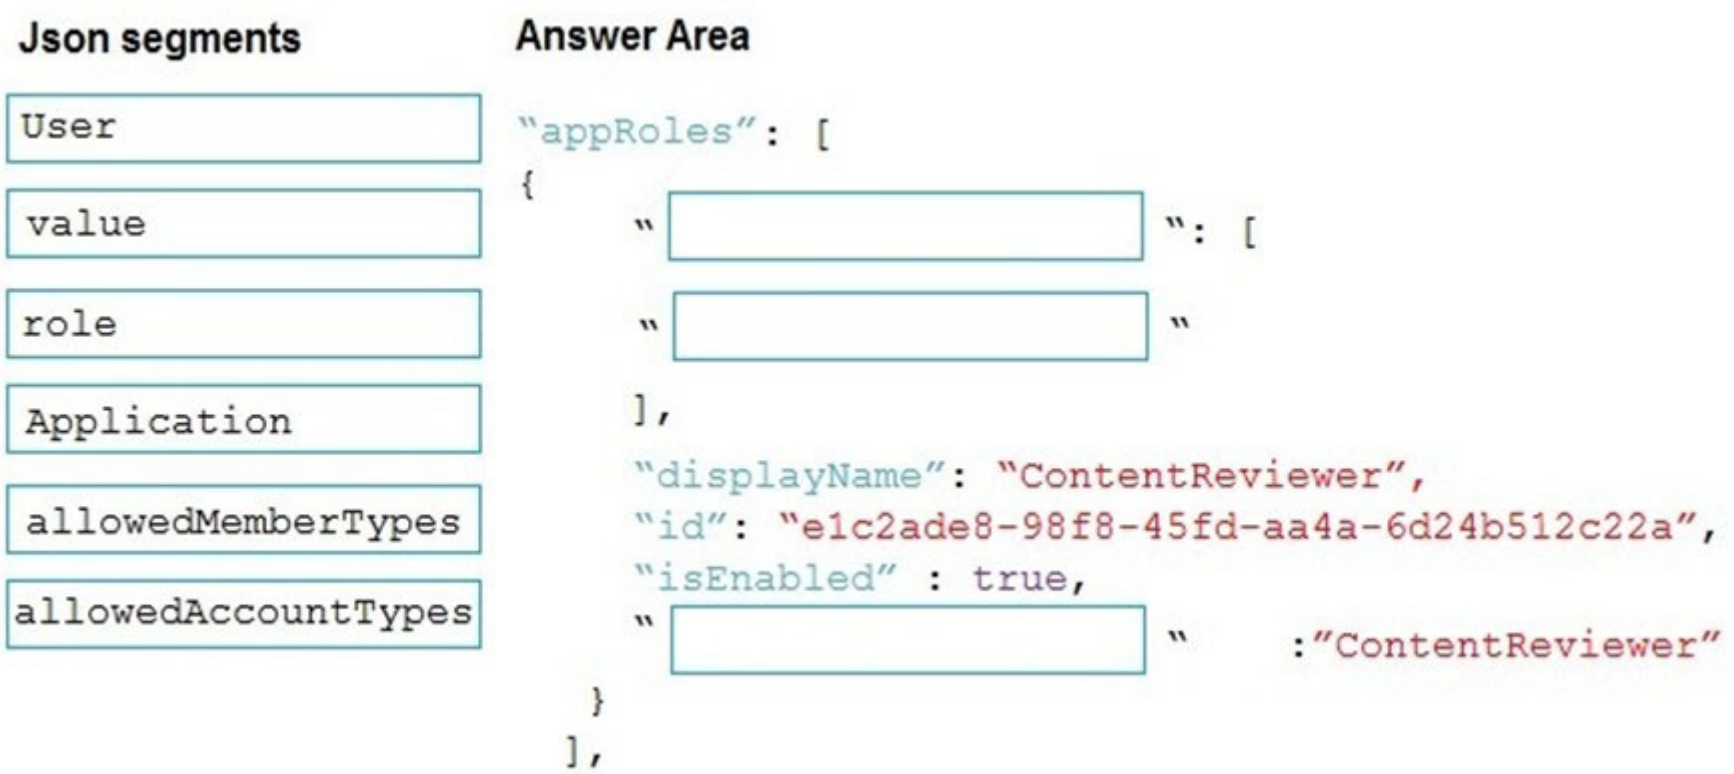

Correct Answer:

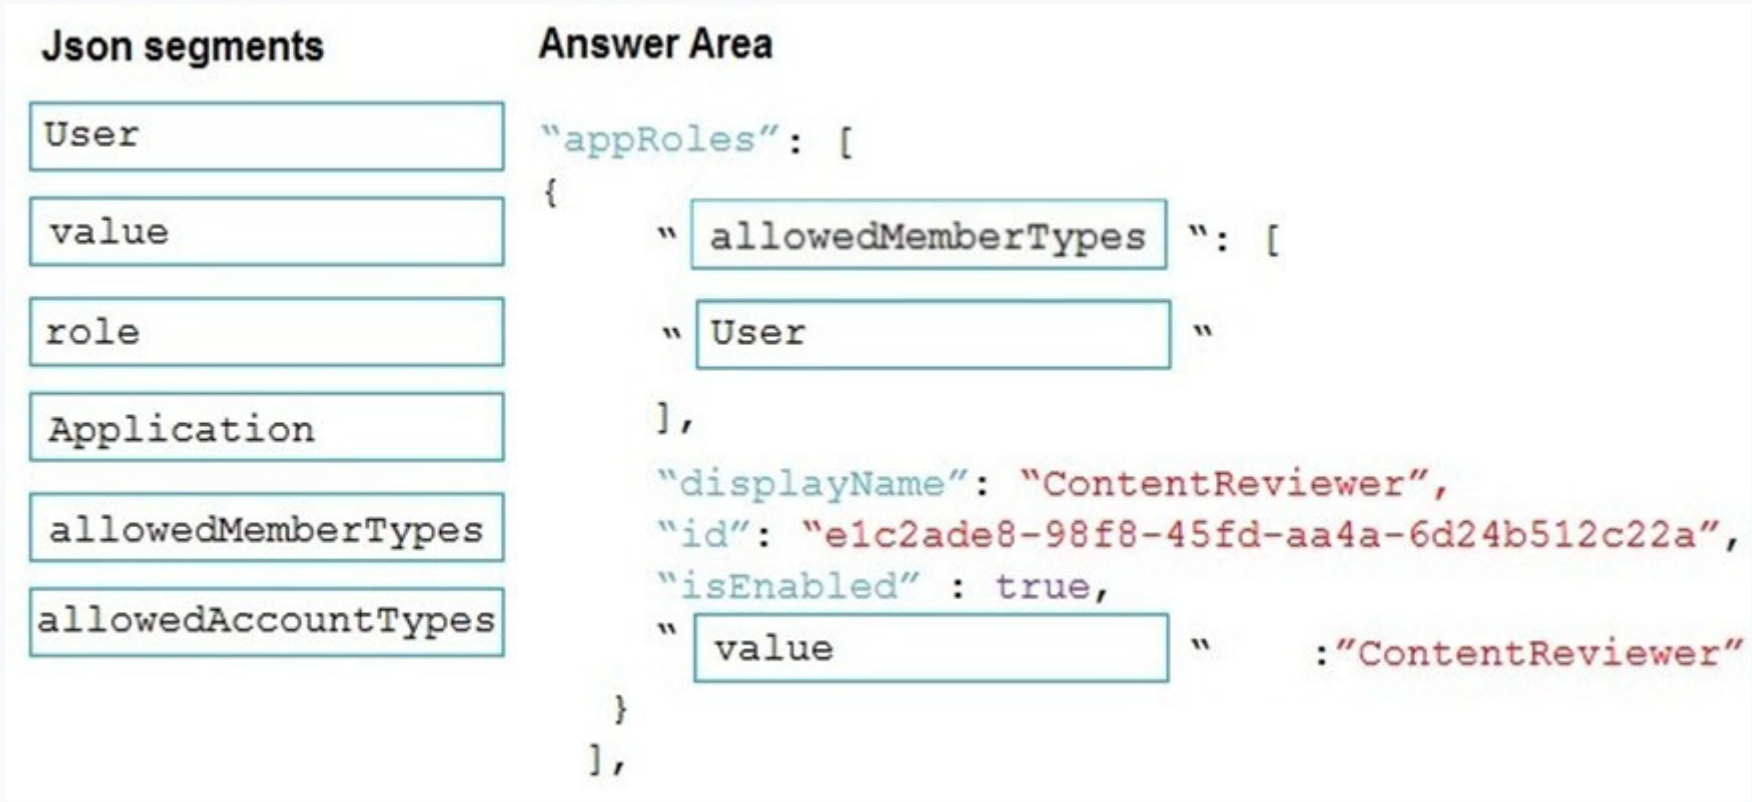

Box 1: allowedMemberTypes -

allowedMemberTypes specifies whether this app role definition can be assigned to users and groups by setting to "User", or to other applications (that are accessing this application in daemon service scenarios) by setting to "Application", or to both.

Note: The following example shows the appRoles that you can assign to users.

"appId": "8763f1c4-f988-489c-a51e-158e9ef97d6a",
"appRoles": [
{
"allowedMemberTypes": [
"User"
],
"displayName": "Writer",
"id": "d1c2ade8-98f8-45fd-aa4a-6d06b947c66f",
"isEnabled": true,
"description": "Writers Have the ability to create tasks.",
"value": "Writer"
}
],
"availableToOtherTenants": false,

Box 2: User -

Scenario: In order to review content a user must be part of a ContentReviewer role.

Box 3: value -

value specifies the value which will be included in the roles claim in authentication and access tokens.

Reference:

https://docs.microsoft.com/en-us/graph/api/resources/approle

ray01 Highly Voted 2 years, 6 months ago

Correct. See MS example here:
https://docs.microsoft.com/de-de/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps#example-user-app-role

upvoted 33 times

coffecold Highly Voted 1 year ago

To prevent reading the cases multiple times:
Please see the spots below where you can find the questions (page/topic/question/subject)
Contoso, Ltd
48 11 1 configure ApplicationManifest
48 11 2 configure ApplicationManifest
49 11 3 type of SSL certificate / proxy used
49 11 4 YAML @Azure Container Instances for mounting
49 11 5 ApplicationManifest optional claims
52 16 1 cli : az monitor metrics alert create
52 16 2 cli : az for http server log output
56 27 1 code : read properties from GRID event
56 27 2 tier of function plan @Vnet

58 31 1 choose best storage solution
58 31 2 code : bindings for function
58 31 3 docker file : Private/Public Windows/Linux
upvoted 13 times

OPT_001122 10 months, 3 weeks ago
Great Help!!!!
upvoted 1 times

CriJim Most Recent 1 month ago
Here below find same case studies topics:
Wide World Importers 7, 12, 23
Litware Inc. 8, 14, 20
Coho Winery produces 9, 22, 28
VanArsdel, Ltd. 10, 15, 21, 29, 30
Contoso, Ltd. 11, 16, 27, 31
City Power & Light 13, 17, 18, 24, 32
Proseware, Inc. 19, 25, 26
upvoted 2 times

applepie 3 months ago
got this case study today, went with provided answer - 7/30/2023, score 895/1000
upvoted 2 times

BaoNguyen2411 3 months ago
got this question on 29/06/2023
upvoted 1 times

BaoNguyen2411 3 months ago
Got this question on 29/06/2023
upvoted 2 times

NightshadeRC 3 months, 1 week ago
Had this question today 2023-07-26
upvoted 2 times

CODE_STS 8 months ago
Got this in the exam today! Feb 28, 2023
upvoted 3 times

AlexeyG 8 months, 2 weeks ago
Got this in 16/02/2023
upvoted 1 times

mabdo 8 months, 2 weeks ago
did you get 2 different case studies? I saw you commenting on another one also. Is that possible?
upvoted 1 times

Net_IT 8 months, 2 weeks ago
Thank you for mentioning the date, would be nice to know your answer and whether you passed as well :)
upvoted 1 times

nvtienanh 11 months ago
On exam December 2, 2022
upvoted 4 times

TheExamMaster2020 11 months, 2 weeks ago
Did my exam on 15th November 2022. This test case and question was on it.
upvoted 3 times

coffecold 1 year ago
Usage see here.
https://stackoverflow.com/questions/57400122/aad-approles-allowedmembertypes-for-both-application-and-user
upvoted 1 times

gmishra88 1 year ago
So, a Microsoft developer is not allowed to refer to documents and has to just memorize the syntax also of a manifest file ! This is a lovely certification exam. Never one more Microsoft certification
upvoted 5 times

petitbilly 1 year, 7 months ago
Got this one 03/2022
upvoted 4 times

edengoforit 1 year, 9 months ago

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using
React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

upvoted 1 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 5 times

mlantonis 2 years, 5 months ago

Correct. Check here: https://docs.microsoft.com/de-de/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps#example-user-app-role

upvoted 11 times

Question #2 *Topic 11*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -

You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14         "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

HOTSPOT -

You need to add code at line AM09 to ensure that users can review content using ContentAnalysisService.

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:

Answer Area

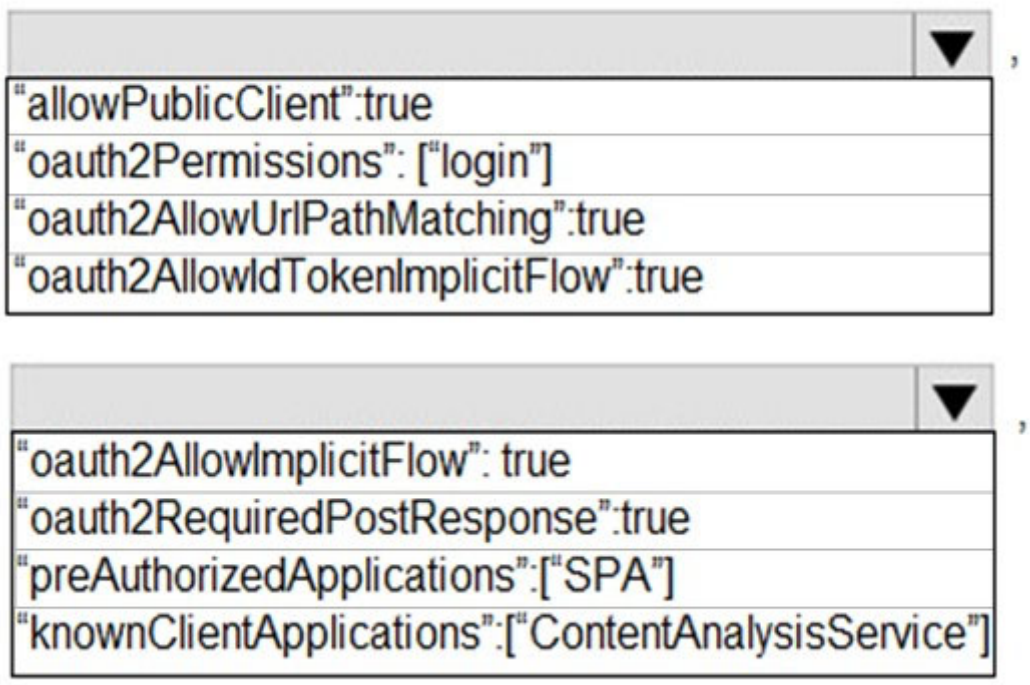

Correct Answer:

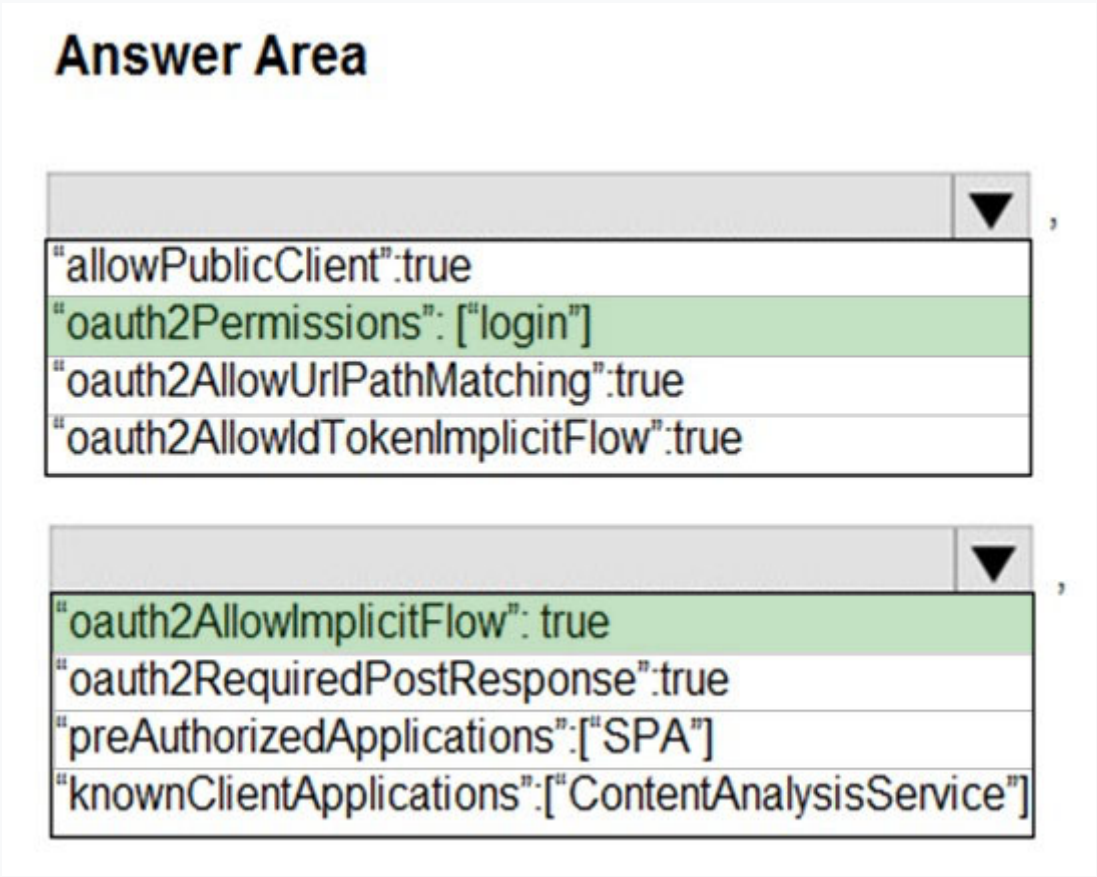

Box 1: "oauth2Permissions": ["login"]

oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

Box 2: "oauth2AllowImplicitFlow":true

For applications (Angular, Ember.js, React.js, and so on), Microsoft identity platform supports the OAuth 2.0 Implicit Grant flow.

Reference:

https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest

Zidimirite Highly Voted 2 years, 6 months ago

https://www.examtopics.com/discussions/microsoft/view/25371-exam-az-204-topic-8-question-2-discussion/
oauth2AllowImplicitFlow = true,
oauth2AllowIdTokenImplicitFlow = true

upvoted 41 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: "oauth2AllowIdTokenImplicitFlow":true
This value indicates whether the web app can request ID tokens of the implicit OAuth 2.0 flow. The default setting is "false".

Box 2: "oauth2AllowImplicitFlow":true
This value indicates whether the web app can request OAuth 2.0 implicit flow access tokens. The default setting is "false".

Reference:

https://docs.microsoft.com/de-de/azure/active-directory/develop/reference-app-manifest#oauth2allowidtokenimplicitflow-attribute

https://docs.microsoft.com/de-de/azure/active-directory/develop/reference-app-manifest#oauth2allowimplicitflow-attribute

upvoted 34 times

edengoforit Most Recent 1 year, 9 months ago

Since it is a browser app (React), Implicit Grant Flow, should be enabled:
oauth2AllowImplicitFlow = true,
oauth2AllowIdTokenImplicitFlow = true

"To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using
React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the "reviewer's email address for auditing purposes.

This is a web app so we need oauth2AllowImplicitFlow.
We need to have information about the user (reviewer email), so we need to have Id Token thus oauth2AllowIdTokenImplicitFlow must be set to true

upvoted 6 times

RajMasilamani 2 years, 1 month ago

I think the answer should be
Oauth2permission login for first box as API to be exposed and email credentials to be saved for auditing purpose.

upvoted 1 times

[Removed] 2 years, 5 months ago

"allowPublicClient": true
> Not required as infered from replyUrlsWithType

"oauth2Permissions": ["login"]
> only for resource server exposing an API

"oauth2AllowUrlPathMatching: true
> Couldn't find in the docu

"oauth2AllowIdTokenImplicitFlow": true
> Correct one to choose!

"oauth2AllowImplicitFlow": true
> Correct one to choose!

"oauth2RequiredPostResponse": true
> Not relevant

"preAuthorizedApplications": ["SPA"]
> Not relevant

"knownClientApplications": ["ContentAnalysisService"]
> Not relevant

https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest?WT.mc_id=Portal-Microsoft_AAD_RegisteredApps#manifest-reference

upvoted 7 times

wtkwsk 2 years, 6 months ago

Zidimirite is correct: oauth2AllowImplicitFlow & oauth2AllowIdTokenImplicitFlow
See here: https://docs.microsoft.com/de-de/azure/active-directory/develop/reference-app-manifest

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

"oauth2AllowIdTokenImplicitFlow":true
"oauth2AllowImplicitFlow":true

upvoted 3 times

trance13 2 years, 6 months ago

I think the first answer should be allowPublicClient: true because this is SPA application and be default is is set to false. implicit flow does not make sense with confidential clients, it would be client credentials flow.

upvoted 4 times

Question #3 *Topic 11*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -

You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14         "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

HOTSPOT -

You need to ensure that network security policies are met.

How should you configure network security? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

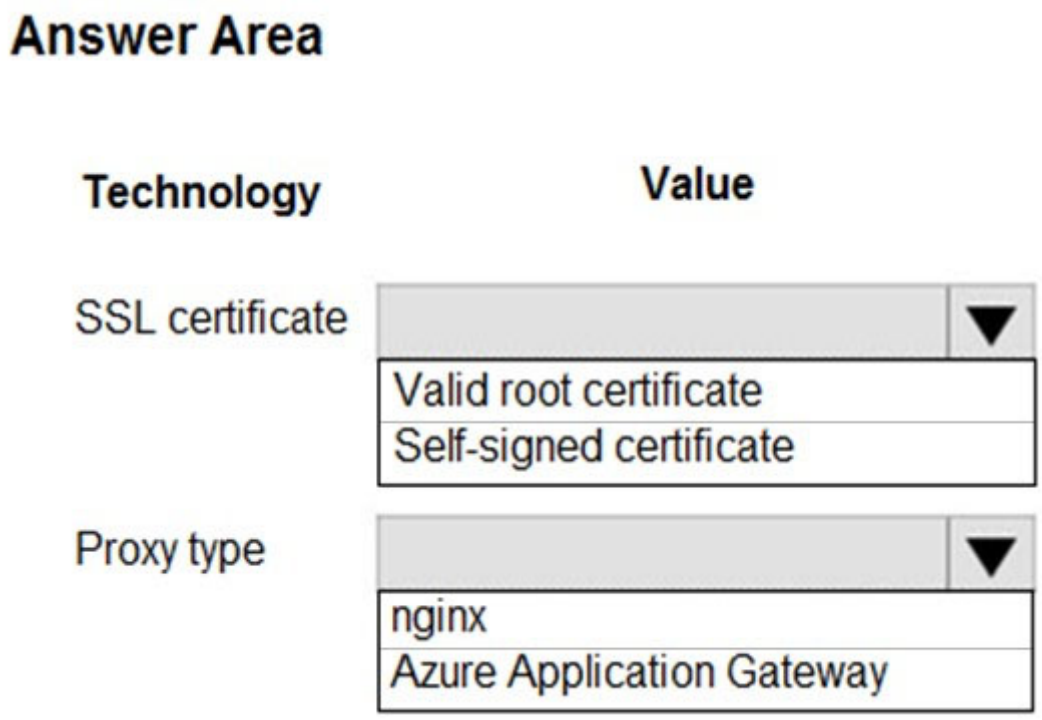

Correct Answer:

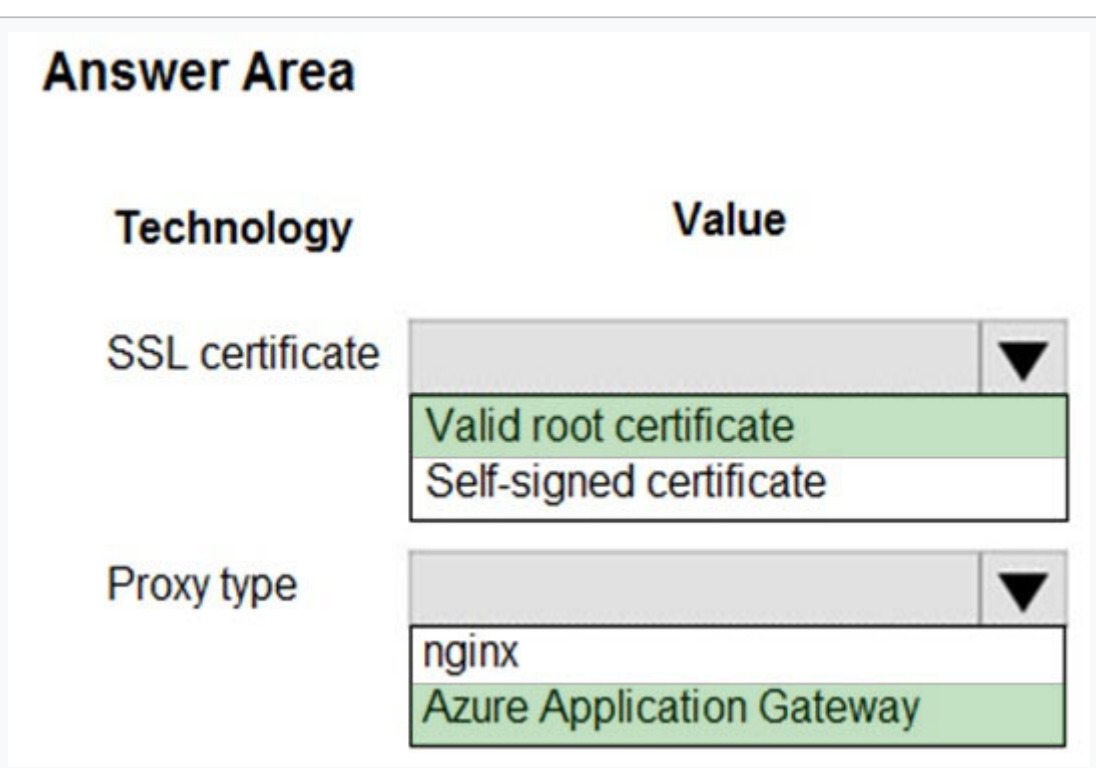

Box 1: Valid root certificate -

Scenario: All websites and services must use SSL from a valid root certificate authority.

Box 2: Azure Application Gateway

Scenario:

☞ Any web service accessible over the Internet must be protected from cross site scripting attacks.
☞ All Internal services must only be accessible from Internal Virtual Networks (VNets)
All parts of the system must support inbound and outbound traffic restrictions.
▪
Azure Web Application Firewall (WAF) on Azure Application Gateway provides centralized protection of your web applications from common exploits and vulnerabilities. Web applications are increasingly targeted by malicious attacks that exploit commonly known vulnerabilities. SQL injection and cross-site scripting are among the most common attacks.
Application Gateway supports autoscaling, SSL offloading, and end-to-end SSL, a web application firewall (WAF), cookie-based session affinity, URL path-based routing, multisite hosting, redirection, rewrite HTTP headers and other features.
Note: Both Nginx and Azure Application Gateway act as a reverse proxy with Layer 7 load-balancing features plus a WAF to ensure strong protection against common web vulnerabilities and exploits.
You can modify Nginx web server configuration/SSL for X-XSS protection. This helps to prevent cross-site scripting exploits by forcing the injection of HTTP headers with X-XSS protection.
Reference:
https://docs.microsoft.com/en-us/azure/web-application-firewall/ag/ag-overview https://www.upguard.com/articles/10-tips-for-securing-your-nginx-deployment

mlantonis Highly Voted 2 years, 5 months ago

Box 1: Valid root certificate
Scenario: All websites and services must use SSL from a valid root certificate authority.

Box 2: Azure Application Gateway
Scenario:
- Any web service accessible over the Internet must be protected from cross site scripting attacks.
- All Internal services must only be accessible from Internal Virtual Networks (VNets).

upvoted 49 times

ExamDev 4 months, 1 week ago

Yeah!
Box 1: Answer is written inside security tab: "must use SSL from a valid root certificate authority" Must read carefully sometime to get the answer :)
Box 2:
nginx as service is not provided by Microsoft azure

upvoted 4 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose：
Valid root certificate
Azure Application Gateway

upvoted 4 times

AlexeyG Highly Voted 8 months, 2 weeks ago

Got this in 16/02/2023

upvoted 6 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 1 times

BaoNguyen2411 3 months ago

Got this question on 29/06/2023

upvoted 1 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 1 times

nvtienanh 11 months ago

On exam December 2, 2022

upvoted 3 times

coffecold 1 year ago

The only thing is : Azure Application Gateway can only be deployed in one region. Don't know if that violates the requirements "All services must run in multiple regions. The failure of any service in a region must not impact overall application availability."
Couldn't find if nginx has the same limitation either.
Would go for Azure Application Gateway

upvoted 1 times

angrybird2007 1 year, 3 months ago

May I know what is the different between Azure Application Gateway Vs NGINX. They are similar. Why we choice Azure Application Gateway instead of NGINX?

upvoted 1 times

gmishra88 1 year ago

Because this is Microsoft exam, just select Application gateway. But I get what you say

upvoted 1 times

Knightie 1 year, 2 months ago

public facing, use an official service to tank the public with Azure Application Gateway, internal calls, just casually use any nginx will do.

upvoted 1 times

Eltooth 1 year, 4 months ago

Valid root cert
Azure App gateway

upvoted 3 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer provided

upvoted 1 times

petitbilly 1 year, 7 months ago

Got this one 03/2022

upvoted 3 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -

You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14         "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

DRAG DROP -

You need to add YAML markup at line CS17 to ensure that the ContentUploadService can access Azure Storage access keys.

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

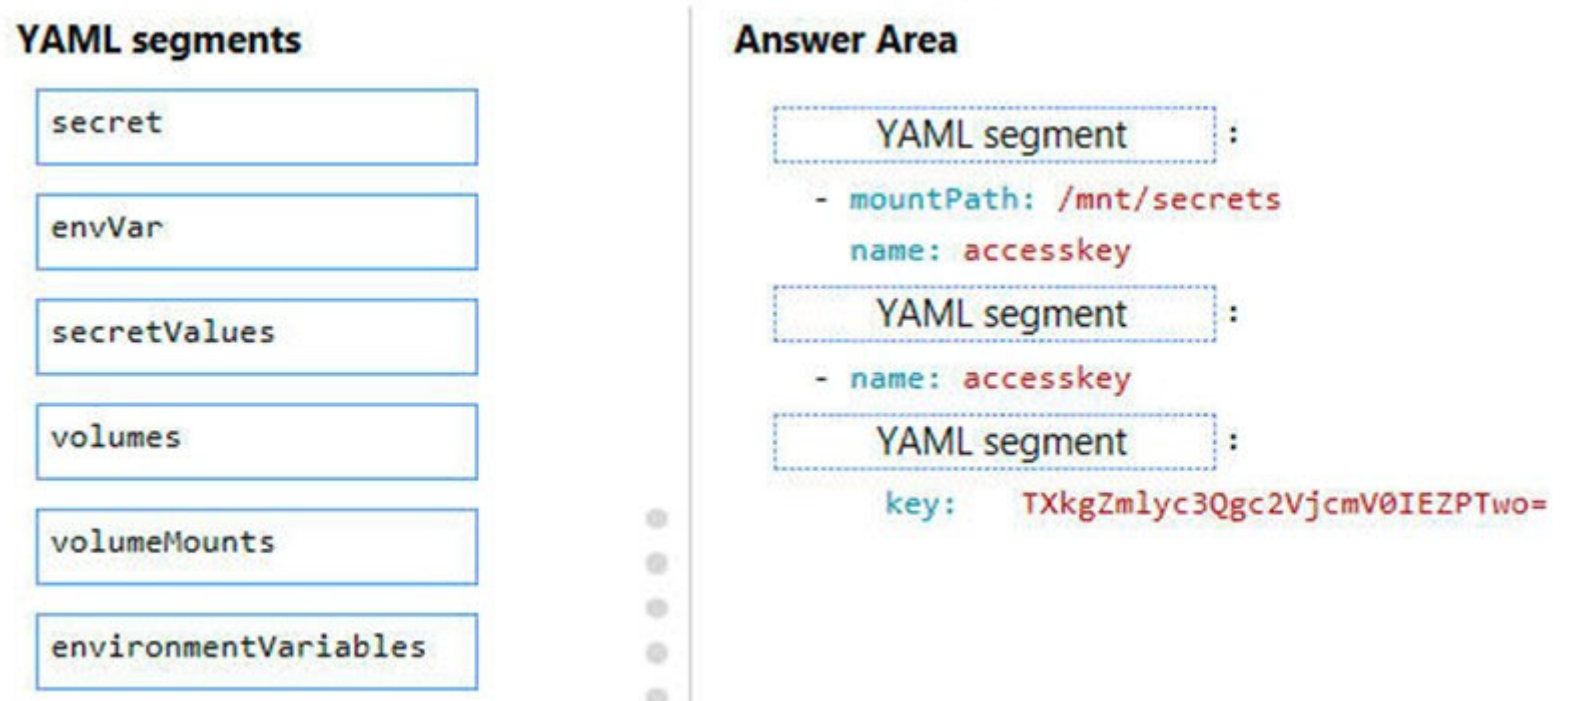

Correct Answer:

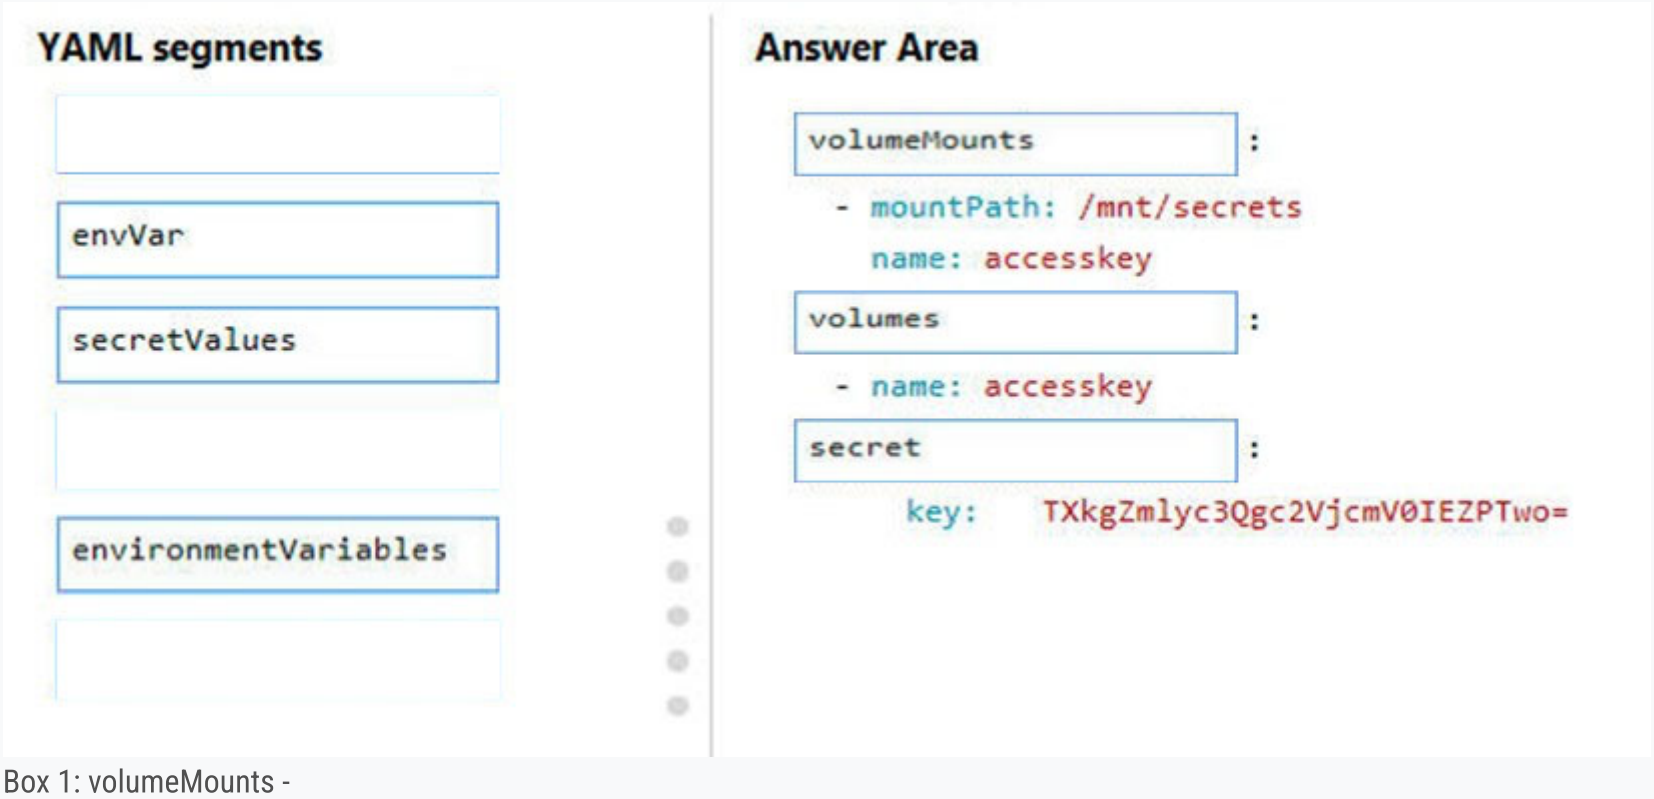

Box 1: volumeMounts -

Example:

volumeMounts:
- mountPath: /mnt/secrets
name: secretvolume1
volumes:
- name: secretvolume1
secret:
mysecret1: TXkgZmlyc3Qgc2VjcmV0IEZPTwo=

Box 2: volumes -

Box 3: secret -
Reference:
https://docs.microsoft.com/en-us/azure/container-instances/container-instances-volume-secret

mlantonis Highly Voted 2 years, 5 months ago

Box 1: volumeMounts

Box 2: volumes

Box 3: secret

Reference:
https://docs.microsoft.com/en-us/azure/container-instances/container-instances-volume-secret

upvoted 28 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose：
volumeMounts
volumes
secert

upvoted 4 times

rdemontis Highly Voted 2 years, 7 months ago

Answers are correct!

upvoted 28 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 2 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 1 times

AlexeyG 8 months, 2 weeks ago

Got this in 16/02/2023

upvoted 3 times

adilkhan 9 months ago

correct! https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-secret#mount-secret-volume---yaml

upvoted 1 times

OPT_001122 10 months, 2 weeks ago

1: volumeMounts
2: volumes
3: secret

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose as below:
1: volumeMounts
2: volumes
3: secret

upvoted 2 times

petitbilly 1 year, 7 months ago

Got this one 03/2022

upvoted 2 times

BrettusMaximus 2 years, 6 months ago

Correct. But a trick question.
accesskey is the name of the volume

upvoted 14 times

coffecold 11 months, 4 weeks ago

Another trick thing:
-volumeMounts belongs to a particular container (of many containers)
-volumes contains a volume that can be shared among all containers
https://learn.microsoft.com/en-us/azure/container-instances/container-instances-reference-yaml#schema

upvoted 2 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -

You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14         "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

HOTSPOT -

You need to add code at line AM10 of the application manifest to ensure that the requirement for manually reviewing content can be met.

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:

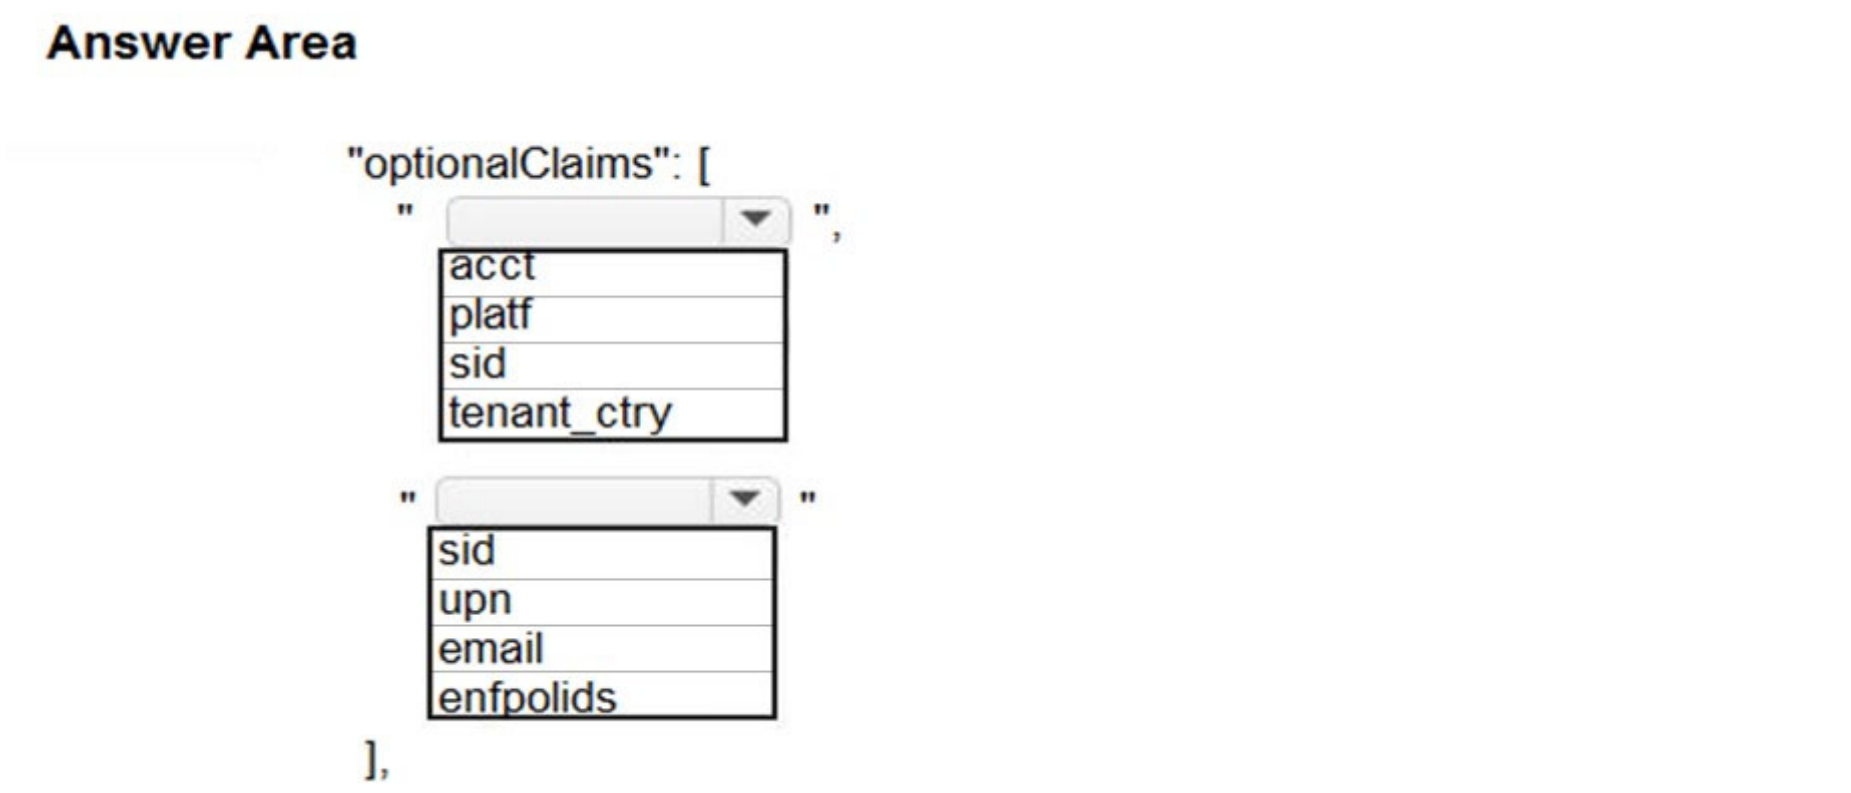

Correct Answer:

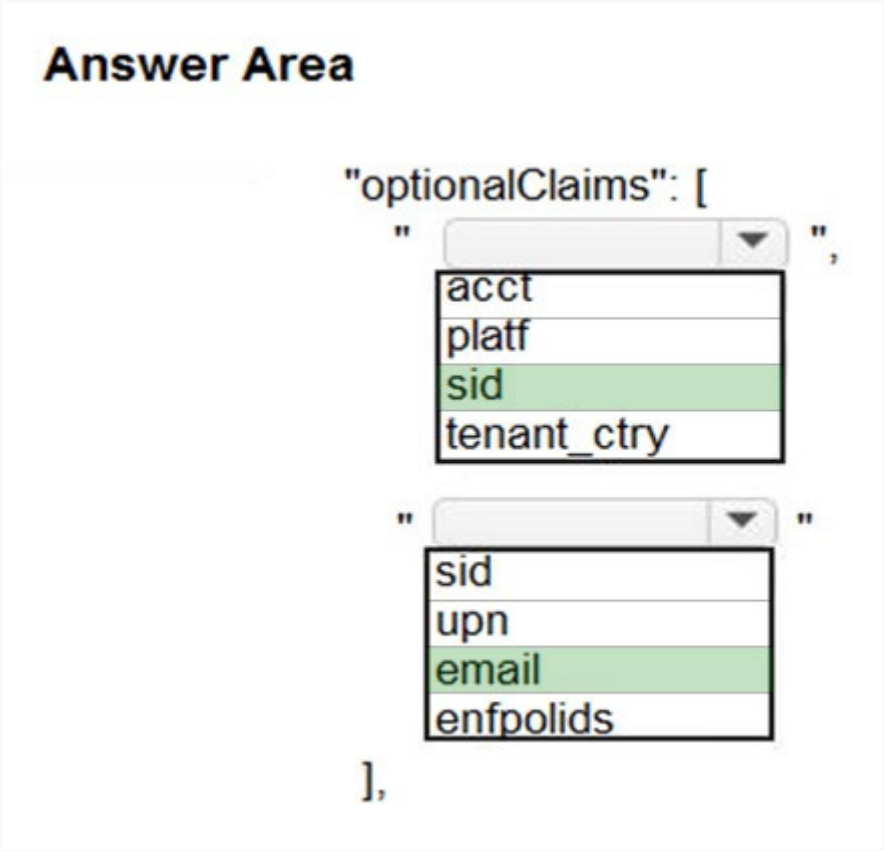

Box 1: sid -

Sid: Session ID, used for per-session user sign-out. Personal and Azure AD accounts.

Scenario: Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role.

Box 2: email -

Scenario: All completed reviews must include the reviewer's email address for auditing purposes.

mlantonis Highly Voted 2 years, 5 months ago

Box 1: sid
Sid: Session ID, used for per-session user sign-out. Personal and Azure AD accounts.
Scenario:
To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role.

Box 2: email
Scenario: All completed reviews must include the reviewer's email address for auditing purposes.

Reference:
https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest
https://docs.microsoft.com/en-us/azure/active-directory/develop/active-directory-optional-claims

upvoted 31 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose:
sid
email

upvoted 3 times

Tom87 Highly Voted 2 years, 7 months ago

The answer is correct. References:
https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest
https://docs.microsoft.com/en-us/azure/active-directory/develop/active-directory-optional-claims

upvoted 9 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 1 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 1 times

surprise0011 6 months, 2 weeks ago

I don't think sid would be necessary but there is no better option here

upvoted 1 times

AlexeyG 8 months, 2 weeks ago

Got this in 16/02/2023

upvoted 2 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This test case and question was on it.

upvoted 2 times

coffecold 1 year ago

Am I missing something here? What can you make up out of a sessionID? Not that the is part of a ContentReviewer role.
I think the users account status in tenant , acct (0= tenant, 1= guest) comes pretty close. If one is user in the tenant he/she could have a ContentReviewer role
(although that is not said).

upvoted 1 times

Eltooth 1 year, 4 months ago

Sid
Email

upvoted 2 times

petitbilly 1 year, 7 months ago

Got this one 03/2022

upvoted 2 times

jose 2 years, 6 months ago

Why sid?

upvoted 1 times

joanbdm 2 years, 5 months ago

"To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials."

sid = Personal and Azure AD accounts Session ID

https://docs.microsoft.com/en-us/azure/active-directory/develop/active-directory-optional-claims

upvoted 7 times

Topic 12 - Testlet 14

Question #1 *Topic 12*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment -

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk Server 2016. The VM runs the following workflows:

Ocean Transport `" This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.

Inland Transport `" This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

Container API `" This API provides container information including weight, contents, and other attributes.

Location API `" This API provides location information regarding shipping ports of call and trucking stops.

Shipping REST API `" This API provides shipping information for use and display on the shipping website.

Shipping Data -

The application uses MongoDB JSON document storage database for all container and transport information.

Shipping Web Site -

The site displays shipping container tracking information and container contents. The site is located at http://shipping.wideworldimporters.com/

Proposed solution -

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:

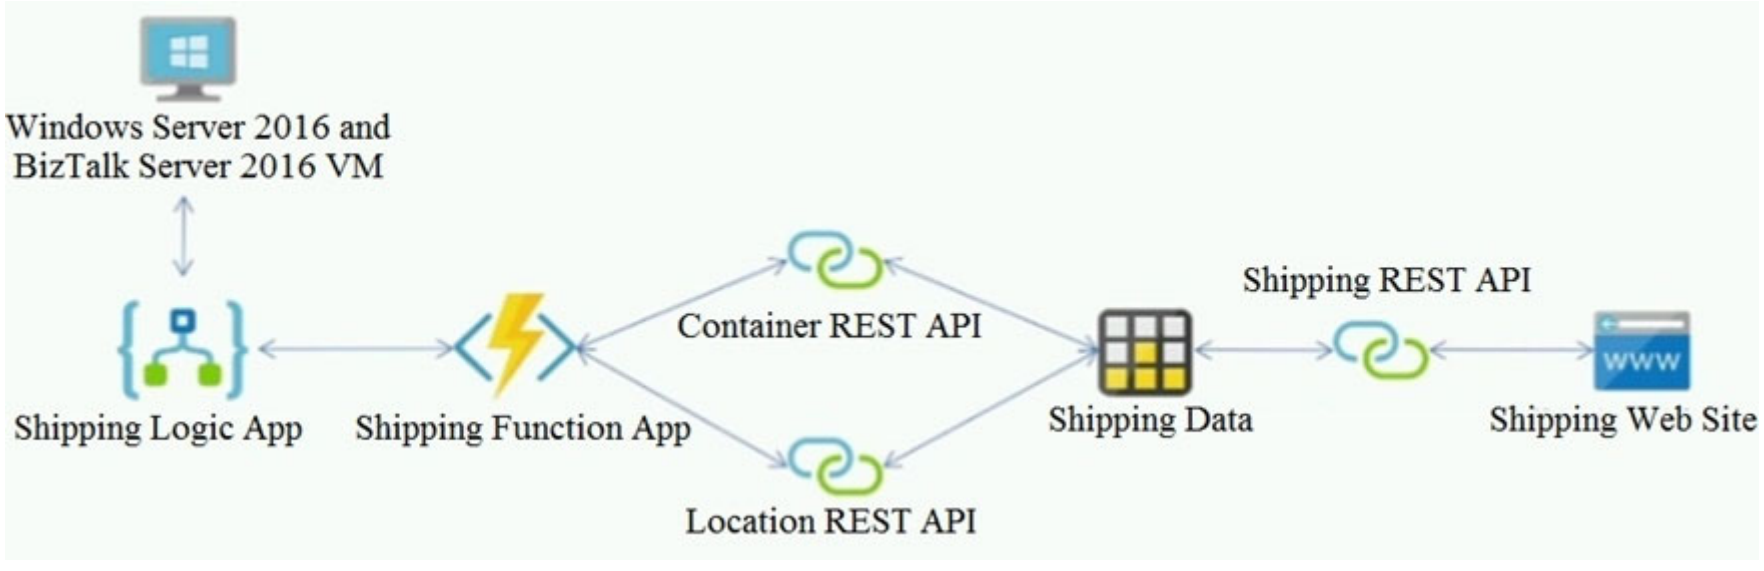

Requirements -

Shipping Logic app -
The Shipping Logic app must meet the following requirements:
Support the ocean transport and inland transport workflows by using a Logic App.
Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app -
Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

REST APIs -
The REST API's that support the solution must meet the following requirements:
Secure resources to the corporate VNet.
Allow deployment to a testing location within Azure while not incurring additional costs.
Automatically scale to double capacity during peak shipping times while not causing application downtime.
Minimize costs when selecting an Azure payment model.

Shipping data -
Data migration from on-premises to Azure must minimize costs and downtime.

Shipping website -
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Issues -

Windows Server 2016 VM -
The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

Shipping website and REST APIs -
The following error message displays while you are testing the website:
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource.
Origin 'http://test.wideworldimporters.com/' is therefore not allowed access.

Question
HOTSPOT -
You need to secure the Shipping Function app.
How should you configure the app? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

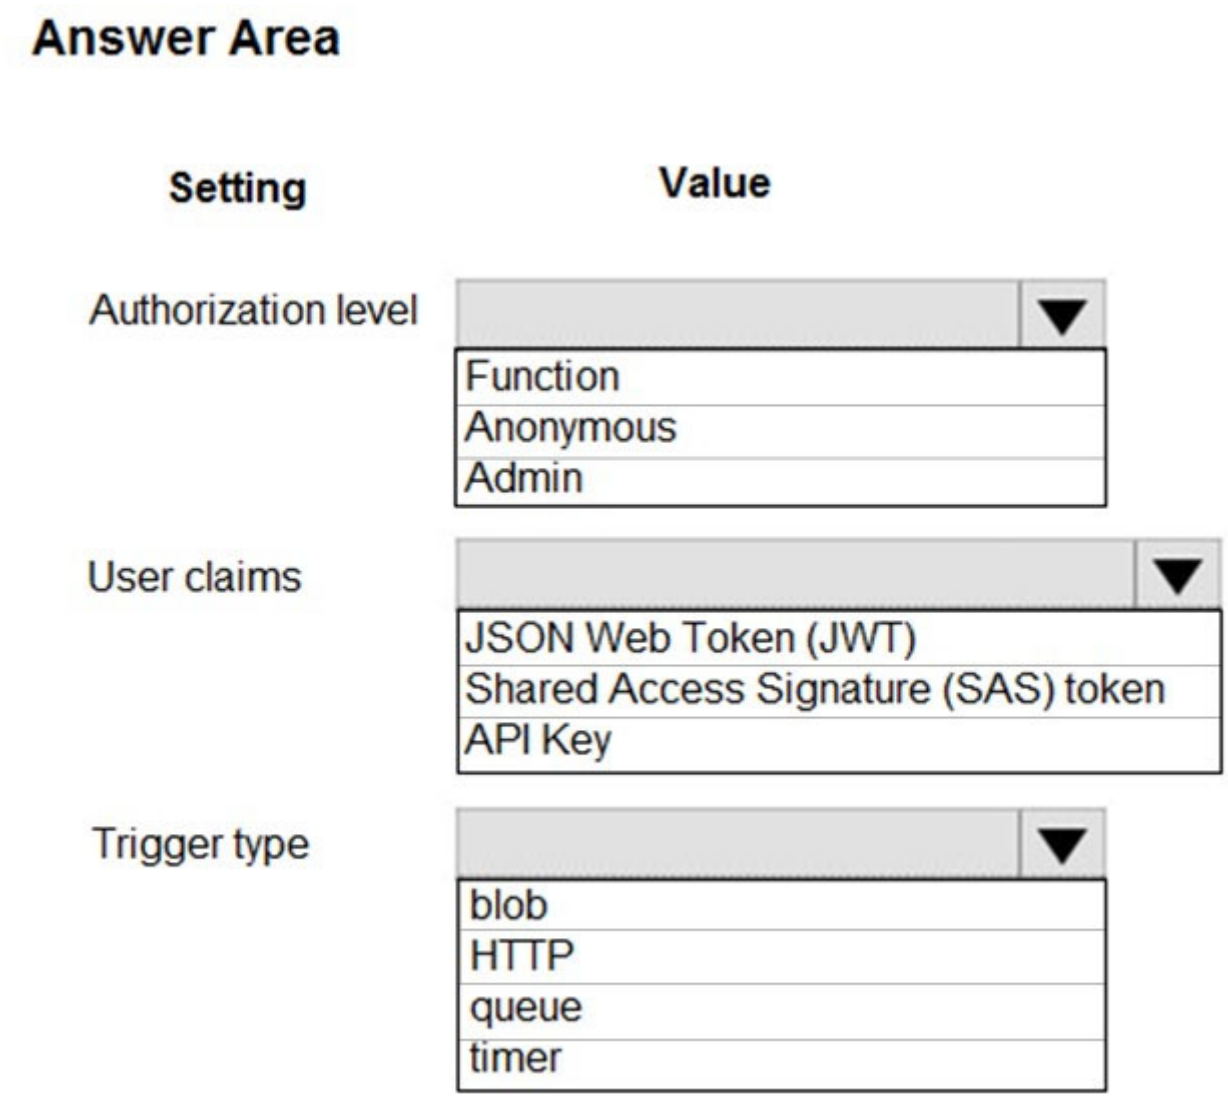

Correct Answer:

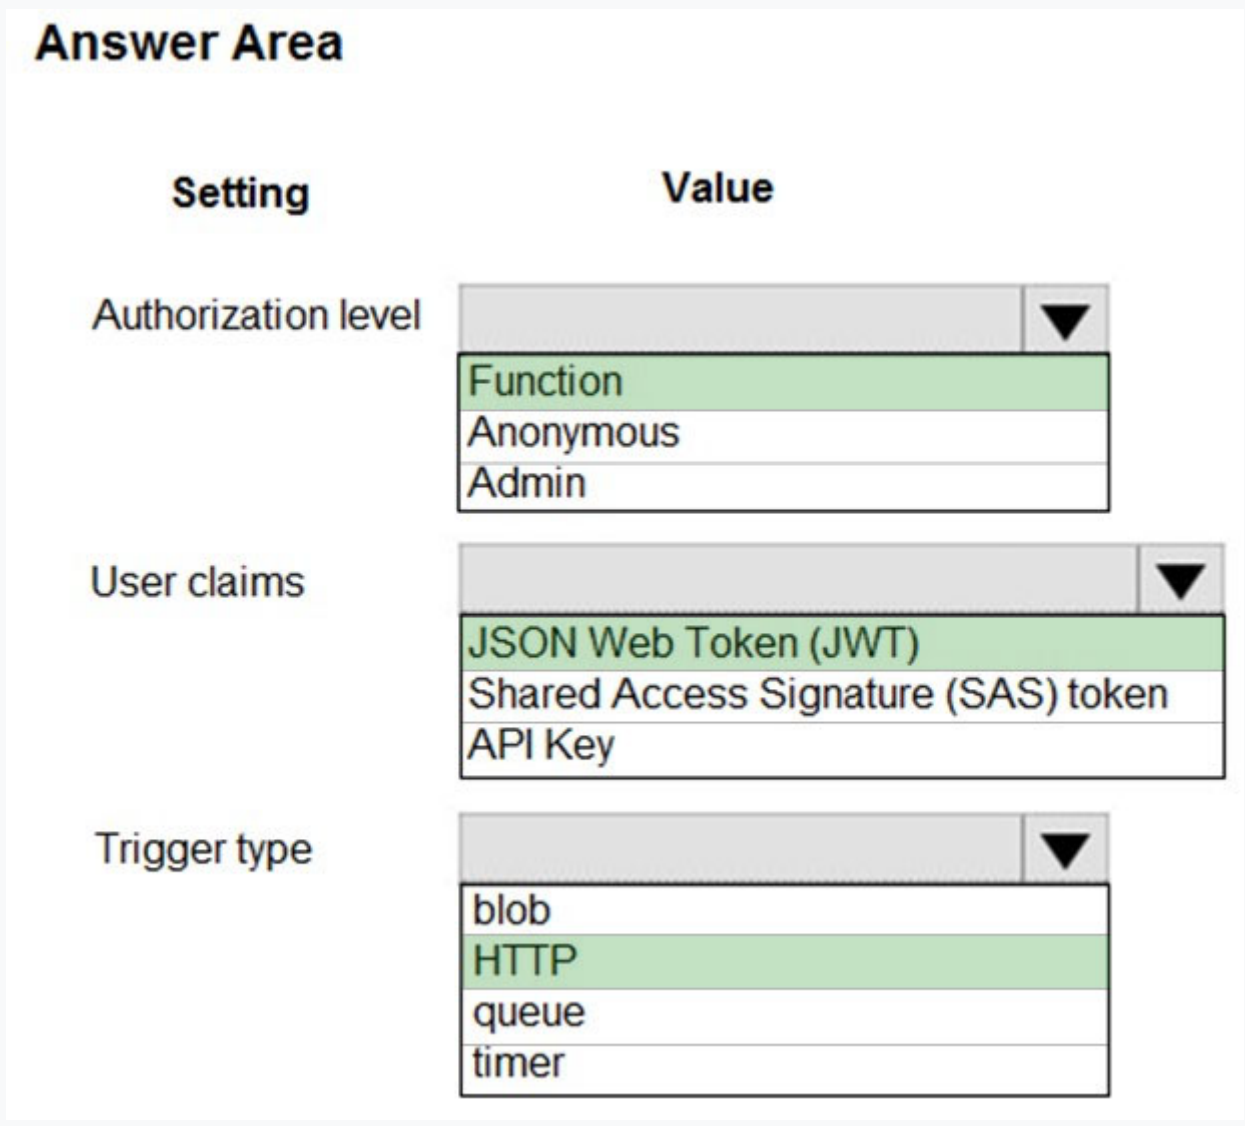

Scenario: Shipping Function app: Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

Box 1: Function -

Box 2: JSON based Token (JWT)

Azure AD uses JSON based tokens (JWTs) that contain claims

Box 3: HTTP -

How a web app delegates sign-in to Azure AD and obtains a token

User authentication happens via the browser. The OpenID protocol uses standard HTTP protocol messages.
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/develop/authentication-scenarios

Eltooth Highly Voted 1 year, 4 months ago

Function
JWT
HTTP

upvoted 17 times

uffuchsi Highly Voted 8 months, 1 week ago

The correct answer is:
Anonymous
JWT
HTTP

Scenario: Shipping Function app: Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

1. Authorization Level must be anonymous to use function app level security methods.
2. User claims must be JWT tokens; API Key is not recommended due to security issues.
3. Function App is triggered from Logic App. So it must be http

https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-http-webhook-trigger?tabs=csharp#secure-an-http-endpoint-in-production

upvoted 16 times

adilkhan Most Recent 8 months, 4 weeks ago

Secure an HTTP endpoint in production
To fully secure your function endpoints in production, you should consider implementing one of the following function app-level security options. When using one of these function app-level security methods, you should set the HTTP-triggered function authorization level to anonymous.
https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-http-webhook-trigger?tabs=in-process%2Cfunctionsv2&pivots=programming-language-csharp#secure-an-http-endpoint-in-production Anonymous JWT and HTTP

upvoted 2 times

Ranzzzan 9 months ago

I will go with anonymous
https://learn.microsoft.com/en-us/answers/questions/801055/authenticate-azure-functions

upvoted 2 times

Knightie 11 months, 1 week ago

the answer is correct, shipping function app is the one at the left.. it will need to access the 2 apis from the internal VM, not public accessing. API key is correct instead of anonymous.

upvoted 1 times

lombster 11 months, 3 weeks ago

Based on description it should be anonymous
https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-http-webhook-trigger?tabs=in-process%2Cfunctionsv2&pivots=programming-language-csharp#secure-an-http-endpoint-in-production

upvoted 4 times

ajayasa 9 months ago

based on the link posted i see it should be anonymous but not tried on it

upvoted 1 times

gmishra88 1 year ago

If using AD is used then Function key should not be used and it should be set to anonymous. Could not find the link, but just think about it and try it

upvoted 2 times

coffecold 12 months ago

here is the link:
Anonymous is the only authorization level that doesn't require a key.
So :

function,
JWT
HTTP

is still valid .
https://learn.microsoft.com/en-us/azure/azure-functions/security-concepts?tabs=v4#function-access-keys

upvoted 3 times

adilkhan 8 months, 4 weeks ago

Functions lets you use keys to make it harder to access your HTTP function endpoints during development. Unless the HTTP access level on an HTTP triggered function is set to anonymous, requests must include an API access key in the request.

NO! its anonymous

upvoted 1 times

qwerty112233 1 year, 1 month ago

"Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD)." How can it be Function is AAD requires "Anonymous"?
My vote:
Anonymous
JWT
HTTP

upvoted 4 times

sghaha 1 year, 6 months ago

in Korean
https://docs.microsoft.com/ko-kr/azure/active-directory/develop/authentication-scenarios

upvoted 3 times

Question #2 *Topic 12*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment -

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk Server 2016. The VM runs the following workflows:

Ocean Transport `" This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.

Inland Transport `" This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

Container API `" This API provides container information including weight, contents, and other attributes.

Location API `" This API provides location information regarding shipping ports of call and trucking stops.

Shipping REST API `" This API provides shipping information for use and display on the shipping website.

Shipping Data -

The application uses MongoDB JSON document storage database for all container and transport information.

Shipping Web Site -

The site displays shipping container tracking information and container contents. The site is located at http://shipping.wideworldimporters.com/

Proposed solution -

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:

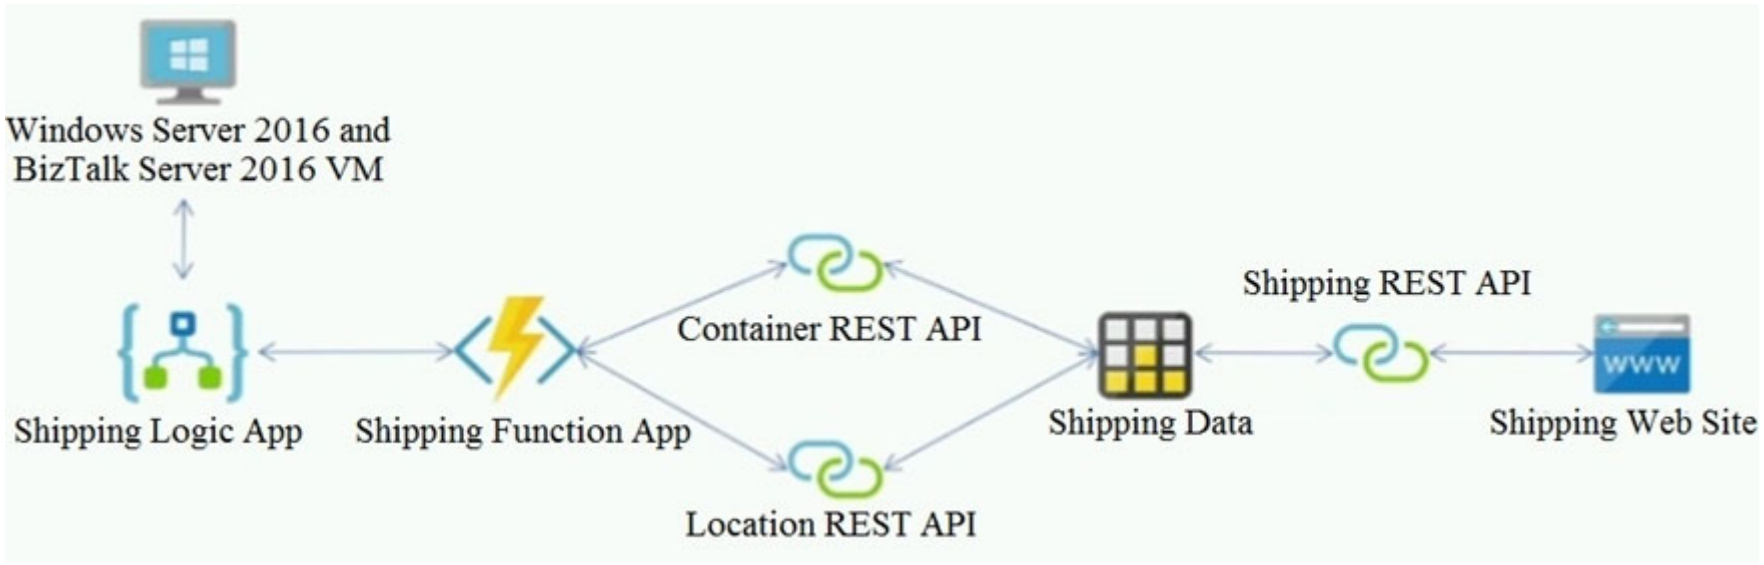

Requirements -

Shipping Logic app -
The Shipping Logic app must meet the following requirements:
Support the ocean transport and inland transport workflows by using a Logic App.
Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app -
Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

REST APIs -
The REST API's that support the solution must meet the following requirements:
Secure resources to the corporate VNet.
Allow deployment to a testing location within Azure while not incurring additional costs.
Automatically scale to double capacity during peak shipping times while not causing application downtime.
Minimize costs when selecting an Azure payment model.

Shipping data -
Data migration from on-premises to Azure must minimize costs and downtime.

Shipping website -
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Issues -

Windows Server 2016 VM -
The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

Shipping website and REST APIs -
The following error message displays while you are testing the website:
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource.
Origin 'http://test.wideworldimporters.com/' is therefore not allowed access.

Question
You need to secure the Shipping Logic App.
What should you use?

A. Azure App Service Environment (ASE)

B. Integration Service Environment (ISE)

C. VNet service endpoint

D. Azure AD B2B integration

Correct Answer: *B*
Scenario: The Shipping Logic App requires secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
You can access to Azure Virtual Network resources from Azure Logic Apps by using integration service environments (ISEs).
Sometimes, your logic apps and integration accounts need access to secured resources, such as virtual machines (VMs) and other systems or services, that are inside an Azure virtual network. To set up this access, you can create an integration service environment (ISE) where you can run your logic apps and create your integration accounts.
Reference:
https://docs.microsoft.com/en-us/azure/logic-apps/connect-virtual-network-vnet-isolated-environment-overview

Community vote distribution
B (80%) C (20%)

NKnab Highly Voted 3 years, 3 months ago

B is correct. confirmed
upvoted 34 times

mlantonis Highly Voted 2 years, 5 months ago

Shipping Logic app:
Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.

For scenarios where your logic apps and integration accounts need access to an Azure virtual network, create an integration service environment (ISE). An ISE is a dedicated environment that uses dedicated storage and other resources that are kept separate from the "global" multi-tenant Logic Apps service. This separation also reduces any impact that other Azure tenants might have on your apps' performance. An ISE also provides you with your own static IP addresses. These IP addresses are separate from the static IP addresses that are shared by the logic apps in the public, multi-tenant service.

Reference:
https://docs.microsoft.com/en-us/azure/logic-apps/connect-virtual-network-vnet-isolated-environment
upvoted 7 times

mlantonis 2 years, 5 months ago

B. Integration Service Environment (ISE)
upvoted 6 times

br00net Most Recent 4 months, 2 weeks ago

Correct answer is A:
App Service Environments are appropriate for application workloads that require:
High scale.
Isolation and secure network access.
High memory utilization.
High requests per second (RPS). You can create multiple App Service Environments in a single Azure region or across multiple Azure regions. This flexibility makes an App Service Environment ideal for horizontally scaling stateless applications with a high RPS requirement.
https://learn.microsoft.com/en-us/azure/app-service/environment/overview

There is nothing about security, apart from that ISE is "fully isolated": https://azure.microsoft.com/en-us/updates/integration-service-environments-ise-are-now-generally-available/?cdn=disable
upvoted 1 times

hubekpeter 11 months, 1 week ago

Selected Answer: C
You can't create ISE starting 1st of November and it'll be retired in 2024, you should create vnet instead. https://learn.microsoft.com/en-us/azure/logic-apps/connect-virtual-network-vnet-isolated-environment-overview
upvoted 1 times

hubekpeter 11 months, 1 week ago

Access to Azure virtual networks from Azure Logic Apps using an integration service environment (ISE) https://learn.microsoft.com/en-us/azure/logic-apps/secure-single-tenant-workflow-virtual-network-private-endpoint
upvoted 1 times

hubekpeter 11 months, 1 week ago

Sorry this it the correct link name - Secure traffic between Standard logic apps and Azure virtual networks using private endpoints
upvoted 1 times

adilkhan 8 months, 4 weeks ago

https://azure.microsoft.com/en-us/updates/integration-services-environment-will-be-retired-on-31-august-2024-transition-to-logic-apps-standard/ logic app is not in the list
upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: B
B is correct answer.
upvoted 4 times

Topic 13 - Testlet 15

Question #1 *Topic 13*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

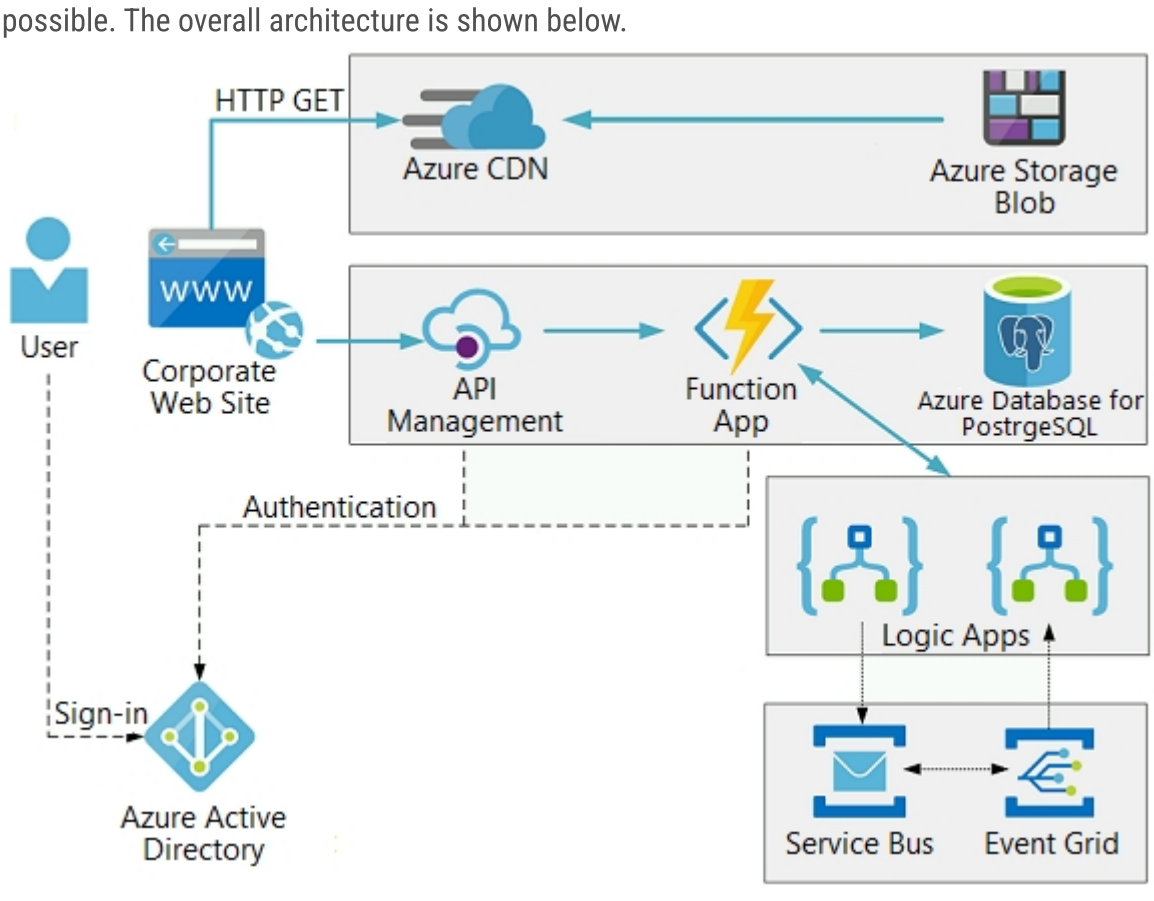

User authentication -

The following steps detail the user authentication process:
1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -

Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

HOTSPOT -

You need to retrieve the database connection string.

Which values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

REST API Endpoint:

https:// [ ] .vault.azure.net/secrets/ [ ] /

cpandlkeyvault
PostgreSQLConn
80df3e46ffcd4f1cb187f79905e9a1e8

cpandlkeyvault
PostgreSQLConn
80df3e46ffcd4f1cb187f79905e9a1e8

Variable type to access Azure Key Vault secret values:

[ ]

Environment
Session
ViewState
Querystring

Correct Answer:

Answer Area

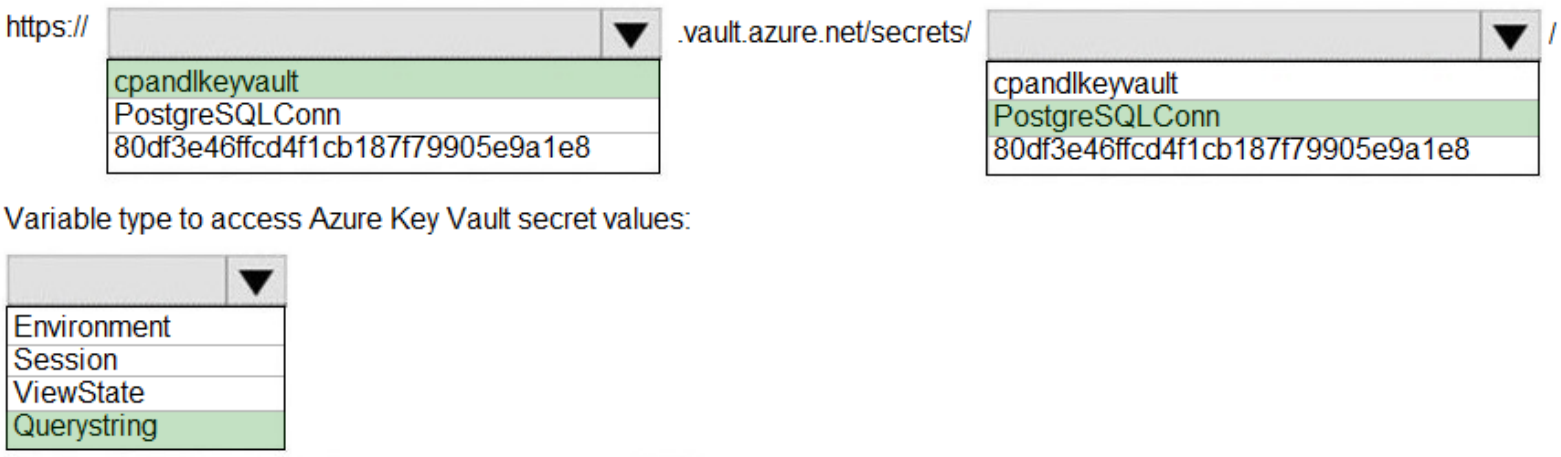

Azure database connection string retrieve REST API vault.azure.net/secrets/

Box 1: cpandlkeyvault -

We specify the key vault, cpandlkeyvault.

Scenario: The database connection string is stored in Azure Key Vault with the following attributes:

Azure Key Vault name: cpandlkeyvault

Secret name: PostgreSQLConn -

Id: 80df3e46ffcd4f1cb187f79905e9a1e8

Box 2: PostgreSQLConn -

We specify the secret, PostgreSQLConn

Example, sample request:

https://myvault.vault.azure.net//secrets/mysecretname/4387e9f3d6e14c459867679a90fd0f79?api-version=7.1

Box 3: Querystring -

Reference:

https://docs.microsoft.com/en-us/rest/api/keyvault/getsecret/getsecret

mlantonis Highly Voted 2 years, 5 months ago

As per requirement:
- Azure Key Vault name: cpandlkeyvault
- Secret name: PostgreSQLConn
- Id: 80df3e46ffcd4f1cb187f79905e9a1e8
https://myvault.vault.azure.net//secrets/mysecretname/4387e9f3d6e14c459867679a90fd0f79?api-version=7.1

Box 1: cpandlkeyvault
We specify the key vault, cpandlkeyvault.

Box 2: PostgreSQLConn
We specify the secret, PostgreSQLConn.

Box 3: Environment
If a reference is not resolved properly, the reference value will be used instead. This means that for application settings, an environment variable would be created

upvoted 44 times

davidkerr7 4 months, 2 weeks ago

correct
https://learn.microsoft.com/en-us/azure/app-service/app-service-key-vault-references?tabs=azure-cli

upvoted 1 times

cool_tool 2 years, 2 months ago

Box 3 is QueryString (query string to specify the API version along with the secret version), the other part is fine. (Answer provided is CORRECT)

upvoted 14 times

coffecold Highly Voted 1 year ago

To prevent reading the cases multiple times:
Please see the spots below where you can find the questions (page/topic/question/subject)
City Power & Light company
50 13 1 API endpoint to Key Vault + variable
50 13 2 create and import certificate in azure web app
50 13 3 configure APIM for authentication with JWT

50 13 4 authenticate user by JWT
50 13 5 MI Authentication to Azure Logic app
51 13 6 Azure Service Bus to Event Grid integration Tier of bus + RBAC role
52 17 1 Application Insights/Monitor/Log Analytics
52 17 2 Azure Blob storage settings SAS or MI, file auditing
52 18 1 solve function timeout
55 24 1 cli for integration Azure Service Bus and Azure Event Grid
55 24 2 ingestion for Grid events
58 32 1 create appropriate storage account + geo settings + cool/hot

upvoted 19 times

Yogendra_examtopic 4 weeks, 1 day ago

bro you are the real champ..... Thanks a lot Man !!!!!

upvoted 1 times

OPT_001122 10 months, 3 weeks ago

Great help!! You did a great help for all for all the case studies !!

upvoted 2 times

AbidooKing Most Recent 3 months, 2 weeks ago

Got this case on my exam 2023July13. Went with highest wote and scored 917.

upvoted 2 times

nekkilodeon 6 months, 1 week ago

Query string makes nonsense.

You'll have to specify a new version every time it updates in the query string. If you use it as environment variable the following will happen.

If a version is not specified in the reference, then the app will use the latest version that exists in the key vault. When newer versions become available, such as with a rotation event, the app will automatically update and begin using the latest version within 24 hours. The delay is because App Service caches the values of the key vault references and refetches it every 24 hours. Any configuration changes to the app that results in a site restart causes an immediate refetch of all referenced secrets.
https://learn.microsoft.com/en-us/azure/app-service/app-service-key-vault-references?tabs=azure-cli

upvoted 1 times

kociuba 10 months ago

got it on my exam 30-12-2022
score: 818

upvoted 2 times

adilkhan 9 months ago

your answer?

upvoted 2 times

gmishra88 1 year ago

The most debated is the third, I guess Microsoft probably mean how to give this secret reference to the app. It is using appsettings, that is a type of Environment variable. So, I will go for that and hope for the best

upvoted 2 times

gmishra88 1 year ago

Because "The connection information is updated frequently. The application must always use the latest information to connect to the database." It has to be a keyvault reference (or not). These are trick questions to be asked in an interview to start an hour long discussion on various options.

upvoted 1 times

gmishra88 1 year, 1 month ago

"variable type to access azure key vault secret value" < this probably does not mean where the response comes. That comes in "value" property. The question creator probably means how to access keyvault key itself. How to pass that authorization. But considering this is functions I do not think query string is correct.

upvoted 2 times

Pize 1 year, 3 months ago

https://docs.microsoft.com/en-us/rest/api/keyvault/secrets/get-secret/get-secret

upvoted 1 times

vavra 1 year, 4 months ago

I think it's a Function App variable, so it should be Environment

upvoted 2 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer:
Box 1: cpandlkeyvault
Name of the key vault.

Box 2: PostgreSQLConn
Name of the secret.

Box 3: QueryString
Variable type to access the secret (100% correct - Not environment)
upvoted 2 times

lugospod 1 year, 9 months ago
Got this one 01/2022. Went with most voted (to avoid writing answers again) @mlantonis
upvoted 2 times

lugospod 1 year, 9 months ago
nah, i read it wrong, they do not have ID in the URL...it is only in the comments.
upvoted 1 times

lugospod 1 year, 9 months ago
I have an issue with the "original question" meaning, they are assuming that we would send ID of the secret version to the HTTP GET. This means that we would ALWAYS retrieve the SAME version of the secret, and would not get the latest version. And they specifically said that the latest version should be used, so based on the information on page
https://docs.microsoft.com/en-us/rest/api/keyvault/getsecret/getsecret
we see that the ID is OPTIONAL, and if we do omit it the latest version will be retrieved.

My comment is just to point out the possibility that the proposed solution by Microsoft is invalid.
upvoted 1 times

lugospod 1 year, 9 months ago
Ignore this, read it wrong. ID is not used in their question.
upvoted 1 times

MiraA 2 years, 1 month ago
I think it could be "Environment".

The question "Variable type to access Azure Key Vault secret values" could mean "what object type instance do I need to call GET request to API endpoint to retrieve the connection string"?

I believe the GET request to retrieve the secret (connection string) must be authorized by the Azure. This means the application had signed in Azure AD already and had retrieved the access token (id_token) which can be added as a HTTP header to GET requests on endpoint API:
GET https://cpandlkeyvault.vault.azure.net/secrets/.....
Authorization: Bearer <bearer-token>

So this means the access token has to be stored somewhere to allow subsequent repeated calls to get secret connection? Maybe Environment is such place?
upvoted 1 times

MiraA 2 years, 1 month ago
I think it could be "Environment".

The question "Variable type to access Azure Key Vault secret values" could mean "what object type instance do I need to call GET request to API endpoint to retrieve the connection string"?

I believe the GET request to retrieve the secret (connection string) must be authorized by the Azure. This means the application had signed in Azure AD already and had retrieved the access token (id_token) which can be added as a HTTP header to GET requests on endpoint API:
GET https://cpandlkeyvault.vault.azure.net/secrets/.....
Authorization: Bearer <bearer-token>

So this means the access token has to be stored somewhere to allow subsequent repeated calls to get secret connection? Maybe Environment is such place?

PS: Not sure if adding the access token to Environment is not a security risk?
upvoted 2 times

MiraA 2 years, 1 month ago
I think it could be "Environment".

The question "Variable type to access Azure Key Vault secret values" could mean "what object type instance do I need to call GET request to API endpoint to retrieve the connection string"?

I believe the GET request to retrieve the secret (connection string) must be authorized by the Azure. This means the application had signed in Azure AD already and had retrieved the access token (id_token) which can be added as a HTTP header to GET requests on endpoint API:
GET https://cpandlkeyvault.vault.azure.net/secrets/.....
Authorization: Bearer <bearer-token>

So this means the access token has to be stored somewhere to allow subsequent repeated calls to get secret connection? Maybe Environment is such place?

PS: Not sure if adding the access token to Environment is not a security risk?
upvoted 1 times

rustycables 2 years, 2 months ago
I'm not happy with answers and discussions on Box 3. Typically this type of setting would be done as an environment variable. The only thing that makes me hesitate is this line: The connection information is updated frequently. The application must always use the latest information to connect to the database. Nothing stops you calling this code each time the connection is needed. In that way the connection string will not be stored in any

of the variable types listed. But that is a big overhead. If you wanted to store the connection string between requests it must be stored in one of these variables types. query string is the least secure, and I would never consider it for this type of value as it exposes parts of the application to an attacker. I would select session, because session can still be updated, and is more secure than both query and viewstate. I would think twice about setting as an environment variable because of the line above that states that the connection changes frequently. I could be wrong, but if I get this question I'm going to choose session.

upvoted 3 times

warchoon 6 months, 3 weeks ago

There is a rotation of environment variable for it
https://learn.microsoft.com/en-us/azure/key-vault/secrets/secrets-best-practices#secrets-rotation:~:text=Secrets%20are%20often%20stored%20in%20application%20memory%20as%20environment%20variables

upvoted 1 times

ning 2 years, 2 months ago

see this link https://azure.microsoft.com/en-us/updates/versions-no-longer-required-for-key-vault-references-in-app-service-and-azure-functions/

upvoted 2 times

Question #2 *Topic 13*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

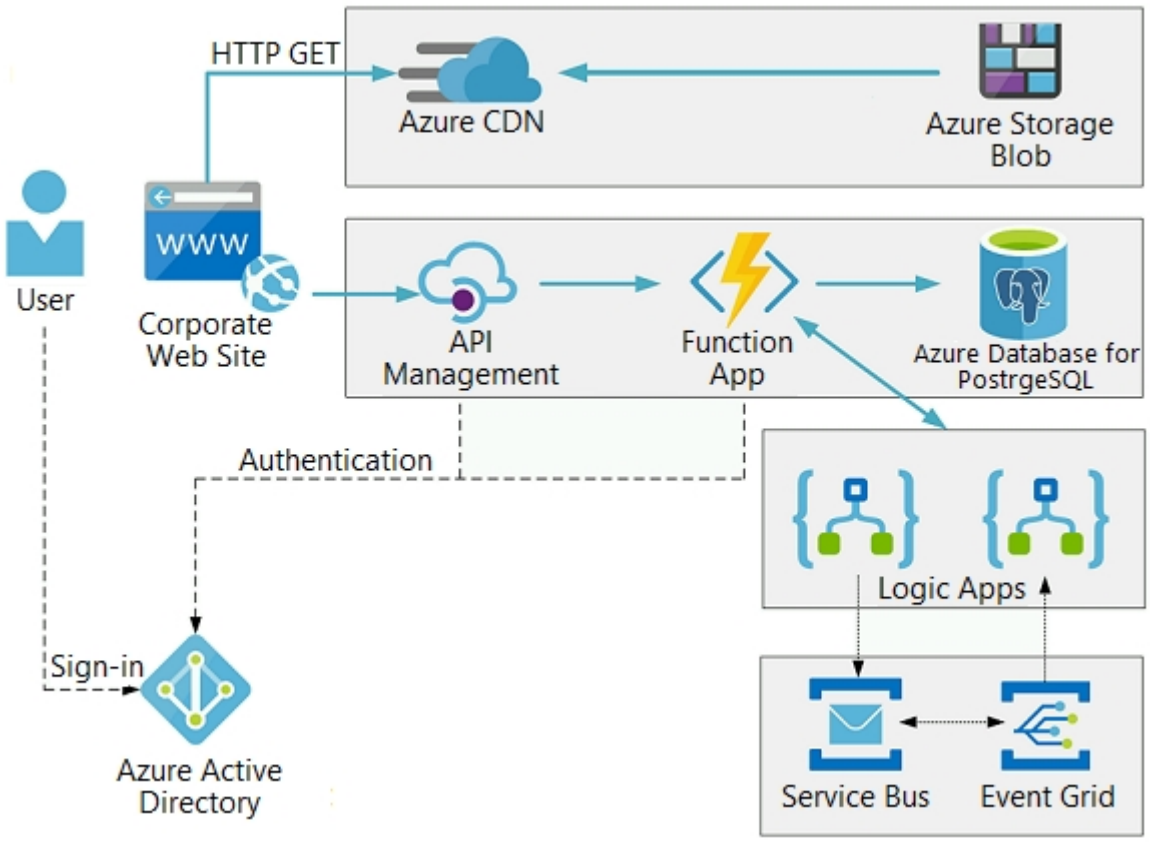

User authentication -

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -
Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

DRAG DROP -

You need to correct the corporate website error.

Which four actions should you recommend be performed 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:

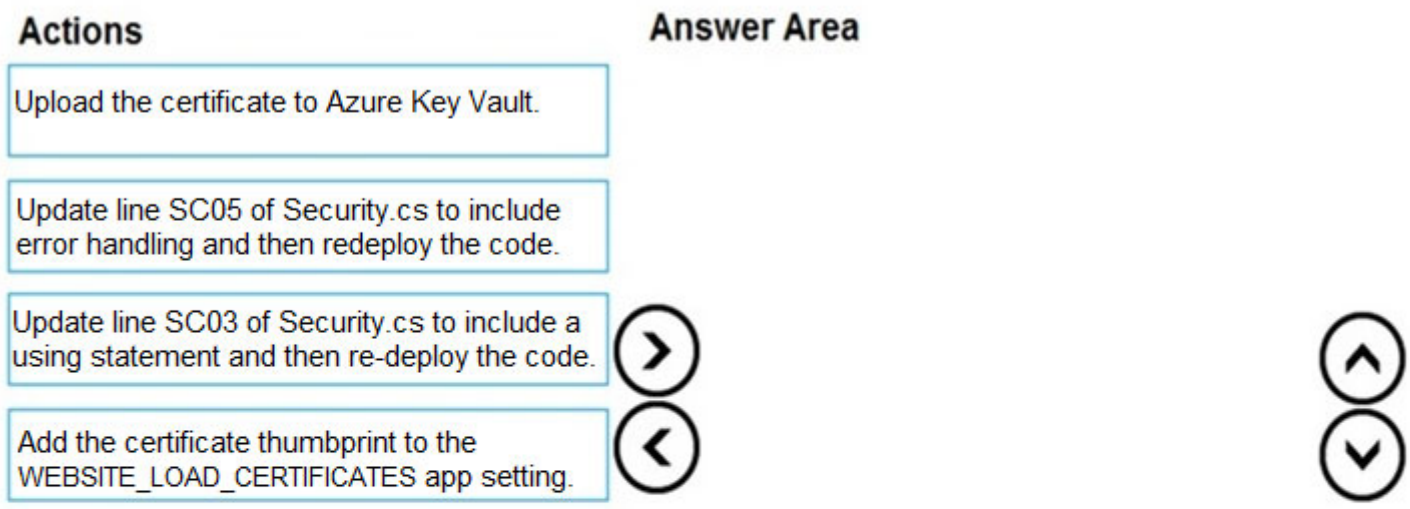

Correct Answer:

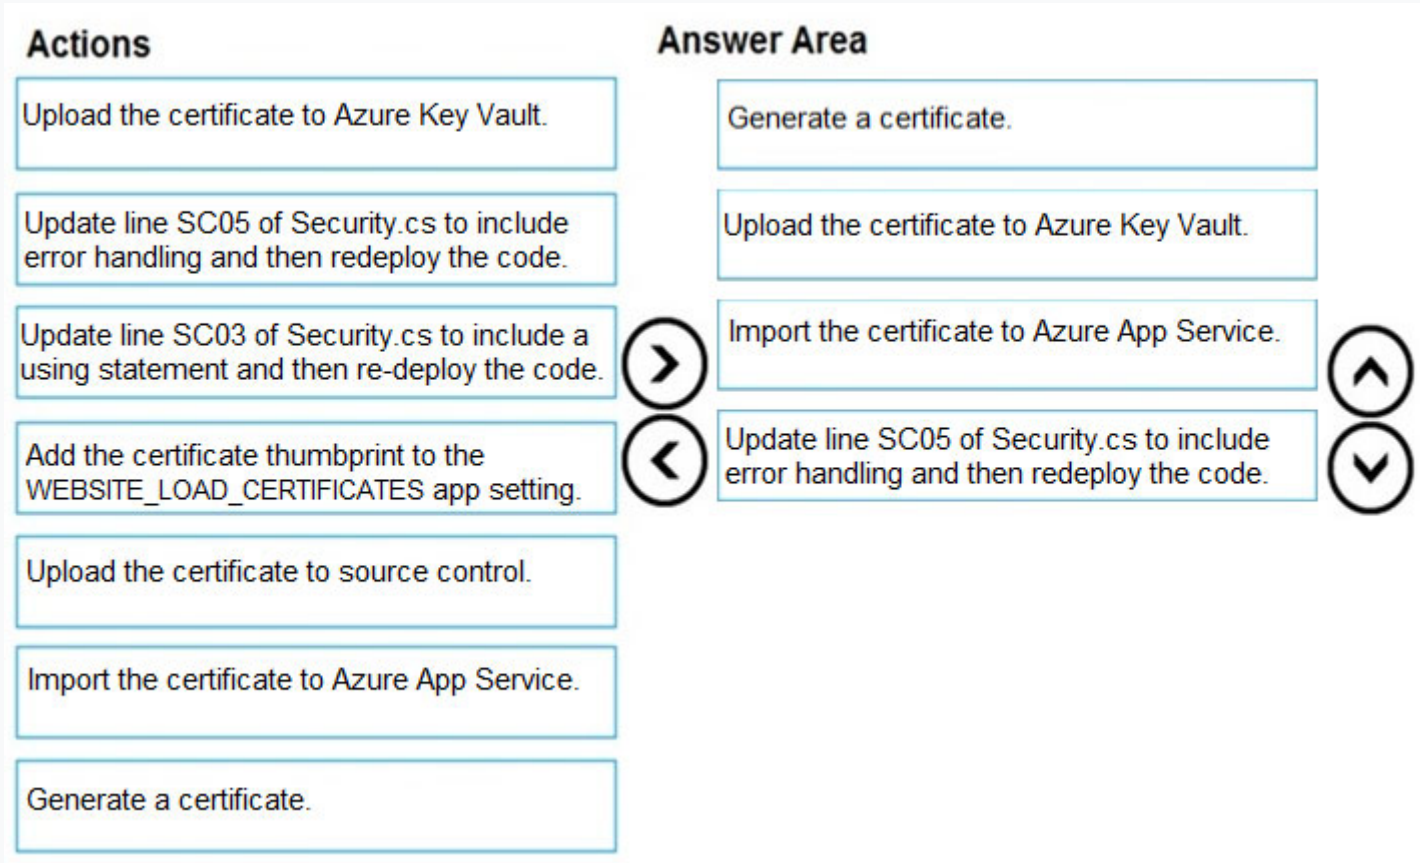

Scenario: Corporate website -

While testing the site, the following error message displays:

CryptographicException: The system cannot find the file specified.

Step 1: Generate a certificate -

Step 2: Upload the certificate to Azure Key Vault

Scenario: All SSL certificates and credentials must be stored in Azure Key Vault.

Step 3: Import the certificate to Azure App Service

Step 4: Update line SCO5 of Security.cs to include error handling and then redeploy the code

Reference:

https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate

Frakandel Highly Voted 2 years, 5 months ago

Answer:
• Generate a certificate
• Upload the certificate to Azure key vault
• Import the certificate to Azure App Service
• Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting

https://ankitvijay.net/2021/04/14/certificate-azure-app-service-linux/

upvoted 73 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with above answer, score 926

upvoted 5 times

SnakePlissken Highly Voted 2 years, 5 months ago

1. Generate a certificate.
2. Upload the certificate to Azure Key Vault.
3. Import the certificate to Azure App Service.
4. Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting.

Prerequisite: Scale up your App Service plan if your app is in a shared infrastructure tier.
https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate
https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate-in-code#load-certificate-in-linuxwindows-containers

upvoted 24 times

acjdev514 Most Recent 1 month ago

My vote: the answer provided is the correct one.
https://learn.microsoft.com/en-us/azure/app-service/configure-ssl-certificate-in-code

upvoted 1 times

surprise0011 6 months, 2 weeks ago

correct
https://learn.microsoft.com/en-us/azure/app-service/configure-ssl-certificate-in-code#when-updating-renewing-a-certificate

upvoted 1 times

st0rmtrooperx 10 months, 2 weeks ago

Got this on Dec 16th, 2022. Scored 921 and went with:
• Generate a certificate
• Upload the certificate to Azure key vault

• Import the certificate to Azure App Service
• Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting

upvoted 7 times

coffecold 1 year ago

Answer of Frakadel OK.
But otherwise : this is bad practice. Azure Key Vault with managed identity should be used for that purpose.

upvoted 3 times

gmishra88 1 year ago

This is an essay question, but I will bet my time on :
Generate
Upload to keyvault
Import to App service (from keyvault)
Add to WEBSITE_LOAD_CERTIFICATES (Important so that it is available in /var/ssl/... or c:\appservies\..\secret path. The code seems to be using that)

upvoted 2 times

Sandeep12093 1 year, 4 months ago

Agree to brettusmaximus
1. Generate a certificate
2. Upload to key vault (mandatory security requirements)
3. Import the certificate to Azure App Service. For Linux container, it keeps the private certificates at the location /var/ssl/private and public certificates at /var/ssl/certs.
4. Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting.

Refer configure-SSL-CERTIFICATE-IN-CODE

upvoted 1 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 6 times

asdasdasg2 1 year, 9 months ago

what a horrible question. Security.cs is not even valid C# and its usage is not even shown in the following snippet that contains the function.

upvoted 5 times

GreenPanda 2 years, 2 months ago

Does the ReadAllBytes method need to specify the file path instead of the directory?

upvoted 1 times

lugospod 1 year, 9 months ago

yes I agree, but I think this is a typo...the same way the path has ~ in front...

upvoted 1 times

GreenPanda 2 years, 2 months ago

https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate-in-code#load-certificate-in-linuxwindows-containers

upvoted 1 times

BrettusMaximus 2 years, 6 months ago

Ok - How does the cert get to be located at the URL (/var/ssl/private)?
1. Generate a certificate
2. Upload to key vault
3. Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting.
4. Add the certificate to Source Control (/var/ssl/private), So when the app deploys it gets there.
It doesn't matter about nice error messages.

upvoted 1 times

BrettusMaximus 2 years, 6 months ago

OK I am slightly wrong here; as in https://ankitvijay.net/2021/04/14/certificate-azure-app-service-linux/
1. Generate a certificate
2. Upload to key vault (mandatory security requirements)
3. Import the certificate to Azure App Service. For Linux container, it keeps the private certificates at the location /var/ssl/private and public certificates at /var/ssl/certs.
4. Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting.

It doesn't matter about nice error messages.

upvoted 15 times

[Removed] 2 years, 6 months ago

General information

* CryptographicException: The system cannot find the file specified
* According to the path we are running a linux container (/var/ssl/private)
* Question seems to have a typo (line SC03, ~ is not needed)

Answer

1. Generate a certificate

2. Import the certificate to Azure App Service
3. Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting
4. Update line SC05 to include error handling ...

Further reading:

https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate-in-code#load-certificate-in-linuxwindows-containers

upvoted 2 times

lugospod 1 year, 9 months ago

Adding exception handling won't help you solve the problem...so not needed. You Must use key vault so you are missing that part as a second step.

upvoted 2 times

anastakasim 2 years, 6 months ago

So what is the correct answer to this?

upvoted 3 times

clarionprogrammer 2 years, 6 months ago

1. Generate a certificate
2. Import the certificate to Azure App Service.
3. Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting.
https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate-in-code#make-the-certificate-accessible
4. Update line SC05 of Security.cs to include error handling and then redeploy the code.

upvoted 4 times

SwatiN 2 years, 6 months ago

Certificates should be imported from azure key vault. its one of the requirement. so it would be :
1. Generate Cert
2. Upload to key vault
3. import from key vault.
4. Update config thumbprint value
5. Handle the error in the code gracefully

https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate
https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate-in-code#load-certificate-in-linuxwindows-containers

upvoted 2 times

lugospod 1 year, 9 months ago

The question is to select FOUR steps..you selected FIVE.

upvoted 1 times

jokergester 2 years, 7 months ago

I think it should be
1. Generate Cert
2. Upload to App Service
3. Update config thumbprint value
4. Handle the error in the code gracefully

Uploading to Key Vault would offer a better security but another step(s) would be needed it is to add access policy for the App Service to use the Certificate

upvoted 2 times

lugospod 1 year, 9 months ago

There is a requirement that states that key vault MUST be used to store certificates.

upvoted 2 times

Zsolt72 2 years, 7 months ago

https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate-in-code#make-the-certificate-accessible

upvoted 1 times

Question #3 *Topic 13*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

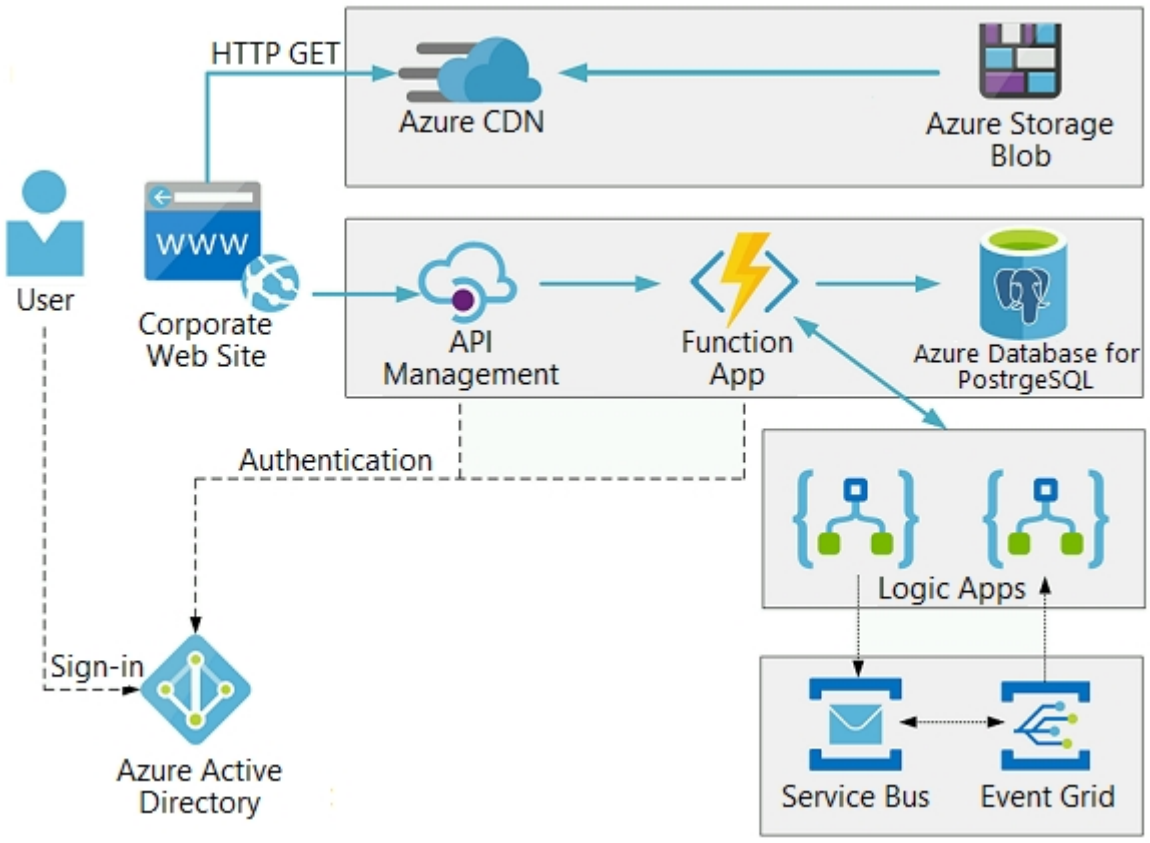

User authentication -

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -
Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

HOTSPOT -

You need to configure API Management for authentication.

Which policy values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Setting	Value
Policy	Check HTTP header Restrict caller IPs Limit call rate by key Validate JWT
Policy section	Inbound Outbound

Correct Answer:

Answer Area

Setting	Value
Policy	Check HTTP header Restrict caller IPs Limit call rate by key **Validate JWT**
Policy section	Inbound **Outbound**

Box 1: Validate JWT -

The validate-jwt policy enforces existence and validity of a JWT extracted from either a specified HTTP Header or a specified query parameter.

Scenario: User authentication (see step 5 below)

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Incorrect Answers:

☞ Limit call rate by key - Prevents API usage spikes by limiting call rate, on a per key basis.

☞ Restrict caller IPs - Filters (allows/denies) calls from specific IP addresses and/or address ranges.

☞ Check HTTP header - Enforces existence and/or value of a HTTP Header.

Box 2: Outbound -

Reference:

https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies

ZiadAZ Highly Voted 2 years, 7 months ago

The second box should be Inbound instead of Outbound.
https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies
Quoting:
This policy can be used in the following policy sections and scopes.
Policy sections: inbound
Policy scopes: all scopes

upvoted 104 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with above answer, score 926

upvoted 3 times

warchoon 6 months, 3 weeks ago

https://learn.microsoft.com/en-us/azure/api-management/validate-jwt-policy#:~:text=Policy%20sections%3A%20inbound

upvoted 2 times

Dani_ac7 1 year, 3 months ago

Correct, you can't pass request to controllers if certificate doesn't exists...

upvoted 1 times

sam365 2 years, 7 months ago

you are correct. it should be INBOUND for Validate JWT
https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies#ValidateJWT

upvoted 10 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: Validate JWT
The validate-jwt policy enforces existence and validity of a JWT extracted from either a specified HTTP Header or a specified query parameter.

Box 2: Inbound
Authentication should be done on Incoming Request and that should be done in Inbound section of the policy of course.

This policy can be used in the following policy sections and scopes.
Policy sections: inbound
Policy scopes: all scopes

Reference:

https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies

https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies#ValidateJWT

https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies

upvoted 26 times

Ciupaz Most Recent 2 weeks, 1 day ago

JWT=Json Web Token (for who does not know).

upvoted 1 times

OlivierPaudex 1 year, 2 months ago

Agreed with inbound. Only inbound direction can be choose when using a JWT token.

upvoted 3 times

Eltooth 1 year, 4 months ago

JWT
Inbound
upvoted 2 times

SivajiTheBoss 1 year, 7 months ago
Correct answer:
1. JWT
2. Inbound
upvoted 5 times

leonidn 1 year, 9 months ago
Validate JWT
Inbound
upvoted 5 times

fr369 1 year, 9 months ago
Scenario, section 'Authentication':
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience (`~aud') claim in the access token.
6. The back-end API validates the access token.

So, shouldn't the first box be "Check HTTP header"?
upvoted 3 times

lugospod 1 year, 9 months ago
Yeap, it could be done that way also... there is a small performance hit when using JWT validation. There is no need to have the validation on both places except in the case one has a "great" number of invalid JWT tokens and you want to remove the burden from the backend server.

So I think that deciding between checking the existence of the header vs validation JWT in this exams boils down to "lucky guess".
upvoted 2 times

MrXBasit 2 years, 3 months ago
Policy Section should be Inbound
upvoted 3 times

MrXBasit 2 years, 3 months ago
Use the Validate JWT policy to pre-authorize requests in API Management, by validating the access tokens of each incoming request. If a request does not have a valid token, API Management blocks it. Validate JWT policy is part of <inbound> policy section of the the API. It checks the audience claim in an access token, and returns an error message if the token is not valid.

https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies#ValidateJWT
upvoted 4 times

anandhprakash 2 years, 5 months ago
The correct answer is:
Validate JWT
Inbound
upvoted 3 times

pavan555manjunath 2 years, 5 months ago
wht is the correct answer
restrict caller ip & inbound
upvoted 1 times

faizalzain 2 years, 6 months ago
the answer should be restrict caller ip & inbound
upvoted 1 times

SnakePlissken 2 years, 5 months ago
IP restriction is only asked for file access, not API calls.
upvoted 1 times

kwaazaar 2 years, 6 months ago
Shouldn't it also restrict by up? Or was that not intended, since the question is about authentication instead of authorization?
upvoted 2 times

gmishra88 1 year ago
A very good point, but I hope Microsoft guys do not understand the difference between Authentication and Authorization and go with the most-voted answer. Majority wins (does it?). All of these questions are substandard
upvoted 1 times

SubhoG 2 years, 7 months ago
It should be Inbound. Authentication should be done on Incoming Request and that should be done in Inbound section of the policy ofcourse.
upvoted 4 times

Question #4 *Topic 13*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

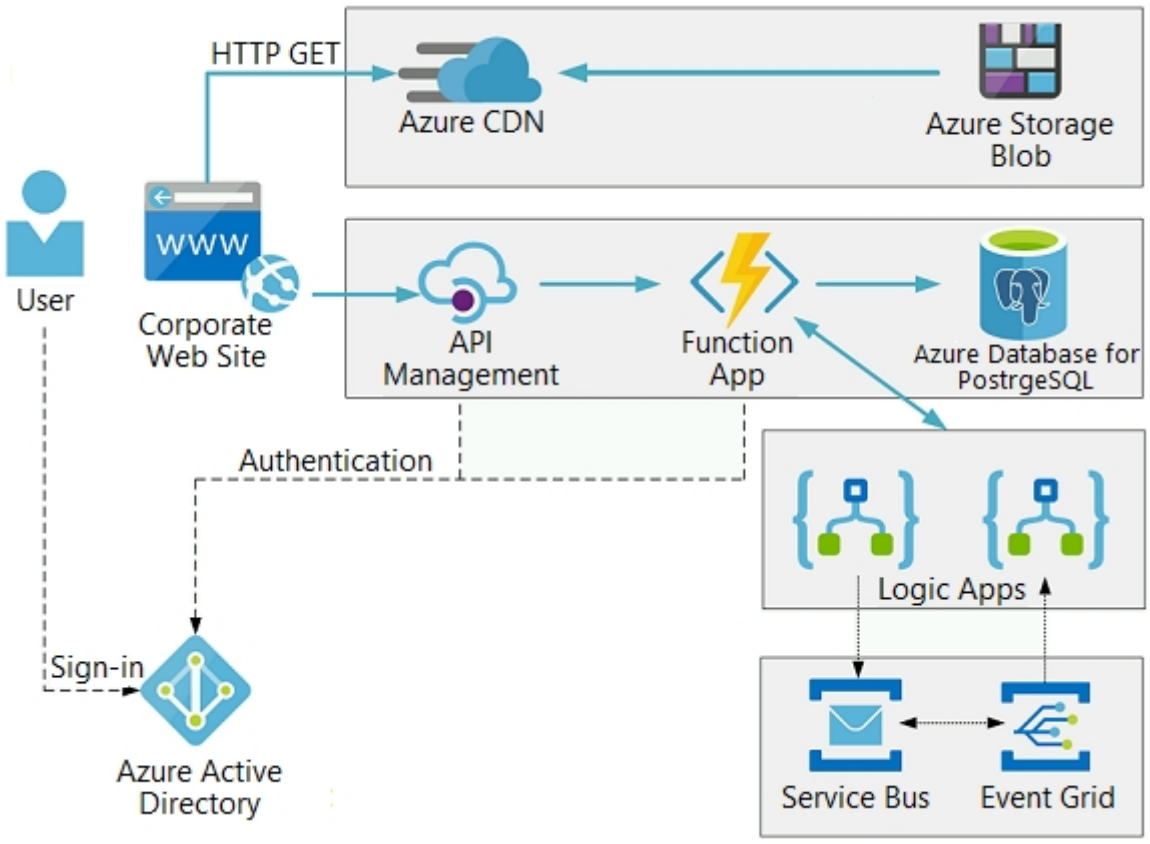

User authentication -

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -
Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

You need to authenticate the user to the corporate website as indicated by the architectural diagram.

Which two values should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. ID token signature

B. ID token claims

C. HTTP response code

D. Azure AD endpoint URI

E. Azure AD tenant ID

Correct Answer: *AD*

A: Claims in access tokens -

JWTs (JSON Web Tokens) are split into three pieces:

☞ Header - Provides information about how to validate the token including information about the type of token and how it was signed.

☞ Payload - Contains all of the important data about the user or app that is attempting to call your service.

☞ Signature - Is the raw material used to validate the token.

E: Your client can get an access token from either the v1.0 endpoint or the v2.0 endpoint using a variety of protocols.

Scenario: User authentication (see step 5 below)

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Reference:

https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies

Community vote distribution

trance13 Highly Voted 2 years, 6 months ago

You need ID Token Claims for authorisation.
You need to validate ID Token signature as a part of authentication.
So the correct answer is A.
Having both D & E does not make sense, URI has tenant ID in it, so answer D.

upvoted 32 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with given answer, score 926

upvoted 4 times

surprise0011 6 months, 2 weeks ago

couldn't agree more.
ID Token signature should be validated to Authenticate User
To get token you need to make a request to Microsoft Identity Platform Endpoint which also contains tenant.
So most appropriate are A and D

upvoted 2 times

faizalzain Highly Voted 2 years, 6 months ago

it should be D & E

upvoted 21 times

Vmwarevirtual Most Recent 5 months ago

Appeared in exam I took at 27-5-2023
I chose A and D

upvoted 4 times

sriazure213 9 months, 3 weeks ago

This question came on jan9,2023. Choose d,e. Score 834

upvoted 2 times

hubekpeter 11 months, 1 week ago

Selected Answer: AE

Guys this is a standard Sign-in flow. I think the correct answer is AE. First you need to know the tenant it so you'll be able to send a sign in request GET https://login.microsoftonline.com/{tenant}/oauth2/v2.0/authorize? and in reply you'll get a id_token (JWT), which in turn can be used to get an access token to UserInfo endpoint. I'm not 100% sure, but this way, it makes sense to me.

upvoted 2 times

warchoon 6 months, 3 weeks ago

If only the solution is multitenant. Which is not.

upvoted 1 times

OPT_001122 11 months, 3 weeks ago

Selected Answer: AD

correct ans

upvoted 3 times

coffecold 1 year ago

Answer A+ B
"Each correct answer presents part of the solution"
From a JWT the header and signature are used to verify the authenticity of the token. The signature is raw information to verify so A is correct.
The claims in the header you need to validate this signature of the token , especially the 'kid' claim. So you need header claims (these are token claims, answer B) as well.

https://learn.microsoft.com/en-us/azure/active-directory/develop/id-tokens
https://learn.microsoft.com/en-us/azure/active-directory/develop/access-tokens

upvoted 1 times

hubekpeter 11 months, 1 week ago

You're getting a claims from the auth endpoint in turn.

upvoted 1 times

gmishra88 1 year ago

"Which two values should you use?" Where, Microsoft, where? While forwarding the request to an authorize endpoint to retrieve an ID token? Or while validating the id-token-signature? The question creator had something in his mind when he suddenly stopped thinking. We have to guess where he stopped. This is a psychology question

upvoted 2 times

gmishra88 1 year ago

I want to go for ID token signature and the token-claims as the answer.
To check authentication, one need to validate first the signature to be sure it is correct. Check the audience is the the application and then of course check the service principal (combination of oid and sub claims). But all these are in the ID token (claims and the signature). So I just need the ID token.

upvoted 1 times

gmishra88 1 year ago

So, it is either AB or DE at least. Just the signature is not enough for authentication, you need also the claims.

upvoted 1 times

gmishra88 1 year ago

wait a minute, I'm overthinking. I do not think this certification department from Microsoft understands difference between Authentication and Authorization. So, right answer at the level of Microsoft intelligence is DE. Look at the "steps to authenticate" they mention access token there. That clearly shows they do not understand the authentications steps. Because it uses an access_token it cannot be A and B.

upvoted 1 times

damianadalid 5 months, 1 week ago

Can't we have claims in access tokens? https://learn.microsoft.com/en-us/azure/active-directory/develop/access-tokens#claims-in-access-tokens

upvoted 1 times

damianadalid 5 months, 1 week ago

..But the possible answers mention ID token signature/claims... Nevermind

upvoted 1 times

OlivierPaudex 1 year, 2 months ago

1. The user have to authenticate clicking on a link in the web app. It means that the "tenant ID" of the AD is used here, not the URI.

2. You have to use the aud claim to ensure that the user intended to call the application. If the identifier of the resource isn't in the aud claim, reject it.

https://docs.microsoft.com/en-us/azure/active-directory/develop/access-tokens#validating-tokens

In my opinion, the right answers are:
B: ID token claims
E: Azure AD tenant ID

Not sure at all, but there was the right options choose by others exam website.

upvoted 2 times

ChrisEvans 1 year, 6 months ago

Selected Answer: AD

Claims are for authorization. The signature is used for authentication, and the endpoint url includes the tenant id, so A&D are correct by process of elimination.

upvoted 5 times

ReniRechner 1 year, 8 months ago

Selected Answer: AD

Authentication is not sign in!
So for example API Management needs to check whether the user is as stated.
So it needs
* D to find the correct authentication "service".
* And to check if the user is as stated "A" is needed, since this is the only option that contains information about the user.

B is not viable since the claims alone are not trustworthy.
I don't know why E would help and it can surely not replace A or D.
C is surely not needed since it contains neither information about the user nor the AD.

upvoted 4 times

warchoon 6 months, 3 weeks ago

https://learn.microsoft.com/en-us/azure/app-service/overview-authentication-authorization#authentication-flow:~:text=Client%20code%20posts%20token%20from%20provider%20to%20/.auth/login/%3Cprovider%3E%20for%20validation.

upvoted 1 times

chingdm 1 year, 8 months ago

To authenticate, you need to know the url of auth service api and the tenant id i.e.
https://login.microsoftonline.com/{tenant}/oauth2/v2.0/authorize?
client_id=6731de76-14a6-49ae-97bc-6eba6914391e
&response_type=code
&redirect_uri=http%3A%2F%2Flocalhost%2Fmyapp%2F
&response_mode=query

after visiting that url, it will be presented the AD consent and login, then auth code is appended to the redirect url, which can be used to get the access token.

so answer is D and E.

upvoted 2 times

leonidn 1 year, 9 months ago

Selected Answer: DE

We are asked about what we need to get authenticated. A and B make sense when a user is authenticated. Hence, both options are invalid. I do not know why we need Response Code. Consequently, options D and E.

upvoted 3 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with DE since they are asking about AUTHENTICATION and not AUTHORIZATION

upvoted 10 times

Loai 1 year, 9 months ago

Selected Answer: DE

this is the correct answer

upvoted 1 times

asdasdasg2 1 year, 9 months ago

AzureAD token is a JWT - therefore A&D are correct. You need the token signature and the claims (for business-logic validation).

https://docs.microsoft.com/en-us/azure/active-directory/develop/access-tokens#validating-tokens

upvoted 1 times

asdasdasg2 1 year, 9 months ago

Sorry, meant A&B

upvoted 1 times

lugospod 1 year, 9 months ago

I would say that the question is about authentication and NOT authorization so this automatically eliminates A & B. C makes no sense, so we are left with two options D & E which you really do need to connect and authenticate via AD (additionally Client/Application ID is needed)

upvoted 3 times

gmishra88 1 year ago

ID token is about authentication and not authorization. Access token is about authorization. But this question is very unclear

upvoted 1 times

Question #5 *Topic 13*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

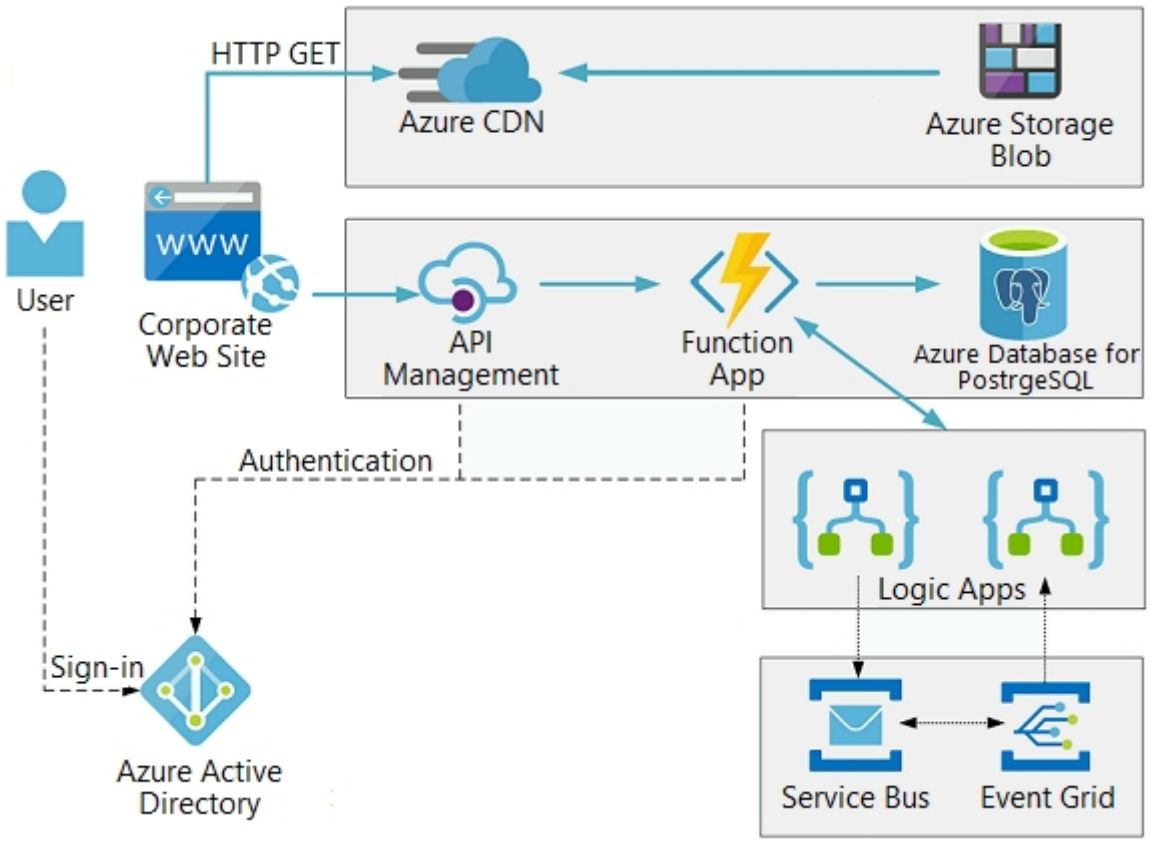

User authentication -

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -
Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

HOTSPOT -

You need to correct the Azure Logic app error message.

Which configuration values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Setting	Value
authentication level	anonymous function admin
managed identity	system-assigned user-assigned

Correct Answer:

Answer Area

Setting	Value
authentication level	anonymous **function** admin
managed identity	**system-assigned** user-assigned

Scenario: You test the Logic app in a development environment. The following error message displays:

'400 Bad Request'

Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Note: If the inbound call's request body doesn't match your schema, the trigger returns an HTTP 400 Bad Request error.

Box 1: function -
If you have an Azure function where you want to use the system-assigned identity, first enable authentication for Azure functions.

Box 2: system-assigned -
Your logic app or individual connections can use either the system-assigned identity or a single user-assigned identity, which you can share across a group of logic apps, but not both.
Reference:
https://docs.microsoft.com/en-us/azure/logic-apps/create-managed-service-identity

mlantonis Highly Voted 2 years, 5 months ago

Logic app:
You test the Logic app in a development environment. The following error message displays: '400 Bad Request'.
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Box 1: anonymous
To use your logic app's managed identity in your function, you must set your function's authentication level to anonymous. Otherwise, your logic app throws a "BadRequest" error.

Box 2: system-assigned
Your logic app or individual connections can use either the system-assigned identity or a single user-assigned identity, which you can share across a group of logic apps, but not both. On your logic app menu, under Settings, select Identity > System assigned

upvoted 42 times

mlantonis 2 years, 5 months ago

Reference:

https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-azure-functions#set-up-anonymous-authentication-in-your-function

upvoted 7 times

rqb11 Highly Voted 2 years, 6 months ago

According to this article Function authz should be set to anonymous as we're using AAD auth. https://adatum.no/azure/azure-ad-authentication-in-azure-functions

upvoted 33 times

micro9000 11 months, 3 weeks ago

agree on this based on the docs,
For your function to use your logic app's managed identity, you must set your function's authentication level to anonymous. Otherwise, your logic app workflow throws a BadRequest error.

upvoted 1 times

anastakasim 2 years, 6 months ago

Not function?

upvoted 1 times

ZodiaC 2 years, 3 months ago

NOPE READ LINK !

upvoted 4 times

lmass 2 years, 6 months ago

Agree, see also this: https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-azure-functions#enable-authentication-for-functions

upvoted 4 times

RaviNikkam 1 year, 5 months ago

Look at line RA05 which clearly mentions, authorization level as function.

upvoted 2 times

gmishra88 1 year ago

That's why it is giving the error and that need to be solved

upvoted 1 times

xRiot007 1 year, 4 months ago

That is for the function in the function app, not the logic app.

upvoted 1 times

gmishra88 Most Recent 1 year ago

For some weird reason Microsoft decided that logic apps cannot use the user-assigned identity when it is given to multiple. Totally random restriction that shows the inter-departmental issues in this organization. And they think that fault-line is a good question to ask.

upvoted 3 times

ChrisEvans 1 year, 6 months ago

We are fixing the logic app, not the function app. If you look at the function app code, you can see it's using function level authentication, so it can't be anonymous. I think it's function + system assigned.

upvoted 1 times

gmishra88 1 year, 1 month ago

I do agree the question is not clear. But if you think deep enough you will see that it is the function app that you need to change for the authorization (Box 1) and in the logic app to use the user-assigned-identity

upvoted 2 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer:
Anonymous
System assigned identity

upvoted 4 times

chingdm 1 year, 8 months ago

Can use combination of auth level=Function + managed identity=user-assigned, this works if you set the Function, Logic App and API Management to use this same user-assigned identity.
otherwise, the simpler approach but less secure is to have auth level=Anonymous set to Function and the Logic App to have the system-assigned managed identity.
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview

upvoted 1 times

edengoforit 1 year, 9 months ago

Why system-assigned?

upvoted 1 times

Jonas_86 2 years, 1 month ago

Anonymous ans system-assigned
check : https://docs.microsoft.com/fr-fr/azure/spring-cloud/tutorial-managed-identities-functions

upvoted 1 times

ning 2 years, 2 months ago

Either system or user identity works in this scenario ...
Without further requirements, it is hard to say ...

upvoted 1 times

ning 2 years, 2 months ago

Only thing I can see "All user accounts and processes must receive only those privileges which are essential to perform their intended function."
maybe system identity then ...

upvoted 3 times

jkes80 2 years, 3 months ago

I think I would answer
Box1: anonymous
Box2: user-assigned -> there are 2 Logic Apps in the picture, so it makes sense to me to create a user-assigned identity that is being used for both Logic Apps?

upvoted 2 times

ZodiaC 2 years, 3 months ago

Nope, thats make it difficult... So stay with system!

upvoted 1 times

anandhprakash 2 years, 5 months ago

Anonymous
System assigned identity

will be correct answer? I am guessing

upvoted 4 times

SnakePlissken 2 years, 5 months ago

Logic App is removed from the study guide on March 26, 2021.
https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B

upvoted 3 times

gmishra88 1 year, 1 month ago

Logic app is removed, but function app is not and this question is about the function app. So, unfortunately this question is still in scope

upvoted 1 times

if54uran 2 years, 4 months ago

It seems to still in ther o_O --> page 6 lower left side "Develop an App Service Logic App"
Am I missing something here?

upvoted 1 times

hstml 2 years, 1 month ago

The left side of the table states the things that are retired from the exam. "Previous Study Guide" :)
So SnakePlissken is correct.

upvoted 3 times

oadigun 2 years, 5 months ago

Answer is Anonymous....To use your logic app's managed identity in your function, you must set your function's authentication level to anonymous. Otherwise, your logic app throws a "BadRequest" error.
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-azure-functions#set-up-anonymous-authentication-in-your-function

upvoted 7 times

Question #6 *Topic 13*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

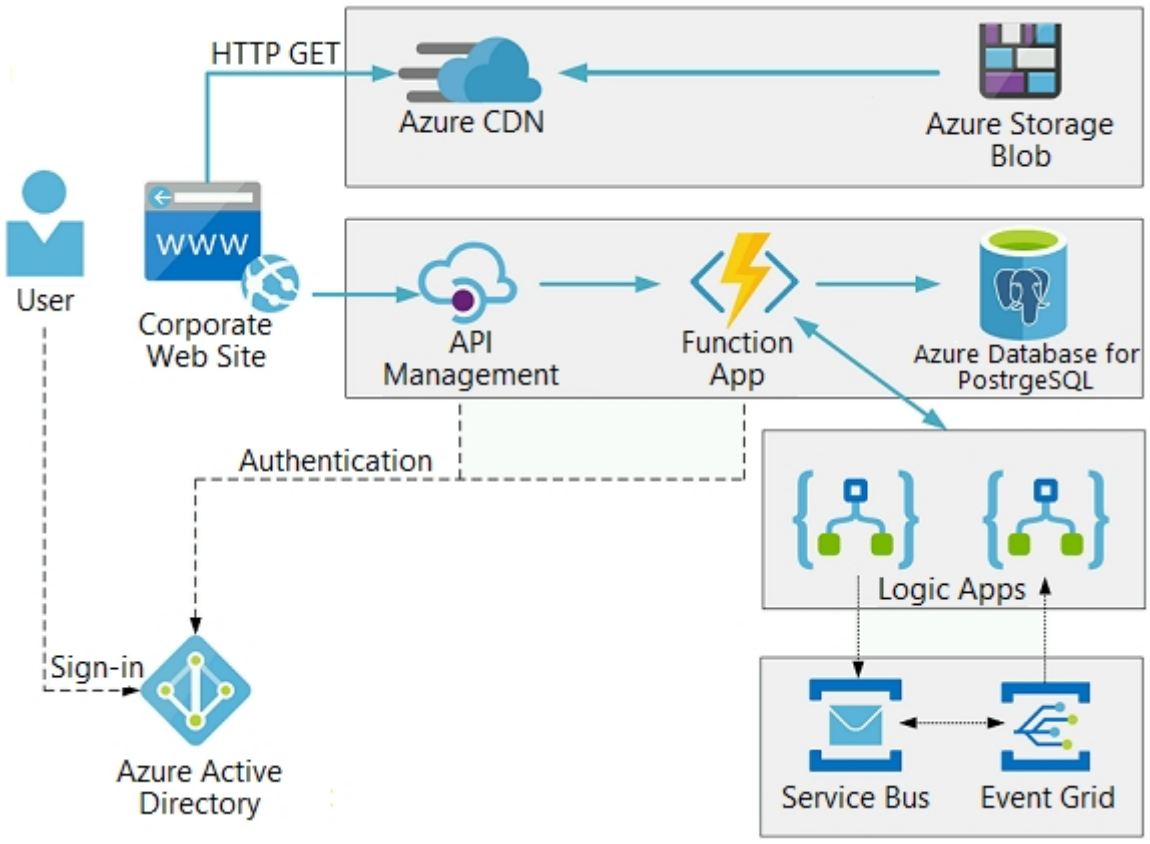

User authentication -

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -
Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

HOTSPOT -

You need to configure Azure Service Bus to Event Grid integration.

Which Azure Service Bus settings should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Setting	Value
Tier	Basic / Standard / Premium
RBAC role	Owner / Contributor / Azure Service Bus Data Owner / Azure Service Bus Data Receiver

Correct Answer:

Answer Area

Setting	Value
Tier	Basic / Standard / **Premium**
RBAC role	Owner / **Contributor** / Azure Service Bus Data Owner / Azure Service Bus Data Receiver

Box 1: Premium -

Service Bus can now emit events to Event Grid when there are messages in a queue or a subscription when no receivers are present. You can create Event Grid subscriptions to your Service Bus namespaces, listen to these events, and then react to the events by starting a receiver. With

this feature, you can use Service
Bus in reactive programming models.
To enable the feature, you need the following items:
A Service Bus Premium namespace with at least one Service Bus queue or a Service Bus topic with at least one subscription.
Contributor access to the Service Bus namespace.

Box 2: Contributor -
Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept

mlantonis Highly Voted 2 years, 5 months ago

To enable the feature, you need the following items:
- A Service Bus Premium namespace with at least one Service Bus queue or a Service Bus topic with at least one subscription.
- Contributor access to the Service Bus namespace.

Box 1: Premium

Box 2: Contributor

Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept

upvoted 31 times

Marusyk Highly Voted 2 years, 7 months ago

The answer is correct

upvoted 27 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with given answer, score 926

upvoted 1 times

TonyMel 7 months, 1 week ago

correct, in 2023Mar24, score: 904/1000.

upvoted 1 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 2 times

gmishra88 1 year, 1 month ago

RBAC role for what the service-bus, event-grid of "you"? I hope Microsoft does not ask such vague questions directly taken from their documentations

upvoted 2 times

gmishra88 1 year ago

"You should have contributor access to setup this". And also service bus should be premium. But what to setup is the issue. The problem with microsoft question is they do not tell what is the functionality that needs to be setup. Just that Azure service bus to Event grid integration. It could be for load-levelling, it could be for "Service Bus can emit events to Event Grid when there are messages in a queue or a subscription when no receivers are present". For the first use, you do not need to do this. It's a guessing game

upvoted 3 times

gmishra88 1 year ago

All these requirements are for integration, but then I guess they are asking for the third case when guessing the options available:
"Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing."

upvoted 1 times

gmishra88 1 year, 1 month ago

"You need to configure Azure Service Bus to Event Grid integration" < for what? there could be many integration possibilities but what exactly is the requirement. Or is Microsoft looking for the page with a similar title as that. Requirements should be functional. Integrate it for what purpose.

upvoted 1 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 5 times

MrXBasit 2 years, 3 months ago

Correct

upvoted 2 times

syedaquib77 2 years, 6 months ago

Azure Service Bus to Event Grid integration overview

To enable the feature, you need the following items:

A Service Bus *Premium* namespace with at least one Service Bus queue or a Service Bus topic with at least one subscription.

Contributor access to the Service Bus namespace. Navigate to your Service Bus namespace in the Azure portal, and then select Access control (IAM), and select Role assignments tab. Verify that you have the contributor access to the namespace.

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept?tabs=event-grid-event-schema

upvoted 7 times

ferut 2 years, 5 months ago

Tier = Premium, because the requirement says "Data must be replicated to a secondary region and three availability zones". Standard and basic don't support Geo Disasster Recover and Availability Zones.

upvoted 15 times

rqb11 2 years, 6 months ago

Premium tier includes Geo-Disaster Recovery (Geo-DR): https://azure.microsoft.com/en-us/pricing/details/service-bus/

upvoted 4 times

atomicicebreaker 2 years, 6 months ago

Microsoft for sure likes your mindset ;) You are correct, but the issue in this question is Service Bus <-> Event Grid communication, not availability or data recovery.

upvoted 4 times

Topic 14 - Testlet 16

Question #1 *Topic 14*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web

API project that is deployed as an Azure Web App.

Overall architecture -

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing -

Employees may upload receipts in two ways:

Uploading using an Azure Files mounted folder

Uploading using the web application

Data Storage -

Receipt and employee information is stored in an Azure SQL database.

Documentation -

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for

Azure File upload, and instructions on how to configure the mounted folder.

Solution details -

Users table -

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application -

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing -

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored

in Azure Blob
Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging -
Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled.
Application Insights must always contain all log messages.

Requirements -

Receipt processing -
Concurrent processing of a receipt must be prevented.

Disaster recovery -
Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security -
User's SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
Receipt data must always be encrypted at rest.
All data must be protected in transit.
User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues -

Upload format issue -
Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File
Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal
Server error page.

Capacity issue -
During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

Log capacity issue -
Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Application code -

Processing.cs -

```
PC01 public static class Processing
PC02 {
PC03   public static class Function
PC04   {
PC05     [FunctionName("IssueWork")]
PC06     public static async Task Run([TimerTrigger("0 */5 * * * *")] TimerInfo timer, ILogger
log)
PC07     {
PC08       var container = await GetCloudBlobContainer();
PC09       foreach (var fileItem in await ListFiles())
PC10       {
PC11         var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12         var ms = new MemoryStream();
PC13         await file.DownloadToStreamAsync(ms);
PC14         var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15         await blob.UploadFromStreamAsync(ms);
PC16
PC17       }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob(CloudBlockBlob sourceBlob)
PC20     {
PC21        . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25       var cloudBlobClient = new CloudBlobClient(new Uri(". . ."), await GetCredentials());
PC26
PC27       await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28       return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32        . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36        . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(". . .");
PC38   }

PC39 }
```

Database.cs -

```
DB01 public class Database
DB02 {
DB03   private string ConnectionString =
DB04
DB05   public async Task<object> LoadUserDetails(string userId)
DB06   {
DB07
DB08     return await policy.ExecuteAsync(async () =>
DB09     {
DB10       using (var connection = new SqlConnection(ConnectionString))
DB11       {
DB12         await connection.OpenAsync();
DB13         using (var command = new SqlCommand("…", connection))
DB14         using (var reader = command.ExecuteReader())
DB15         {
DB16           …
DB17         }
DB18       }
DB19     });
DB20   }

DB21 }
```

ReceiptUploader.cs -

```
RU01 public class ReceiptUploader
RU02 {
RU03   public async Task UploadFile(string file, byte[] binary)
RU04   {
RU05     var httpClient = new HttpClient();
RU06     var response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU07     while (ShouldRetry(response))
RU08     {
RU09       response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU10     }
RU11   }
RU12   private bool ShouldRetry(HttpResponseMessage response)
RU13   {
RU14
RU15   }
RU16 }
```

ConfigureSSE.ps1 -

```
CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "..." -AccountName "..."
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "..."
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "..."
CS04 Set-AzureRmKeyVaultAccessPolicy `
CS05  -VaultName $keyVault.VaultName `
CS06  -ObjectId $storageAccount.Identity.PrincipalId `
CS07
CS08
CS09 Set-AzureRmStorageAccount `
CS10  -ResourceGroupName $storageAccount.ResourceGroupName `
CS11  -AccountName $storageAccount.StorageAccountName `
CS12  -EnableEncryptionService File `
CS13  -KeyvaultEncryption `
CS14  -KeyName $key.Name
CS15  -KeyVersion $key.Version `
CS16  -KeyVaultUri $keyVault.VaultUri
```

Question

HOTSPOT -

You need to add code at line PC26 of Processing.cs to ensure that security policies are met.

How should you complete the code that you will add at line PC26? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
var resolver = new KeyVaultKeyResolver(_keyVaultClient);
var keyBundle = await _keyVaultClient.GetKeyAsync("...", "...");
```

var key = keyBundle.Key;
var key = keyBundle.KeyIdentifier.Identifier;
var key = await resolver.ResolveKeyAsync("encrypt", null);
var key = await resolver.ResolveKeyAsync(keyBundle.KeyIdentifier.Identifier, CancellationToken.None);

var x = keyBundle.Managed;
var x = AuthenticationScheme.SharedKey;
var x = new BlobEncryptionPolicy(key, resolver);
var x = new DeleteRetentionPolicy {Enabled = key.Kid != null};

cloudBlobClient.AuthenticationScheme = x;
cloudBlobClient.DefaultRequestOptions.RequireEncryption = x;
cloudBlobClient.DefaultRequestOptions.EncryptionPolicy = x;
cloudBlobClient.SetServiceProperties(new ServiceProperties(deleteRetentionPolicy:x));

Correct Answer:

Answer Area

```
var resolver = new KeyVaultKeyResolver(_keyVaultClient);
var keyBundle = await _keyVaultClient.GetKeyAsync("...", "...");
```

var key = keyBundle.Key;
var key = keyBundle.KeyIdentifier.Identifier;
var key = await resolver.ResolveKeyAsync("encrypt", null);
var key = await resolver.ResolveKeyAsync(keyBundle.KeyIdentifier.Identifier, CancellationToken.None);

var x = keyBundle.Managed;
var x = AuthenticationScheme.SharedKey;
var x = new BlobEncryptionPolicy(key, resolver);
var x = new DeleteRetentionPolicy {Enabled = key.Kid != null};

cloudBlobClient.AuthenticationScheme = x;
cloudBlobClient.DefaultRequestOptions.RequireEncryption = x;
cloudBlobClient.DefaultRequestOptions.EncryptionPolicy = x;
cloudBlobClient.SetServiceProperties(new ServiceProperties(deleteRetentionPolicy:x));

Box 1: var key = await Resolver.ResolveKeyAsyn(keyBundle,KeyIdentifier.CancellationToken.None);

Box 2: var x = new BlobEncryptionPolicy(key,resolver);

Example:

// We begin with cloudKey1, and a resolver capable of resolving and caching Key Vault secrets.

BlobEncryptionPolicy encryptionPolicy = new BlobEncryptionPolicy(cloudKey1, cachingResolver);

client.DefaultRequestOptions.EncryptionPolicy = encryptionPolicy;

Box 3: cloudblobClient. DefaultRequestOptions.EncryptionPolicy = x;

Reference:

https://github.com/Azure/azure-storage-net/blob/master/Samples/GettingStarted/EncryptionSamples/KeyRotation/Program.cs

wumingshi Highly Voted 2 years, 7 months ago

The answer is correct

upvoted 29 times

gmishra88 Highly Voted 1 year ago

Also Microsoft says : you need not know .Net for this certification and proceed to ask questions deep in the dark corners of .Net libraries

upvoted 12 times

macobuzi 1 month, 3 weeks ago

To be honest, if you are an experienced programmer you can guess the right answer to this question by pure logic, no .Net code or even Azure knowledge is needed.

upvoted 3 times

AndySmith 1 week, 3 days ago

Exactly! Did it from 1st try, but originally was very confused with plenty of optins, LOL.
So, 2nd and 3rd steps are easy to connect. And in 2nd step you should use "key" from Step1, but it only occurs in options 3 and 4. Since, 4 is about deletion policy, which doesn't make sense - then go with 3 which is Encryption Policy.
And Encryption Policy is 1:1 mapping to last step 3 :)

upvoted 1 times

gmishra88 1 year ago

The requirement says "Receipt data must always be encrypted at rest." That might make a person think do not use client side encryption. But then proceeds to make the person guess from the answer choices that it is client side encryption. Why does this organization think this is a good way to ask questions?

upvoted 1 times

gmishra88 1 year ago

"Consider using service-side encryption features instead of client-side encryption. For more information about service-side encryption features, see Azure Storage encryption for data at rest."

upvoted 1 times

SaintBahamut Most Recent 1 year, 5 months ago

Answer is correct, checked that in code,
If we take second box as correct then only 4rd option from 1st box fits

upvoted 2 times

Yazhu 1 year, 10 months ago

Given Answer is correct.
Refer the below URL.
https://csharp.hotexamples.com/examples/Microsoft.WindowsAzure.Storage.Blob/BlobEncryptionPolicy/-/php-blobencryptionpolicy-class-examples.html

upvoted 2 times

UnknowMan 2 years, 5 months ago

BlobEncryptionPolicy accept "Ikey" on constructor

=> https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.blobencryptionpolicy.-ctor?view=azure-dotnet-legacy#Microsoft_Azure_Storage_Blob_BlobEncryptionPolicy__ctor_Microsoft_Azure_KeyVault_Core_IKey_Microsoft_Azure_KeyVault_Core_IKeyResolver_

(keyBundle.Key return a Microsoft.Azure.KeyVault.WebKey.JsonWebKey).

The answer is correct !

upvoted 9 times

UnknowMan 2 years, 5 months ago

BlobEncryptionPolicy accept "Ikey" on constructor

=> https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.blobencryptionpolicy.-ctor?view=azure-dotnet-legacy#Microsoft_Azure_Storage_Blob_BlobEncryptionPolicy__ctor_Microsoft_Azure_KeyVault_Core_IKey_Microsoft_Azure_KeyVault_Core_IKeyResolver_

(keyBundle.Key return a Microsoft.Azure.KeyVault.WebKey.JsonWebKey).

The answer is correct !

upvoted 1 times

SnakePlissken 2 years, 5 months ago

1. var key = keyBundle.Key;
2. var x = new BlobEncryptionPolicy(key, resolver);
3. cloudBlobClient.DefaultRequestOptions.EncryptionPolicy = x;

But I'm afraid I've wasted my time. As you can see in the links, it's all legacy code.

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.keyvault.keyvaultclientextensions.getkeyasync?view=azure-dotnet-legacy&viewFallbackFrom=azure-dotnet
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.blobencryptionpolicy.-ctor?view=azure-dotnet-legacy#Microsoft_Azure_Storage_Blob_BlobEncryptionPolicy__ctor_Microsoft_Azure_KeyVault_Core_IKey_Microsoft_Azure_KeyVault_Core_IKeyResolver_
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.blobrequestoptions.encryptionpolicy?view=azure-dotnet-legacy#Microsoft_Azure_Storage_Blob_BlobRequestOptions_EncryptionPolicy

upvoted 6 times

coffecold 1 year ago

yes legacy code..

for version 12 use KeyClient, ClientSideEncryptionOptions ...
see
https://stackoverflow.com/questions/64644174/encryption-with-azure-bob-storage-v12-sdk-for-net

upvoted 1 times

ning 2 years, 2 months ago

This is NOT correct, keyBundle.Key() is correct, it is a method, not a property, so no correct

upvoted 1 times

vokep77043 2 years, 2 months ago

No, it's a property. Read documenation. https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.keyvault.models.keybundle.key?view=azure-dotnet-legacy#Microsoft_Azure_KeyVault_Models_KeyBundle_Key

upvoted 3 times

jvyas 2 years, 5 months ago

According to Udemy instructor Alan Rodriguez answer for first box is
var key = keyBundle.key

upvoted 4 times

jokergester 2 years, 7 months ago

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-encrypt-decrypt-blobs-key-vault?tabs=dotnet11#encrypt-blob-and-upload

upvoted 4 times

Question #2 *Topic 14*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web

API project that is deployed as an Azure Web App.

Overall architecture -

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing -

Employees may upload receipts in two ways:

Uploading using an Azure Files mounted folder

Uploading using the web application

Data Storage -

Receipt and employee information is stored in an Azure SQL database.

Documentation -

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for

Azure File upload, and instructions on how to configure the mounted folder.

Solution details -

Users table -

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application -

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing -

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob

Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must

remain valid if the email is forwarded to another user.

Logging -
Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled.
Application Insights must always contain all log messages.

Requirements -

Receipt processing -
Concurrent processing of a receipt must be prevented.

Disaster recovery -
Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security -
User's SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
Receipt data must always be encrypted at rest.
All data must be protected in transit.
User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues -

Upload format issue -
Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File
Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal
Server error page.

Capacity issue -
During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

Log capacity issue -
Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Application code -

Processing.cs -

```
PC01 public static class Processing
PC02 {
PC03   public static class Function
PC04   {
PC05     [FunctionName("IssueWork")]
PC06     public static async Task Run([TimerTrigger("0 */5 * * * *")] TimerInfo timer, ILogger
log)
PC07     {
PC08       var container = await GetCloudBlobContainer();
PC09       foreach (var fileItem in await ListFiles())
PC10       {
PC11         var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12         var ms = new MemoryStream();
PC13         await file.DownloadToStreamAsync(ms);
PC14         var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15         await blob.UploadFromStreamAsync(ms);
PC16
PC17       }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob(CloudBlockBlob sourceBlob)
PC20     {
PC21        . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25       var cloudBlobClient = new CloudBlobClient(new Uri(". . ."), await GetCredentials());
PC26
PC27       await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28       return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32        . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36        . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(". . .");
PC38   }

PC39 }
```

Database.cs -

```
DB01 public class Database
DB02 {
DB03   private string ConnectionString =
DB04
DB05   public async Task<object> LoadUserDetails(string userId)
DB06   {
DB07
DB08     return await policy.ExecuteAsync(async () =>
DB09     {
DB10       using (var connection = new SqlConnection(ConnectionString))
DB11       {
DB12         await connection.OpenAsync();
DB13         using (var command = new SqlCommand("…", connection))
DB14         using (var reader = command.ExecuteReader())
DB15         {
DB16           …
DB17         }
DB18       }
DB19     });
DB20   }

DB21 }
```

ReceiptUploader.cs -

```
RU01 public class ReceiptUploader
RU02 {
RU03   public async Task UploadFile(string file, byte[] binary)
RU04   {
RU05     var httpClient = new HttpClient();
RU06     var response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU07     while (ShouldRetry(response))
RU08     {
RU09       response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU10     }
RU11   }
RU12   private bool ShouldRetry(HttpResponseMessage response)
RU13   {
RU14
RU15   }
RU16 }
```

ConfigureSSE.ps1 -

```
CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "..." -AccountName "..."
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "..."
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "..."
CS04 Set-AzureRmKeyVaultAccessPolicy `
CS05  -VaultName $keyVault.VaultName `
CS06  -ObjectId $storageAccount.Identity.PrincipalId `
CS07
CS08
CS09 Set-AzureRmStorageAccount `
CS10  -ResourceGroupName $storageAccount.ResourceGroupName `
CS11  -AccountName $storageAccount.StorageAccountName `
CS12  -EnableEncryptionService File `
CS13  -KeyvaultEncryption `
CS14  -KeyName $key.Name
CS15  -KeyVersion $key.Version `
CS16  -KeyVaultUri $keyVault.VaultUri
```

Question

You need to ensure the security policies are met.

What code do you add at line CS07 of ConfigureSSE.ps1?

A. ג€"PermissionsToKeys create, encrypt, decrypt

B. ג€"PermissionsToCertificates create, encrypt, decrypt

C. ג€"PermissionsToCertificates wrapkey, unwrapkey, get

D. ג€"PermissionsToKeys wrapkey, unwrapkey, get

Correct Answer: *B*

Scenario: All certificates and secrets used to secure data must be stored in Azure Key Vault.

You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.

The Set-AzureRmKeyValutAccessPolicy parameter -PermissionsToKeys specifies an array of key operation permissions to grant to a user or service principal.

The acceptable values for this parameter: decrypt, encrypt, unwrapKey, wrapKey, verify, sign, get, list, update, create, import, delete, backup, restore, recover, purge

Incorrect Answers:

A, C: The Set-AzureRmKeyValutAccessPolicy parameter -PermissionsToCertificates specifies an array of certificate permissions to grant to a user or service principal. The acceptable values for this parameter: get, list, delete, create, import, update, managecontacts, getissuers, listissuers, setissuers, deleteissuers, manageissuers, recover, purge, backup, restore

Reference:

https://docs.microsoft.com/en-us/powershell/module/azurerm.keyvault/set-azurermkeyvaultaccesspolicy

Community vote distribution

D (100%)

Zsolt72 Highly Voted 2 years, 7 months ago

My opinion is that the answer is D.

The policy should belong to a key. In the case study the code retrieve the key so the GET access policy is mandatory. The wrap/unwrap is used for symmetric encryption and in this case study the task is to encrypt the blobs.

upvoted 54 times

azurelearner666 2 years, 4 months ago

Yes! it's D
PermissionsToKeys wrapkey, unwrapkey, get
https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccount?view=azps-5.8.0#example-5--set-encryption-keysource-to-keyvault
code example at line 7

upvoted 11 times

mlantonis Highly Voted 2 years, 5 months ago

Correct Answer: D

PS C:\>Set-AzKeyVaultAccessPolicy -VaultName "MyKeyVault" -ObjectId $account.Identity.PrincipalId -PermissionsToKeys wrapkey,unwrapkey,get

Reference:

https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccount?view=azps-5.8.0#example-5--set-encryption-keysource-to-keyvault

upvoted 12 times

alejary Most Recent 6 months, 1 week ago

What is ג€" ?

upvoted 1 times

adilkhan 6 months, 1 week ago

are these questions enough to pass the exam?

upvoted 2 times

dddddd111 1 week, 4 days ago

Some questions here appeared to the actual exam. But the problem is the answer here are not accurate. Same in some highly voted answers. I failed on my first attempt (646/1000) even though I have contributor access.

upvoted 1 times

uffuchsi 8 months, 1 week ago

Selected Answer: D

100% D - All certificates and secrets used to secure data must be stored in Azure Key Vault.

You need to retrieve the keys so get permission is required. The wrapkey and unwrapkey will be used for symmetric encryption to encrypt the blobs.

Below link contains an example of same scenario.

https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccount?view=azps-8.0.0#example-5-set-encryption-keysource-to-keyvault

https://docs.microsoft.com/en-us/azure/key-vault/keys/about-keys-details#key-access-control

upvoted 3 times

AlexeyG 8 months, 2 weeks ago

Got this in 16/02/2023

upvoted 1 times

mabdo 8 months, 2 weeks ago

same case study?

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: D

PermissionsToKeys wrapkey, unwrapkey, get

upvoted 1 times

gmishra88 1 year ago

This page shows the example: https://learn.microsoft.com/en-us/powershell/module/azurerm.storage/set-azurermstorageaccount?view=azurermps-6.13.0#example-5-set-encryption-keysource-to-keyvault

upvoted 1 times

aruni_mishra 1 year, 4 months ago

Selected Answer: D

https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccount?view=azps-8.0.0&viewFallbackFrom=azps-5.8.0#example-5--set-encryption-keysource-to-

keyvault:~:text=PS%20C%3A%5C%3E%20Set%2DAzKeyVaultAccessPolicy%20%2DVaultName%20%24keyvaultName%20%2DResourceGroupName%20%24resourceGroupName%20%2DObjectId%20%24userId.PrincipalId%20%2DPermissionsToKeys%20get%2Cwrapkey%2Cunwrapkey%20%2DBypassObjectIdValidation

upvoted 1 times

pandrer 1 year, 6 months ago

B and C wrong parameters

upvoted 1 times

Bogdan75 1 year, 7 months ago

Selected Answer: D

https://docs.microsoft.com/en-us/azure/key-vault/keys/about-keys-details#key-access-control

upvoted 1 times

leonidn 1 year, 9 months ago

Selected Answer: D

PermissionsToKeys wrapkey, unwrapkey, get

upvoted 4 times

edengoforit 1 year, 9 months ago

The Set-AzureRmKeyValutAccessPolicy parameter -PermissionsToKeys specifies an array of key operation permissions to grant to a user or service principal.
According to the reference, the answer is D

https://docs.microsoft.com/es-es/powershell/module/azurerm.storage/set-azurermstorageaccount?view=azurermps-6.13.0

upvoted 1 times

RajMasilamani 2 years, 1 month ago

Answer is D.
Wrap,Unwrap,encrypt,decrypt available only for -PermissionsToKeys
https://docs.microsoft.com/en-us/powershell/module/az.keyvault/set-azkeyvaultaccesspolicy?view=azps-6.4.0#parameters

upvoted 3 times

ReniRechner 1 year, 8 months ago

This site also clearly states that PermissionsToCertificates only has these options:
all, get, list, delete, create, import, update, managecontacts, getissuers, listissuers, setissuers, deleteissuers, manageissuers, recover, purge, backup, restore

So B and C are not even valid

upvoted 1 times

ning 2 years, 2 months ago

D is correct from https://docs.microsoft.com/es-es/powershell/module/azurerm.storage/set-azurermstorageaccount?view=azurermps-6.13.0

upvoted 2 times

anandhprakash 2 years, 5 months ago

Refer
https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccount?view=azps-5.8.0#example-5--set-encryption-keysource-to-keyvault
Example 5: Set Encryption KeySource to Keyvault

Answer should be D: wrapkey,unwrapkey,get

PS C:\>Set-AzKeyVaultAccessPolicy -VaultName "MyKeyVault" -ObjectId $account.Identity.PrincipalId -PermissionsToKeys wrapkey,unwrapkey,get

upvoted 3 times

[Removed] 2 years, 6 months ago

D is correct:

https://docs.microsoft.com/en-us/powershell/module/az.storage/set-azstorageaccount?view=azps-5.8.0#example-5--set-encryption-keysource-to-keyvault

upvoted 5 times

Topic 15 - Testlet 17

Question #1 *Topic 15*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -
The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
You need to reduce read latency for the retail store solution.
What are two possible ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Create a new composite index for the store location data queries in Azure Cosmos DB. Modify the queries to support parameterized SQL and update the Azure Function app to call the new queries.

B. Provision an Azure Cosmos DB dedicated gateway. Update the Azure Function app connection string to use the new dedicated gateway endpoint.

C. Configure Azure Cosmos DB consistency to session consistency. Cache session tokens in a new Azure Redis cache instance after every write. Update reads to use the session token stored in Azure Redis.

D. Provision an Azure Cosmos DB dedicated gateway. Update blob storage to use the new dedicated gateway endpoint.

E. Configure Azure Cosmos DB consistency to strong consistency. Increase the RUs for the container supporting store location data.

Correct Answer: *BC*
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.
B: A dedicated gateway is server-side compute that is a front-end to your Azure Cosmos DB account. When you connect to the dedicated gateway, it both routes requests and caches data.
You can provision a dedicated gateway to improve performance at scale.
You must connect to Azure Cosmos DB using the dedicated gateway in order to use the integrated cache. The dedicated gateway has a different endpoint from the standard one provided with your Azure Cosmos DB account. When you connect to your dedicated gateway endpoint, your application sends a request to the dedicated gateway, which then routes the request to different backend nodes. If possible, the integrated cache will serve the result.
C: Azure Cache for Redis perfectly complements Azure database services such as Cosmos DB. It provides a cost-effective solution to scale read and write throughput of your data tier. Store and share database query results, session states, static contents, and more using a common cache-aside pattern.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/data-cache-with-redis-cache https://docs.microsoft.com/en-us/azure/cosmos-db/dedicated-gateway

Community vote distribution
AC (44%) AB (44%) 11%

AbdulMannan Highly Voted 1 year, 1 month ago
Got this question on 30-Sep-2022 exam.
Correct answer is A and B. Passed with 870 score.
upvoted 14 times

ReyPirata 2 months, 1 week ago
Correct answer is A and B.
This was on the exam (08/20/2023). Scored 925

upvoted 1 times

surprise0011 6 months, 2 weeks ago

I think is correct. If your wondering about B: https://learn.microsoft.com/en-us/azure/cosmos-db/dedicated-gateway

And some info from ChatGpt:
A dedicated gateway provides a dedicated endpoint for client applications to communicate with Azure Cosmos DB. By using a dedicated gateway, client applications such as the Azure Function app can reduce network latency and improve overall performance when communicating with Azure Cosmos DB.

When a dedicated gateway is provisioned, it creates a Virtual Network (VNet) peering between the gateway and the Azure Cosmos DB account. This ensures that all traffic between the client application and Azure Cosmos DB remains within the same network, reducing the network latency.

upvoted 1 times

finnishr Highly Voted 1 year, 1 month ago

C makes no sense... so A and B

upvoted 8 times

cecho123 1 year, 1 month ago

"Cache session tokens in a new Azure Redis cache instance after every write" in C. Does not sound right to me either.

upvoted 2 times

EliteAllen Most Recent 3 months ago

Selected Answer: AB

A & B
A. Create a new composite index for the store location data queries in Azure Cosmos DB. Modify the queries to support parameterized SQL and update the Azure Function app to call the new queries:
By creating a composite index tailored to the specific queries used, you can potentially reduce the query cost and improve performance. Parameterized queries can further improve efficiency.

B. Provision an Azure Cosmos DB dedicated gateway. Update the Azure Function app connection string to use the new dedicated gateway endpoint:
A dedicated gateway in Azure Cosmos DB provides improved performance for query execution and can reduce latency. Connecting the Azure Function app to this dedicated gateway can leverage these performance improvements.

upvoted 2 times

phucngueyn 5 months ago

AB is correct

upvoted 1 times

winterthor4 7 months, 1 week ago

Got this vanarsdelltd case study on 26-Mar-2023 exam. Go with A and C. Score 890.

upvoted 1 times

adilkhan 7 months, 3 weeks ago

A and C. ChatGPT

upvoted 1 times

comoon 8 months, 4 weeks ago

Selected Answer: AC

A and C

upvoted 1 times

rotimislaw 9 months, 4 weeks ago

Selected Answer: AC

A and C
B & D are off due to Azure Cosmos DB dedicated gateway not supporting SQL queries. See limitations on:
https://learn.microsoft.com/en-us/azure/cosmos-db/dedicated-gateway#dedicated-gateway-in-multi-region-accounts
E. Makes no sense as increasing consistency to Strong only makes the reads slower

Thus A & C which makes sense: Composite Index in Cosmos DB and Redis Cache for reads

upvoted 3 times

NK203 8 months, 1 week ago

https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/query/getting-started
In Azure Cosmos DB for NoSQL accounts, there are two ways to read data： Point reads& SQL queries
https://learn.microsoft.com/en-us/azure/cosmos-db/dedicated-gateway#dedicated-gateway-in-multi-region-accounts
Dedicated gateways are only supported on API for NoSQL accounts
So dedicated gateway is support NoSQL account.How do you found that Azure Cosmos DB dedicated gateway not supporting SQL queries?

upvoted 3 times

toysky731 10 months, 1 week ago

Selected Answer: AB

composite index improves the sql performace

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: BC

As per the description B and C look correct

upvoted 1 times

gmishra88 1 year, 1 month ago

Redis is clearly a correct option if using Session consistency. From the following document it is clear that the session-id can be used in a cookie for multiple client-instances to share the same session. That is a clear case for redis-cache. So, if Session consistency is correct then redis cache is also correct. https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/data-cache-with-redis-cache

upvoted 2 times

gmishra88 1 year ago

"Parameterized SQL provides robust handling and escaping of user input, and prevents accidental exposure of data through SQL injection"
So, it does not give performance improvements. Maybe a trick to confuse the traditional sql guys (at least the Java ones, who use that for performance). A trick to punish the Java guys (non-microsoft, I guess. Don't be evil, Microsoft).
Dedicated gateway is a fine option with integrated cache. But otherwise increases latency with an extra hop if using a gateway (either standard or dedicated).
So, another Microsoft gem to create total confusion. An ill conceived question that is half baked

upvoted 1 times

warchoon 6 months, 3 weeks ago

Of course parametrization gives performance when it uses cached query plans.

upvoted 1 times

gmishra88 1 year ago

"Connecting to Azure Cosmos DB with the dedicated gateway provides lower and more predictable latency than connecting to Azure Cosmos DB with the standard gateway." But worse than a direct connection unless the cache is used. Probably that is the guessing game we have to do with a lot of assumption thinking that the location data can be cached.

upvoted 1 times

gmishra88 1 year ago

I will bet on B (dedicated gateway) and E (Strong consistency with increased RUs). Strong consistency because that will make the writes do synchronous and so reads can be with lesser consistency level.
This is very innovative answer, but then I can only hope. It cannot be session consistency. It cannot be parameterized sql (unless the Microsoft guys does not understand that it does not increase performance)

upvoted 2 times

gmishra88 1 year, 1 month ago

B and C are correct.
Not using Strong consistency (or bounded staleness is good for read latency), I
Redis cache can be used to store session information. But that is about web sessions. But I guess this also can be stored in Redis.

upvoted 3 times

gmishra88 1 year ago

Strong consistency as default will make sure the writes are synchronized. But this is about the read latency. While reading any consistency less than Bounded staleness will be better for read latency (read on read quoroms).

upvoted 1 times

le129 1 year, 1 month ago

why not A composite index

upvoted 4 times

adilkhan 8 months, 4 weeks ago

A = https://learn.microsoft.com/en-us/azure/cosmos-db/index-policy is correct

upvoted 1 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
You need to audit the retail store sales transactions.
What are two possible ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Update the retail store location data upload process to include blob index tags. Create an Azure Function to process the blob index tags and filter by store location.

B. Process the change feed logs of the Azure Blob storage account by using an Azure Function. Specify a time range for the change feed data.

C. Enable blob versioning for the storage account. Use an Azure Function to process a list of the blob versions per day.

D. Process an Azure Storage blob inventory report by using an Azure Function. Create rule filters on the blob inventory report.

E. Subscribe to blob storage events by using an Azure Function and Azure Event Grid. Filter the events by store location.

Correct Answer: *BE*
Scenario: Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.
"Process the change feed logs of the Azure Blob storage account by using an Azure Function. Specify a time range for the change feed data": Change feed support is well-suited for scenarios that process data based on objects that have changed. For example, applications can:
Store, audit, and analyze changes to your objects, over any period of time, for security, compliance or intelligence for enterprise data management.
"Subscribe to blob storage events by using an Azure Function and Azure Event Grid. Filter the events by store location": Azure Storage events allow applications to react to events, such as the creation and deletion of blobs. It does so without the need for complicated code or expensive and inefficient polling services. The best part is you only pay for what you use.
Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener. Event Grid provides reliable event delivery to your applications through rich retry policies and dead-lettering.
Incorrect Answers:
"Enable blob versioning for the storage account. Use an Azure Function to process a list of the blob versions per day": You can enable Blob storage versioning to automatically maintain previous versions of an object. When blob versioning is enabled, you can access earlier versions of a blob to recover your data if it is modified or deleted.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

Community vote distribution
BE (100%)

OPT_001122 Highly Voted 10 months, 3 weeks ago
Selected Answer: BE
B and E
upvoted 6 times

ReyPirata 2 months, 1 week ago
Correct B and E.
This was on the exam (08/20/2023). Scored 925
upvoted 1 times

motekim Highly Voted 7 months, 2 weeks ago

B. Process the change feed logs of the Azure Blob storage account by using an Azure Function. Specify a time range for the change feed data.
The change feed provides a way to log changes to blobs in a storage account. By using an Azure Function to process the change feed logs, it is possible to track and audit sales transaction information. The time range for the change feed data can be specified to capture the transactions within a specific time period.

E. Subscribe to blob storage events by using an Azure Function and Azure Event Grid. Filter the events by store location.
Azure Event Grid allows subscribing to events raised by Azure services or third-party services. By using an Azure Function and Event Grid, it is possible to filter the events for a specific store location and track sales transactions. This approach can help to monitor sales transactions in real-time and provide an audit trail for reconciliation.

upvoted 5 times

nardk Most Recent 1 month, 2 weeks ago

Selected Answer: BE

Must be B and E

upvoted 1 times

juanckar 3 months, 3 weeks ago

This was on the exam (July 2023). Went with highly voted. Scored 917

upvoted 2 times

gmishra88 1 year ago

Also note that it is possible to react to change in storage blob the following way:
* Blob storage trigger (Functions), They call this sometimes Blob storage trigger standard. This uses polling and latency is high
* Blob storage trigger event based (function). Uses event grid but not the same as event-grid-trigger
* Event grid trigger (function or others) Also called Storage Events: Not same as above. This uses an event grid subscription directly (for functions) and others as storage events that can be subscribed to
* Change feed : Avro and event grid format (see articles for details, I never used). Latency higher than events. Transactional log, ordered
* Azure storage blob inventory : Records of all read, write, list, and delete operations with successful and failed requests across all operations. Analytics logs are best-effort and no ordering is guaranteed. Not transactional log

upvoted 3 times

Mousavi 1 year, 1 month ago

why not D? can someone explain that?

upvoted 1 times

gmishra88 1 year, 1 month ago

Sorry, I did not read the question fully or forgot. It says clearly it needs to audit the transactions. The option to use Azure storage blob inventory does not give the transactions, but the properties for the blob containers. So, it is clearly not the option to audit the transactions. Hope it is clear

upvoted 4 times

gmishra88 1 year, 1 month ago

Unfortunately that cannot be explained because it is not clear what is being audited. Microsoft keeps that vague so that you can guess what the person who wrote this question is thinking. But if the audit is about content other two choices are fine.

upvoted 1 times

sghaha 1 year, 6 months ago

in Korean
https://docs.microsoft.com/ko-kr/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal

upvoted 3 times

mkahmann 1 month, 2 weeks ago

In Spanish
https://learn.microsoft.com/es-es/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal

Greetz, Señor developer :P

upvoted 1 times

Billabongs 12 months ago

in English
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal

upvoted 4 times

Younes364 1 year, 1 month ago

Sceen shots in English :D

upvoted 2 times

Topic 16 - Testlet 18

Question #1 *Topic 16*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -
You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         -  port: 80
CS12            protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14        "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

You need to monitor ContentUploadService according to the requirements.

Which command should you use?

A. az monitor metrics alert create ג€"n alert ג€"g ג€¦ - -scopes ג€¦ - -condition "avg Percentage CPU > 8"

B. az monitor metrics alert create ג€"n alert ג€"g ג€¦ - -scopes ג€¦ - -condition "avg Percentage CPU > 800"

C. az monitor metrics alert create ג€"n alert ג€"g ג€¦ - -scopes ג€¦ - -condition "CPU Usage > 800"

D. az monitor metrics alert create ג€"n alert ג€"g ג€¦ - -scopes ג€¦ - -condition "CPU Usage > 8"

Correct Answer: *B*

Scenario: An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Reference:

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

Community vote distribution

C (96%)

anirbanzeus Highly Voted 2 years, 4 months ago

C is the correct answer. We are dealing with containers here not VM so "CPU usage" is a valid condition. Had it been VM then it should have been "Percentage CPU usage". 800 is also correct since for containers its measured in millicores.
Ref : https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported#microsoftcontainerinstancecontainergroups

upvoted 46 times

beonsoft 1 year, 12 months ago

It is "avg" and not "usage". Reason: it is generalized as multicores could be there.
=>Answer C

upvoted 5 times

beonsoft 1 year, 12 months ago

...B...

upvoted 4 times

Illumielle 2 years ago

Where does it say contentuploadservice uses containers? The main problem with B is that 800 is not a percentage.

upvoted 2 times

Illumielle 2 years ago

It's in the code. There is still the problem with 800 not being a percentage

upvoted 1 times

MiraA 2 years ago

See ContentUploadService, line CS02.

upvoted 3 times

RaviNikkam 1 year, 5 months ago

Correct. 800 percent does not make sense. CPU usage of 800 seems to be correct. C is the answer.

upvoted 2 times

robjanssen Highly Voted 2 years, 7 months ago

Shouldn't it be > 80 (instead of 800)?

upvoted 21 times

insanewriters 2 years, 4 months ago

The CPU Usage measurement is in milicores (1/1000 of a core). So, 80% of a core would be 800 milicores.

upvoted 23 times

Pozz4ever 2 years, 7 months ago

agree with you

upvoted 2 times

SaNagh 2 years, 2 months ago

The CPU usage measurement is in milicores. 80% of a core would be 800 milicores.

upvoted 11 times

uffuchsi Most Recent 8 months, 1 week ago

B - Scenario: An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU-cores

Azure Monitor provides the following metrics for Azure Container Instances. These metrics are available for a container group and individual containers. By default, the metrics are aggregated as averages.

CPU Usage - measured in millicores. One millicore is 1/1000th of a CPU core, so 500 millicores represents usage of 0.5 CPU core.

Memory Usage - in bytes.

Network Bytes Received Per Second and Network Bytes Transmitted Per Second.

https://docs.microsoft.com/en-us/azure/container-instances/container-instances-monitor

upvoted 1 times

warchoon 6 months, 3 weeks ago

Usage measured in millicores, percentage does not
https://learn.microsoft.com/en-us/cli/azure/monitor/metrics/alert?view=azure-cli-latest#:~:text=%2D%2Dcondition%20%22-,avg%20Percentage%20CPU%20%3E%2090,-%22%20%2D%2Dwindow%2Dsize

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 2 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This test case was on it, but not this specific question.

upvoted 2 times

Knightie 1 year, 2 months ago

All wrong.
1. It needs to be avg CPU, it doen't make sense for sudden spike of any individual computation, but if Average is over, it needs the alert.
2. I don't see the syntax "CPU Usage" and the words do not match the percentage.
3. Percentage > 80, not 8 nor 800, see from the official example.

az monitor metrics alert create -n alert1 -g {ResourceGroup} --scopes {VirtualMachineID1} {VirtualMachineID2} {VirtualMachineID3} \
--condition "avg Percentage CPU > 90" --description "High CPU" --region westus

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: C

C is correct answer.

upvoted 4 times

xRiot007 1 year, 4 months ago

We want an alert for an exceeded limit.
A,B are invalid from the start, we do not care about averages.
D has a wrong number in the evaluation expression
C is the correct answer - that "800" is units per thousand. 800/1000 is 80%

upvoted 3 times

Sandeep12093 1 year, 4 months ago

ContentUploadService is in ACI.
when reading Microsft document found that CPU Usage is used for ACI and Avg percentage is used for VM

Also we don't want average just 80% of available CPU
SO ANSWER IS C

upvoted 2 times

Elsheimy 1 year, 5 months ago

According to
https://docs.azure.cn/zh-cn/cli/monitor/metrics/alert?view=azure-cli-latest
The answer is B. However, B states 800, I don't know if it's a typo

upvoted 1 times

pandrer 1 year, 6 months ago

Selected Answer: B

similar alerts on https://docs.microsoft.com/en-us/cli/azure/monitor/metrics/alert?view=azure-cli-latest#az-monitor-metrics-alert-create

upvoted 1 times

ReniRechner 1 year, 8 months ago

Selected Answer: C

Unit avg CPU: per hundred (%)
Unit CPU: per thousand
A:8%
B:800%
C:80%
D:0.8%

upvoted 8 times

Netspud 1 year, 8 months ago

Selected Answer: C

It is a ACI not a VM so "CPU usage" is correct

upvoted 4 times

heisenberg33 1 year, 8 months ago

Selected Answer: C

I believe its C based on this Ref: https://docs.microsoft.com/en-us/azure/container-instances/container-instances-monitor

upvoted 4 times

asdasdasg2 1 year, 9 months ago

It's not A or B because we don't care about average usage, we want an alert when the usage goes above 80%. Therefore it mus be C, because as others have stated, CPU usage measurement is in millicores, so 800 would be 80%.

Correct answer: C

upvoted 3 times

gfiorini 1 year, 11 months ago

"The ContentAnalysisService is deployed with Azure Container Instances"
"The solution will use eight CPU cores."
CPU usage is measured in millicore (https://docs.microsoft.com/en-us/azure/container-instances/container-instances-monitor)
so correct answer should be 'C' (80% = 800millicore)

upvoted 2 times

mlantonis 2 years, 5 months ago

Monitoring:
An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.
The metric "CPU Usage" does not exist, avg Percentage CPU is valid.

Note: It should be “avg Percentage CPU > 80”

Reference:

https://docs.microsoft.com/en-us/cli/azure/monitor/metrics/alert?view=azure-cli-latest#az_monitor_metrics_alert_create

upvoted 9 times

Renwa 2 years, 4 months ago

As per Udemy, C is the correct answer - https://docs.microsoft.com/en-us/azure/container-instances/container-instances-monitor

upvoted 3 times

Netspud 1 year, 8 months ago

Because it is an ACI "CPU Usage" is valid. If it were a VM then it would not be. In fact "avg Percentage CPU" is not valid for an ACI. Check the various links in the comments, all the ones I checked point to this.

upvoted 1 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -

You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14        "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

You need to investigate the http server log output to resolve the issue with the ContentUploadService.

Which command should you use first?

A. az webapp log

B. az ams live-output

C. az monitor activity-log

D. az container attach

Correct Answer: *C*

Scenario: Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

"502 bad gateway" and "503 service unavailable" are common errors in your app hosted in Azure App Service.

Microsoft Azure publicizes each time there is a service interruption or performance degradation.

The az monitor activity-log command manages activity logs.

Note: Troubleshooting can be divided into three distinct tasks, in sequential order:

1. Observe and monitor application behavior
2. Collect data
3. Mitigate the issue

Reference:

https://docs.microsoft.com/en-us/cli/azure/monitor/activity-log

Community vote distribution

Whirly Highly Voted 2 years, 6 months ago

Answer is az container attach

Ref: https://docs.microsoft.com/en-us/azure/container-instances/container-instances-get-logs

upvoted 55 times

ray01 2 years, 6 months ago

Correct. If one will take a look at "CS09" line, it's clear that "ContentUploadService" is a container, so the first step should be "az container attach"

upvoted 11 times

maukaba 1 year, 11 months ago

A web app can be deployed both as source code and as a docker container.
C it's correct see link https://docs.microsoft.com/en-us/azure/app-service/troubleshoot-http-502-http-503

upvoted 1 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose： D

upvoted 3 times

maukaba 1 year, 11 months ago

C it's correct see link https://docs.microsoft.com/en-us/azure/app-service/troubleshoot-http-502-http-503

upvoted 2 times

Zidimirite 2 years, 6 months ago

It's not deployed to a container. This is a regular webapp so A, az webapp log is the correct answer.

upvoted 3 times

somenkr 2 years, 3 months ago

It is , Read the question again.

upvoted 1 times

clarionprogrammer 2 years, 6 months ago

It is. Note the type in the ContentUploadService log.

upvoted 1 times

SnakePlissken 2 years, 5 months ago

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

upvoted 8 times

DV007 1 year, 9 months ago

The question is about the Content*Upload*Service, not the Content*Analysis*Service. The upload service is a container according to the CS## source code

upvoted 3 times

clarionprogrammer Highly Voted 2 years, 6 months ago

It is not A. "az webapp log" by itself is not a command. It would need to be "az webapp log tail" or "az webapp log download".
https://docs.microsoft.com/en-us/cli/azure/webapp/log?view=azure-cli-latest

upvoted 13 times

BaoNguyen2411 Most Recent 3 months ago

got this question on 29/06/2023

upvoted 1 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 1 times

JH81 4 months, 1 week ago

Selected Answer: C

Another wonderfully obtuse question. I'm not going to disagree that ContentUploadService is hosted in a container but I think MS wants answer C.

upvoted 2 times

adilkhan 6 months, 3 weeks ago

A is the right answer. ContentUploadService has nothing to do with the CONTAINER!

upvoted 1 times

uffuchsi 8 months, 1 week ago

D - Scenario: The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named AmyImages.
The solution will use eight CPU cores.

When you have a misbehaving container in Azure Container Instances, start by viewing its logs with az container logs, and stream its standard out and standard error with az container attach.

https://docs.microsoft.com/en-us/azure/container-instances/container-instances-get-logs

upvoted 3 times

AlexeyG 8 months, 2 weeks ago

Got this in 16/02/2023

upvoted 1 times

hubekpeter 11 months, 1 week ago

Selected Answer: D

https://learn.microsoft.com/en-us/azure/container-instances/container-instances-get-logs

upvoted 3 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: D

az container attach

upvoted 2 times

TheExamMaster2020 11 months, 2 weeks ago

Did my exam on 15th November 2022. This test case and question was on it.

upvoted 1 times

gmishra88 1 year ago

Without knowing what Microsoft thinks it is difficult to answer this question.
* az monitor activity-log alert create can be used to create an alert. If you follow the page then that is the first step. Monitor the service health and for that create alerts. Thi tool can be used for that. https://learn.microsoft.com/en-us/azure/app-service/troubleshoot-http-502-http-503
* az container attach : To get the current log output, but that is not a good option to debug this issue.

I will bet on az monitor activity-log. And I'm changing my original answer from the first iteration of looking at the options

upvoted 1 times

gmishra88 1 year ago

az monitor activity-log alert create -n {AlertName} -g {ResourceGroup} \
--condition category=ServiceHealth and level=Error

upvoted 1 times

gmishra88 1 year ago

az container attach can be used in the second step of collect data (observer and monitor, collect data and mitigate) to get the log stream.

upvoted 1 times

gmishra88 1 year ago

But it asks which you will use first then I guess Microsoft wants me to follow that step exactly like a machine and reading their manuals as soon as I face an issue. So it is definitely alerts in the first step; Observe and monitor. This is a horrible question to ask what I will do first. I will go get a coffee first, microsoft

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: D

D is correct answer.
az container attach

upvoted 5 times

heisenberg33 1 year, 8 months ago

Selected Answer: D

ContentUploadService is a container, see line CS06-CS09. Ref: https://docs.microsoft.com/en-us/azure/container-instances/container-instances-get-logs#attach-output-streams

upvoted 3 times

leonidn 1 year, 9 months ago

Selected Answer: D

ConentCotainerService is hosted in container services. Hense, the correct answer is az container attach.
https://docs.microsoft.com/en-us/azure/container-instances/container-instances-get-logs#attach-output-streams

upvoted 4 times

lugospod 1 year, 9 months ago

My two cents...
Based on info on 502 and 503 on page https://docs.microsoft.com/en-us/azure/app-service/troubleshoot-http-502-http-503
It is possible that the root cause is "our" app or unavailability of Azure service

For Azure services: Microsoft Azure publicizes each time there is a service interruption or performance degradation. Which leads to answer C
For your own applications: A or D. A has no info which command it is using. If using DOWNLOAD it will give us some history info. If using TAIL it will give us real time info similar to D.
So if this happens occasionally, there is a small chance that it will happen right when we are watching so I would rather go and check the history info - A.

Now between A and C - both are correct because the problem could be caused by Azure also since it is occasionally occuring.

So we need more info to decide if the problem is A or C.

upvoted 1 times

ning 2 years, 2 months ago

A is correct, application level log ...
az container attach is redirect of STD input, output and error ...

upvoted 2 times

Topic 17 - Testlet 19

Question #1 *Topic 17*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

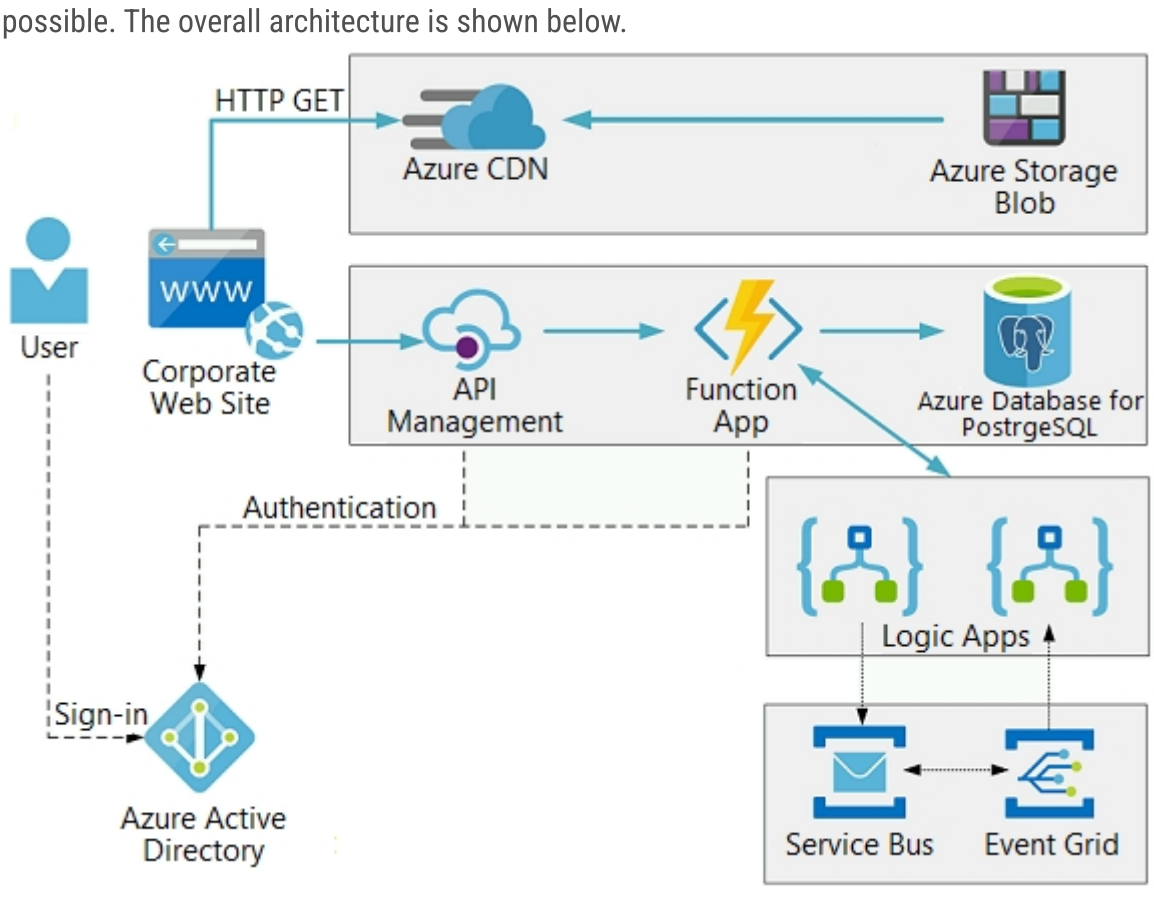

User authentication -

The following steps detail the user authentication process:
1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -

Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

You need to investigate the Azure Function app error message in the development environment.

What should you do?

A. Connect Live Metrics Stream from Application Insights to the Azure Function app and filter the metrics.

B. Create a new Azure Log Analytics workspace and instrument the Azure Function app with Application Insights.

C. Update the Azure Function app with extension methods from Microsoft.Extensions.Logging to log events by using the log instance.

D. Add a new diagnostic setting to the Azure Function app to send logs to Log Analytics.

Correct Answer: *A*

Azure Functions offers built-in integration with Azure Application Insights to monitor functions.

The following areas of Application Insights can be helpful when evaluating the behavior, performance, and errors in your functions:

Live Metrics: View metrics data as it's created in near real-time.

Failures -

Performance -

Metrics -

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/functions-monitoring

Community vote distribution

A (100%)

Deep007 Highly Voted 2 years, 9 months ago

Given answer is correct.
Live Metrics Stream:
when your function app is connected to Application Insights, you can view log data and other metrics in near real time in the Azure portal using Live Metrics Stream. Open below link for this statement.
https://docs.microsoft.com/en-us/azure/azure-functions/functions-monitoring

upvoted 31 times

satyadharma 2 years ago

not a good option :- "If you enable Applications Insights during development, you might hit this limit during testing. Azure provides portal and email notifications when you're approaching your daily limit. If you miss those alerts and hit the limit, new logs won't appear in Application Insights queries. Be aware of the limit to avoid unnecessary troubleshooting time
"

upvoted 3 times

27close Highly Voted 2 years, 11 months ago

Live Metrics Stream: when your function app is connected to Application Insights, you can view log data and other metrics in near real time in the Azure portal using Live Metrics Stream. Use this method when monitoring functions running on multiple-instances or on Linux in a Consumption plan. This method uses sampled data.. Answer is A
https://docs.microsoft.com/en-us/azure/azure-functions/functions-monitoring

upvoted 14 times

surprise0011 Most Recent 6 months, 2 weeks ago

I believe D is might be an option now also
https://learn.microsoft.com/en-us/azure/azure-functions/functions-monitor-log-analytics?tabs=csharp

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: A

Connect Live Metrics Stream from Application Insights to the Azure Function app and filter the metrics.

upvoted 1 times

gmishra88 1 year ago

If it is a timeout I might not go check the metrics but of course it depends on what I expect. I might want to see the application logs. But I wonder how can one question be so confusing and totally dependent on the person checking the issue and his knowledge on the application. It totally depends, Microsoft

upvoted 1 times

gmishra88 1 year ago

It has to be about Metrics because that is mentioned as one of the first steps in one of the obscure microsoft pages

upvoted 1 times

ning 2 years, 2 months ago

Should not B happen before A happen? I would think correct answer is B then A ...

If there is no log analytics, nor application insights connected, how can live stream happens?

upvoted 4 times

DefaultName2 1 year, 11 months ago

Current environment:
“The solution uses Application Insights, Azure Monitor, and Azure Key Vault.”

A is correct

upvoted 1 times

Zidimirite 2 years, 6 months ago

I thought: Why not B? But then realised it is in a development environment and those generally don't log to AppInsights. So A is indeed correct.

upvoted 3 times

ning 2 years, 2 months ago

"when your function app is connected to Application Insights, you can view log data and other metrics in near real time in the Azure portal using Live Metrics Stream. " If there are no application insights, there cannot be any Live Steaming for A

upvoted 1 times

hstml 2 years, 1 month ago

Why do they not log to AppInsights in a development environment? https://docs.microsoft.com/en-us/azure/azure-functions/functions-monitoring

upvoted 3 times

idrisfl 2 years, 7 months ago

The question specifically says "in a development environment". Live Metrics Stream is recommended for production applications. I would go for D.
https://docs.microsoft.com/en-us/azure/azure-monitor/app/live-stream

upvoted 6 times

rajwit 2 years, 9 months ago

shouldn't be D as mentioned in Issues section it's 400 error
https://docs.microsoft.com/en-us/azure/azure-functions/functions-diagnostics

upvoted 5 times

[Removed] 2 years, 9 months ago

I think the answer should be (D) - Add new diagnostics settings.
https://docs.microsoft.com/en-us/azure/azure-functions/functions-monitor-log-analytics?tabs=csharp
"From the Monitoring section of your function app in the Azure portal, select Diagnostic settings, and then select Add diagnostic setting."

Live Metrics is used for monitoring a live system and should not be necessary in a development environment, because there are not that many events to use the "big canon".
https://docs.microsoft.com/en-us/azure/azure-monitor/app/live-stream

upvoted 12 times

Deep007 2 years, 9 months ago

Streaming Logs
While developing an application, you often want to see what's being written to the logs in near real time when running in Azure.

There are two ways to view a stream of the log data being generated by your function executions.

Built-in log streaming: the App Service platform lets you view a stream of your application log files. This stream is equivalent to the output seen when you debug your functions during local development and when you use the Test tab in the portal. All log-based information is displayed. For more information, see Stream logs. This streaming method supports only a single instance, and can't be used with an app running on Linux in a Consumption plan.

Live Metrics Stream: when your function app is connected to Application Insights, you can view log data and other metrics in near real time in the Azure portal using Live Metrics Stream. Use this method when monitoring functions running on multiple-instances or on Linux in a Consumption plan. This method uses sampled data.
https://docs.microsoft.com/en-us/azure/azure-functions/functions-monitoring

upvoted 4 times

JVTM 2 years, 11 months ago

I do not understand the explanation: CONNECT Live Metrics Stream from Application Insights TO THE AZURE FUNCTION app and filter the metrics... Functions have InSights integrated by default. I presume, I should check collected data somehow. But not connect anything... ?

upvoted 5 times

Question #2 *Topic 17*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

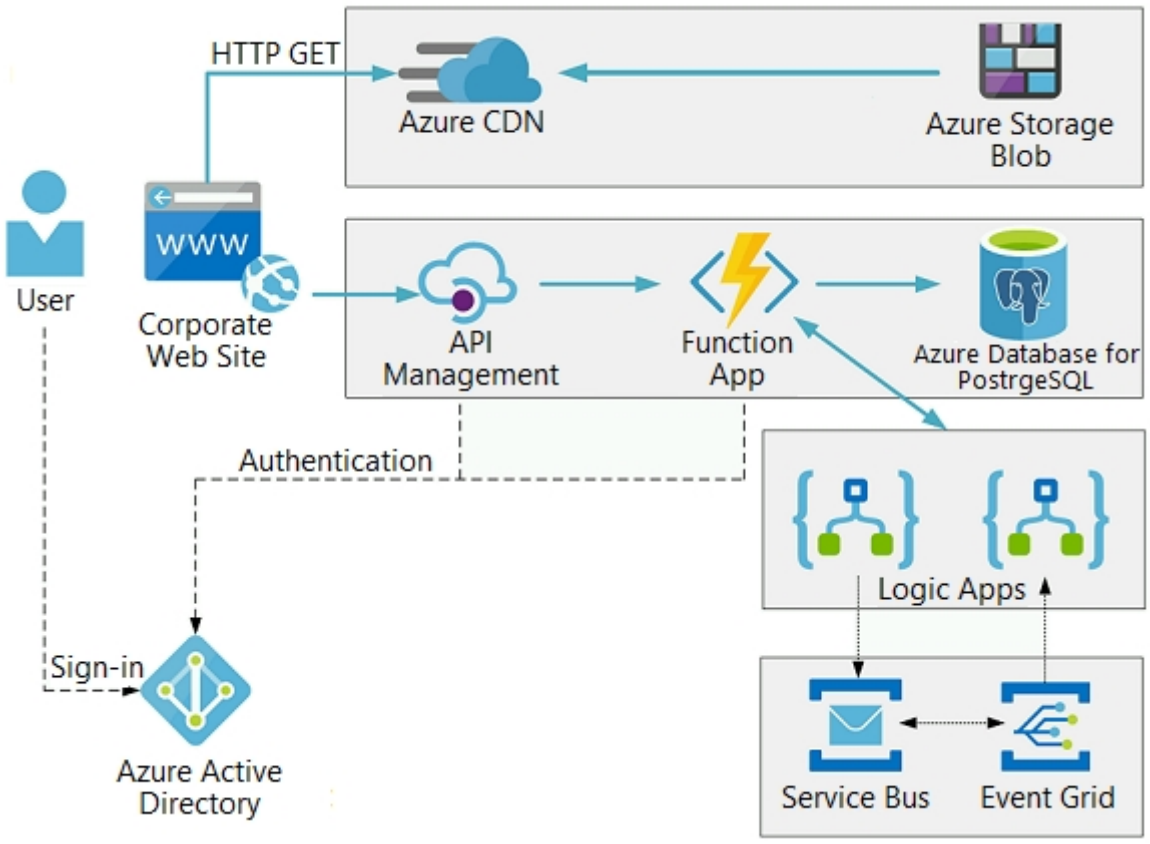

User authentication -

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -
Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

HOTSPOT -

You need to configure security and compliance for the corporate website files.

Which Azure Blob storage settings should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Action	Setting
Restrict file access	role-based access control (RBAC) managed identity shared access signature (SAS) token connection string
Enable file auditing	access tier change feed blob indexer storage account type

Correct Answer:

Answer Area

Action	Setting
Restrict file access	role-based access control (RBAC)
Enable file auditing	storage account type

Box 1: role-based access control (RBAC)

Azure Storage supports authentication and authorization with Azure AD for the Blob and Queue services via Azure role-based access control (Azure RBAC).

Scenario: File access must restrict access by IP, protocol, and Azure AD rights.

Box 2: storage account type -

Scenario: The website uses files stored in Azure Storage

Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR).

Creating a diagnostic setting:

1. Sign in to the Azure portal.
2. Navigate to your storage account.
3. In the Monitoring section, click Diagnostic settings (preview).

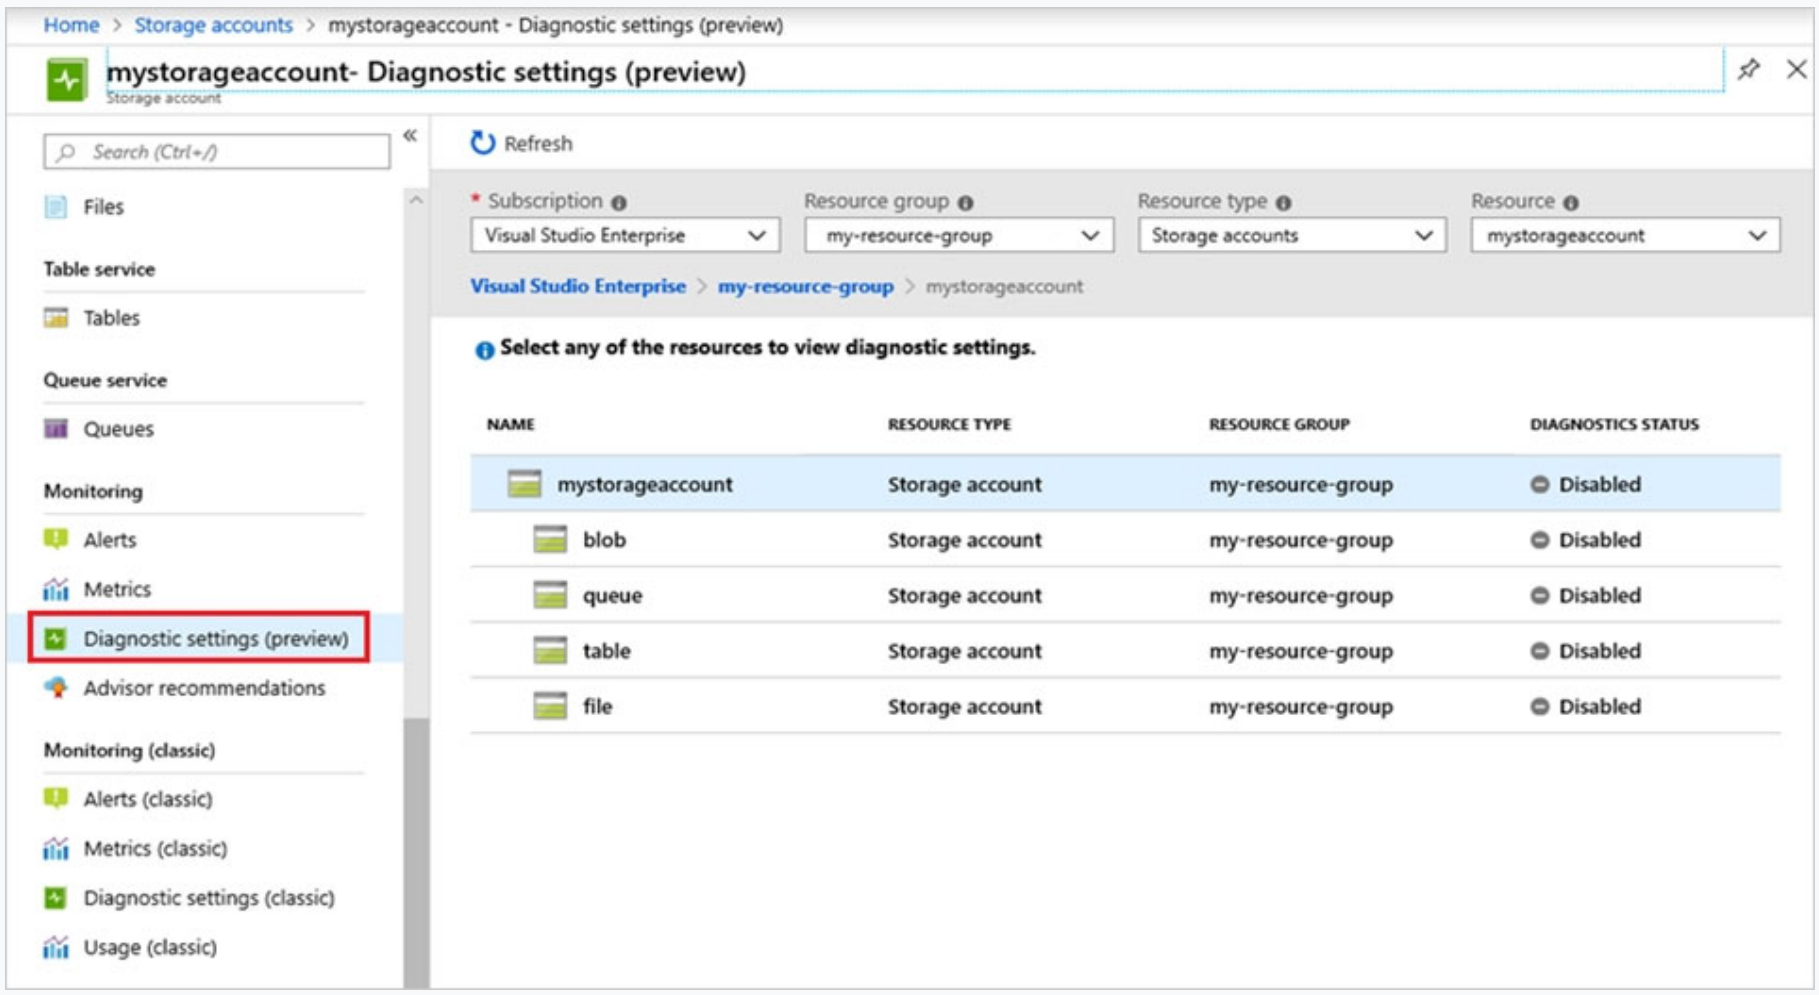

4. Choose file as the type of storage that you want to enable logs for.
5. Click Add diagnostic setting.

Reference:

https://docs.microsoft.com/en-us/azure/storage/common/storage-introduction https://docs.microsoft.com/en-us/azure/storage/files/storage-files-monitoring

clarionprogrammer Highly Voted 2 years, 6 months ago

shared access signature (SAS) token
change feed

upvoted 91 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with above, score 926

upvoted 6 times

mlantonis Highly Voted 2 years, 5 months ago

Box 1: shared access signature (SAS) token
According to the diagram, blob storage is accessed from Azure CDN. Azure CDN doesn't support authentication with managed identity. If you want to grant limited access to private storage containers, you can use the Shared Access Signature (SAS) feature of your Azure storage account. Also, using a managed identity you can't restrict access by IP as requested.

Box 2: change feed
The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account.
The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Reference:
https://docs.microsoft.com/en-us/azure/cdn/cdn-sas-storage-support
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal

upvoted 57 times

huhezculynvhzaljgs 1 year, 11 months ago

Professor is back :)))

upvoted 9 times

edengoforit 1 year, 9 months ago

File access must restrict access by IP, protocol, and Azure AD rights.

Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

upvoted 1 times

applepie Most Recent 3 months, 3 weeks ago

An example of Access storage blobs using an Azure CDN custom domain
It's using SAS.
https://learn.microsoft.com/en-us/azure/cdn/cdn-storage-custom-domain-https

upvoted 1 times

Vmwarevirtual 5 months ago

Appeared the exam I toke at 27-5-2023
I chose SAS and change feed

https://docs.microsoft.com/en-us/azure/cdn/cdn-sas-storage-support
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal

upvoted 1 times

aragones 5 months, 3 weeks ago

Got this 2023-05-12.
Make sure to prepare VanArsdel Inc Canada study case

upvoted 2 times

st0rmtrooperx 10 months, 2 weeks ago

Got this on Dec 16th, 2022. Scored 921 and answered SAS token and change feed.

upvoted 6 times

OPT_001122 11 months, 2 weeks ago

SAS
change feed

upvoted 2 times

coffecold 1 year ago

RBAC and change feed.
why RBAC? Triggering keywords for me are "Azure AD" and "restrict File Access". It seems that some kind of authorization is set for groups.

upvoted 2 times

Eltooth 1 year, 4 months ago

SAS token
change feed

upvoted 3 times

AzureDJ 1 year, 7 months ago

shared access signature (SAS) token
change feed

upvoted 1 times

kozchris 1 year, 8 months ago

Answer: SAS/Change Feed

From problem description:

"Security -
File access must restrict access by IP, protocol, and Azure AD rights."

The keyword here is IP.

From https://docs.microsoft.com/en-us/azure/cdn/cdn-sas-storage-support
"With a SAS, you can define various parameters of access to a blob, such as start and expiry times, permissions (read/write), and IP ranges. "

SAS is from AD so you get the AD rights.

For Change Feed see: https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal
upvoted 2 times

leonidn 1 year, 9 months ago
Agree on RBAC.
Change feed

The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes. https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal
upvoted 2 times

cool_tool 2 years, 2 months ago
RBAC
change feed
upvoted 7 times

ning 2 years, 2 months ago
Correct, file access is AD User based rights. IP and Protocol, can be configured separately
upvoted 1 times

Kuna_Lambo 2 years, 7 months ago
managed identity
change feed
upvoted 4 times

inputoutput 2 years, 7 months ago
According to the diagram, blob storage is accessed from Azure CDN. Azure CDN doesn't support authentication with managed identity. I think the correct answer is Shared Access Token. https://docs.microsoft.com/en-us/azure/cdn/cdn-sas-storage-support
upvoted 10 times

Kuna_Lambo 2 years, 7 months ago
Thanks, I think you are right.
upvoted 1 times

jay158 2 years, 3 months ago
See the arrow -- Flow is from Storage to CDN.
Diagram does not show, how Storage is populated.
No one will populate storage via CDN
upvoted 2 times

rdemontis 2 years, 7 months ago
Exactly, and using a managed identity you can't restrict access by IP as requested. User delegation SAS is the right choice in this case (you need AAD integration) and change feed is the service designed for file audits.
https://docs.microsoft.com/en-us/rest/api/storageservices/create-user-delegation-sas
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal
upvoted 16 times

kwaazaar 2 years, 6 months ago
But RBAC is supported on file shares too. It needs Azure AD Domain Services, I think.
upvoted 1 times

Topic 18 - Testlet 2

Question #1 *Topic 18*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

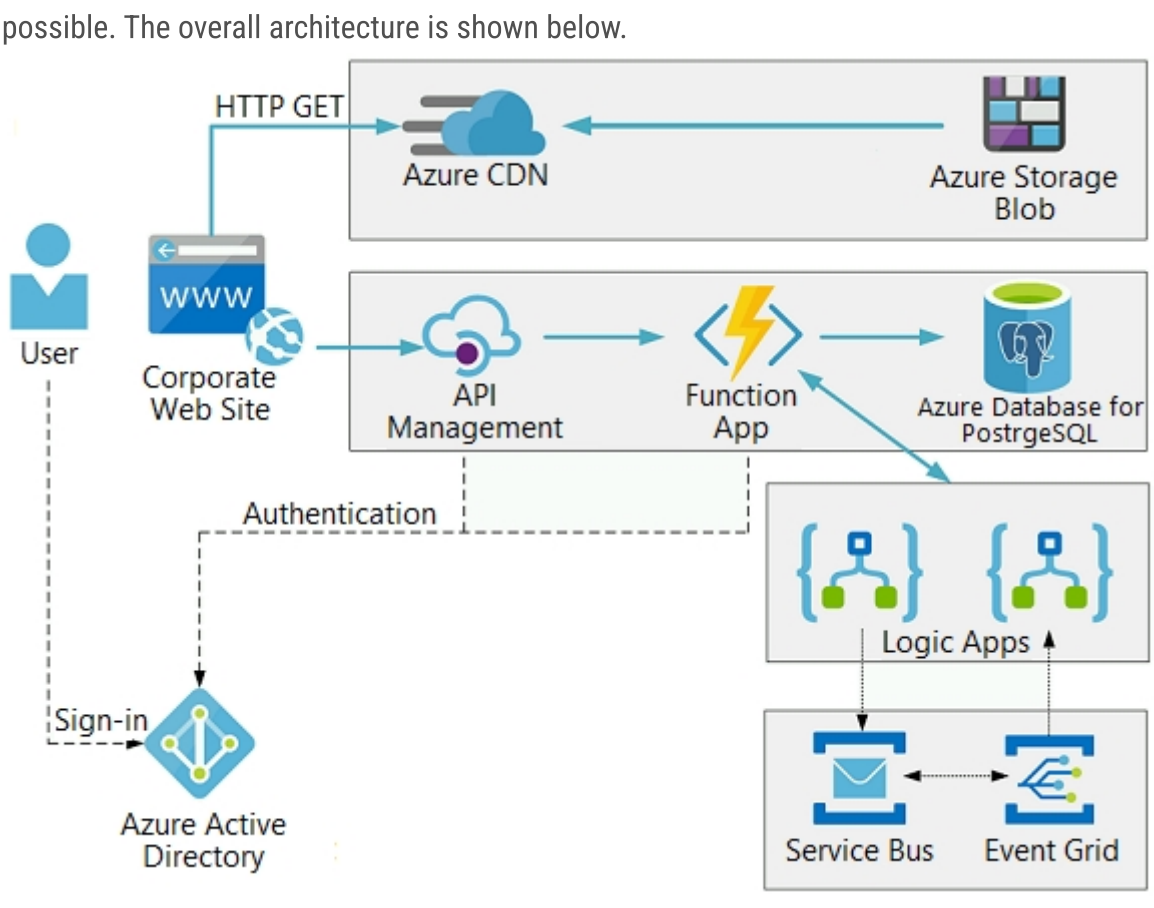

User authentication -

The following steps detail the user authentication process:
1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -

Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

You need to correct the RequestUserApproval Function app error.

What should you do?

A. Update line RA13 to use the async keyword and return an HttpRequest object value.

B. Configure the Function app to use an App Service hosting plan. Enable the Always On setting of the hosting plan.

C. Update the function to be stateful by using Durable Functions to process the request payload.

D. Update the functionTimeout property of the host.json project file to 15 minutes.

Correct Answer: *C*

Async operation tracking -

The HTTP response mentioned previously is designed to help implement long-running HTTP async APIs with Durable Functions. This pattern is sometimes referred to as the polling consumer pattern.

Both the client and server implementations of this pattern are built into the Durable Functions HTTP APIs.

Function app -

You perform local testing for the RequestUserApproval function. The following error message displays:

'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'

The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -

| where FunctionName = = "RequestUserApproval"

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-http-features

Community vote distribution

C (100%)

rdemontis Highly Voted 2 years, 7 months ago

Answer is correct. In addition you can see this pattern in microosoft documentation and it's known as Human Interaction.
https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-overview?tabs=csharp#human

upvoted 34 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with given answer, score 926

upvoted 2 times

nflappo89 Highly Voted 2 years, 4 months ago

http trigger timeout is maxed to 230 seconds, so the only available option is to change to a durable one

upvoted 10 times

maukaba 1 year, 11 months ago

"Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response." REF: https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale

upvoted 2 times

OPT_001122 Most Recent 11 months, 2 weeks ago

Selected Answer: C

Durable Function.

upvoted 2 times

Eltooth 1 year, 4 months ago

Selected Answer: C

C is correct answer.
Durable Function.

upvoted 4 times

MrXBasit 2 years, 3 months ago

Answer is 100% correct

upvoted 3 times

farich 2 years, 5 months ago

I think here the answer is just to increase the functionTimeout.

Durable Functions are used when you want to introduce a state, it is not just "if you have a long running function switch to Durable Functions".

Another argument against Durable Functions is that architectural change seem to be out of scope for this case study.

Another argument is that function can run for up to 30 mins.

upvoted 3 times

maukaba 1 year, 11 months ago

Maximum timeout you can set is 10 min in consumption plan. Unless you go to premium which is unlimited: https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale

upvoted 1 times

xRiot007 1 year, 4 months ago

"if you have a long running function switch to Durable Functions" - actually, that is one of the features of DF. Modifying the function app timeout will do nothing. There is a hard limit of 230 seconds of timeout. Anything larger than that, use DF, state or no state.

upvoted 1 times

Topic 19 - Testlet 20

Question #1 *Topic 19*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements -

Policy service -

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Policies -

Log policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events -

Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

PolicyLib -

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Other -

Anomaly detection service -

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring -

All web applications and services have health monitoring at the /health service endpoint.

Issues -

Policy loss -
When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue -
When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency -
Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code -

EventGridController.cs -
Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

EventGridController.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "inputl",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
EG55   }
EG56   private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57   {
EG58     . . .
EG59   }
EG60   private string GetLogData(string uri)
EG61   {
EG62     . . .
EG63   }
EG64   static string BlobStoreAccountSAS(string containerName)
EG65   {
EG66     . . .
EG67   }
EG68 }
```

LoginEvent.cs -

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

Question

DRAG DROP -

You need to implement the Log policy.

How should you complete the Azure Event Grid subscription? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment may be used once, more than once, or not at all. You may need to drag the split bar between panes to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Code segment

- All
- WebHook
- EventHub
- subjectEndsWith
- Mictosoft.Storage
- subjectBeginsWith
- Microsoft.Storage.BlobCreated

Answer Area

```
{
  "name": "newlogs",
  "properties": {
    "topic": "/subscriptions/. . ./providers/Microsoft.EventGrid/topics/. . .",
    "destination": {
    "endpointType"  : " [ code segment ] " },
    "filter": {
       " [ code segment ] ": "/blobServices/default/containers/logdrop/",
       "includedEventTypes": [ " [ code segment ] "   ] },
    },
    "labels": [],
    "eventDeliverySchema": "EventGridSchema"
```

Correct Answer:

Code segment

- All
- WebHook
- EventHub
- subjectEndsWith
- Mictosoft.Storage
- subjectBeginsWith
- Microsoft.Storage.BlobCreated

Answer Area

```
{
  "name": "newlogs",
  "properties": {
    "topic": "/subscriptions/. . ./providers/Microsoft.EventGrid/topics/. . .",
    "destination": {
    "endpointType"  : " [WebHook] " },
    "filter": {
       " [subjectBeginsWith] ": "/blobServices/default/containers/logdrop/",
       "includedEventTypes": [ " [Microsoft.Storage.BlobCreated] "   ] },
    },
    "labels": [],
    "eventDeliverySchema": "EventGridSchema"
```

Box 1:WebHook -

Scenario: If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook. endpointType: The type of endpoint for the subscription (webhook/HTTP, Event Hub, or queue).

Box 2: SubjectBeginsWith -

Box 3: Microsoft.Storage.BlobCreated

Scenario: Log Policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Example subscription schema -

{

"properties": {

"destination": {

"endpointType": "webhook",
"properties": {
"endpointUrl": "https://example.azurewebsites.net/api/HttpTriggerCSharp1?
code=VXbGWce53l48Mt8wuotr0GPmyJ/nDT4hgdFj9DpBiRt38qqnnm5OFg=="
}
},
"filter": {
"includedEventTypes": ["Microsoft.Storage.BlobCreated", "Microsoft.Storage.BlobDeleted"],
"subjectBeginsWith": "blobServices/default/containers/mycontainer/log",
[1]
"isSubjectCaseSensitive ": "true"
}
}
}
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/subscription-creation-schema

clarionprogrammer Highly Voted 2 years, 6 months ago

Looks correct based on the provided reference.
https://docs.microsoft.com/en-us/azure/event-grid/subscription-creation-schema

upvoted 27 times

gmishra88 1 year ago

You mean to say that documentation link has eventType as "Webhook" as an example. But isn't that an example? Where is the reference to say it is a webhook? The question itself is not clear

upvoted 3 times

coffecold Highly Voted 1 year ago

To prevent reading the cases multiple times:
Please see the spots below where you can find the questions (page/topic/question/subject)
Proseware, Inc.
53 19 1 ARM for Azure Event Grid subscription, enable + filter logging
53 19 2 use Application Insights for scaling
53 19 3 code : telemetry filter/ telemetry processor
53 19 4 code : telemetry initializer
55 25 1 code : if-then condition on properties of incomming event from GRID
56 25 2 code : class for login event, properties needed
56 25 3 code : programatically update application settings for blob
56 26 1 azure function delay : always on and tier

upvoted 15 times

ks321 3 months, 1 week ago

Dude, you are awesome, Seems like you have provided this for all case study questions. Thanks a ton!

upvoted 1 times

OPT_001122 10 months, 3 weeks ago

Thanks again for this great help!!

upvoted 2 times

gmishra88 Most Recent 1 year ago

It looks correct purely based on options and there is no need to read the big scenario. But I wonder why it is reacting on logdrop container events. That's where the service itself is writing log files to. Won't it create an endless loop?

upvoted 3 times

coffecold 11 months, 4 weeks ago

"All Azure App Service Web Apps must write logs to Azure Blob storage"
-There are n (production) Web apps that write logs to a blob storage.
-On that blob storage an event grid is put to trigger the event when blobs are written.
-From that event a function is triggered, doing great thing with policies.

upvoted 2 times

gmishra88 1 year ago

marvelous microsoftiness, probably this is for the subscription for the anomaly detection service and not the policy service where they explicIty mentioned event grid to confuse you. Once you pass this exam, the certification can be used when applying for a job at CSI new york. For finding clues where they do not exist.

upvoted 7 times

acjdev514 1 month ago

hahaha gmishra88

upvoted 1 times

edengoforit 1 year, 9 months ago

What is a webhook and how do you use it?
Webhooks are automated messages sent from apps when something happens. They have a message—or payload—and are sent to a unique URL—essentially the app's phone number or address. Webhooks are almost always faster than polling, and require less work on your end. They're much like SMS notifications.

upvoted 3 times

gfiorini 1 year, 11 months ago

The question is really worded poorly. The subscription is for events that are inserted into the container and a webhook is invoked when a log file is created in the storage container. I assume that the webhook point to a function (?) that analyze the log with a machine learning model and THEN if there is an anomaly call an azure function to notify adminstrator via mail. Am I missing something ?

upvoted 4 times

coffecold 1 year ago

It's in the text:
"You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events"
The controller of "Policy service" is given in the code EventGridController.cs.

So the subscription in Grid triggers a webhook to the Policy service and executes the 'Process' controller function there .

upvoted 1 times

mc0re 2 years, 2 months ago

> an Azure Function that emails administrators is called by using an HTTP WebHook

The question is about the EventGrid subscription, not how another Function is called. How do they fit together?

upvoted 3 times

MiraA 2 years ago

I am missing "properties" section within "destination" key to define either "endpointUrl" or "resourceId".

https://docs.microsoft.com/en-us/azure/event-grid/subscription-creation-schema#example-subscription-schema
https://docs.microsoft.com/en-us/azure/event-grid/handler-event-hubs#event-hub

upvoted 1 times

MiraA 2 years ago

Endpoint properties based on the "destination" type:
https://docs.microsoft.com/en-us/azure/templates/microsoft.eventgrid/eventsubscriptions?tabs=json#endpointproperties-object

upvoted 1 times

UnknowMan 2 years, 5 months ago

Correct !

upvoted 2 times

UnknowMan 2 years, 5 months ago

Correct !

upvoted 2 times

jokergester 2 years, 7 months ago

https://docs.microsoft.com/en-us/azure/templates/microsoft.eventgrid/eventsubscriptions?tabs=json#template-format

upvoted 1 times

Question #2 *Topic 19*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements -

Policy service -

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Policies -

Log policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events -

Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

PolicyLib -

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Other -

Anomaly detection service -

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring -

All web applications and services have health monitoring at the /health service endpoint.

Issues -

Policy loss -
When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue -
When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency -
Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code -

EventGridController.cs -
Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

EventGridController.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "inputl",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
EG55   }
EG56   private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57   {
EG58     . . .
EG59   }
EG60   private string GetLogData(string uri)
EG61   {
EG62     . . .
EG63   }
EG64   static string BlobStoreAccountSAS(string containerName)
EG65   {
EG66     . . .
EG67   }
EG68 }
```

LoginEvent.cs -

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

Question

You need to ensure that the solution can meet the scaling requirements for Policy Service.

Which Azure Application Insights data model should you use?

A. an Application Insights dependency

B. an Application Insights event

C. an Application Insights trace

D. an Application Insights metric

Correct Answer: *D*

Application Insights provides three additional data types for custom telemetry:

Trace - used either directly, or through an adapter to implement diagnostics logging using an instrumentation framework that is familiar to you, such as Log4Net or

System.Diagnostics.

Event - typically used to capture user interaction with your service, to analyze usage patterns.

Metric - used to report periodic scalar measurements.

Scenario:

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/data-model

Community vote distribution

D (100%)

Kitkit Highly Voted 2 years, 8 months ago

Answer is correct. User can use metric telemetry to get different application metrics like: requestsPerSecond, requestsInQueue, and use these values to know when to scale

https://docs.microsoft.com/en-us/azure/azure-monitor/app/data-model-metric-telemetry

upvoted 24 times

mlantonis Highly Voted 2 years, 5 months ago

Application Insights provides three additional data types for custom telemetry:

- Trace: used either directly, or through an adapter to implement diagnostics logging using an instrumentation framework that is familiar to you, such as Log4Net or System.Diagnostics.
- Event: typically used to capture user interaction with your service, to analyze usage patterns.
- Metric: used to report periodic scalar measurements.

Scenario:

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

So, it is D.

upvoted 18 times

mlantonis 2 years, 5 months ago

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/data-model

https://docs.microsoft.com/en-us/azure/azure-monitor/app/data-model-metric-telemetry

upvoted 5 times

NombreFalso Most Recent 8 months, 1 week ago

Selected Answer: D

Always metric
upvoted 3 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: D

Answer is correct
upvoted 2 times

coffecold 1 year ago

"Policy service" is a web app that needs to scale automatically depending on the usage
I agree metrics will provide you the information (so answer D is best), but that is only part of the solution. What about the rest ? How is it scaled and how is it scaled automatically?
I'm missing the words "Monitor-respond-autoscale-rules"
upvoted 1 times

macobuzi 1 month, 3 weeks ago

Agree, this is confusing as hell!
upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

çorrect answer provided: Application Insights metric
upvoted 1 times

UnknowMan 2 years, 5 months ago

Correct and useful for Scaling requirements
upvoted 2 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements -

Policy service -

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Policies -

Log policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events -

Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

PolicyLib -

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Other -

Anomaly detection service -

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring -

All web applications and services have health monitoring at the /health service endpoint.

Issues -

Policy loss -
When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue -
When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency -
Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code -

EventGridController.cs -
Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

EventGridController.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "inputl",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
EG55   }
EG56   private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57   {
EG58     . . .
EG59   }
EG60   private string GetLogData(string uri)
EG61   {
EG62     . . .
EG63   }
EG64   static string BlobStoreAccountSAS(string containerName)
EG65   {
EG66     . . .
EG67   }
EG68 }
```

LoginEvent.cs -

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

```
LoginEvent.cs
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

Question

DRAG DROP -

You need to implement telemetry for non-user actions.

How should you complete the Filter class? 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:

Code segments

- /health
- /status
- RequestTelemetry
- PageViewTelemetry
- ITelementryProcessor
- ITelementryInitializer

Answer Area

```
public class Filter : [ code segment ]
{
   private readonly [ code segment ] _next;
   public (Filter [ code segment ] next)
   {
   _next = next;
   }
   public void Process(ITelemetry item)
   {
     var x = item as [ code segment ];
      if (x?.Url.AbsolutePath == " [ code segment ] " )
     {
       return;
     }
     _next.Process(item);
   }
}
```

Correct Answer:

Code segments

- /health
- /status
- RequestTelemetry
- PageViewTelemetry
- ITelementryProcessor
- ITelementryInitializer

Answer Area

```
public class Filter : ITelementryProcessor
{
   private readonly ITelementryProcessor _next;
   public (Filter ITelementryProcessor next)
   {
   _next = next;
   }
   public void Process(ITelemetry item)
   {
     var x = item as RequestTelemetry ;
      if (x?.Url.AbsolutePath == " /health " )
     {
       return;
     }
     _next.Process(item);
   }
}
```

Scenario: Exclude non-user actions from Application Insights telemetry.

Box 1: ITelemetryProcessor -
To create a filter, implement ITelemetryProcessor. This technique gives you more direct control over what is included or excluded from the telemetry stream.

Box 2: ITelemetryProcessor -

Box 3: ITelemetryProcessor -

Box 4: RequestTelemetry -

Box 5: /health -
To filter out an item, just terminate the chain.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-filtering-sampling

Sukon_Desknot Highly Voted 2 years, 2 months ago
(∩^o^)⊃—☆
upvoted 21 times

surmistry Highly Voted 2 years, 6 months ago
"ITelemetryProcessor" typo, no "n"
https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-filtering-sampling#create-a-telemetry-processor-c
upvoted 14 times

fr369 1 year, 9 months ago
Also the constructor is incorrect. Should be "public Filter(" instead of "public (Filter".
upvoted 8 times

Mev4953 1 year, 9 months ago
Yes, you are right. It should be;
public Filter (ITelemetryProcessor next)
upvoted 5 times

coffecold Most Recent 1 year ago
The question is so wrong
"You need to implement telemetry for non-user actions"
that MUST be:
"You need to implement telemetry to exclude non-user actions" (according to the requirements).

In that case you have to exclude RequestTelemetry in combination with /health
upvoted 3 times

gmishra88 1 year ago

This is chaining of ITelemetryProcessors where the Filter (the one created here) is the first one. The ITelemetryProcessor instance passed in the constructor is the next in the chain. The code checks if the url is /health and it does not do anything otherwise it fowards to the next (which is the one passed in the Constructor).

Look at the line: "item as RequestTelemetry" : What kind of people work in Microsoft that will do the casting without checking the instance of the item? No wonder Microsoft things are so buggy

upvoted 1 times

coffecold 1 year ago

"The as operator is used to perform conversion between compatible reference types or Nullable types. This operator returns the object when they are compatible with the given type and return null if the conversion is not possible instead of raising an exception"

So the code is correct

upvoted 1 times

xRiot007 1 year, 4 months ago

In code, public Filter() not public (Filter ...)

upvoted 3 times

SivajiTheBoss 1 year, 7 months ago

Answer: ITelemetryProcessor, ITelemetryProcessor,ITelemetryProcessor,RequestTelemetry, /health

Reference: https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-filtering-sampling

public class SuccessfulDependencyFilter : ITelemetryProcessor
{
private ITelemetryProcessor Next { get; set; }

// next will point to the next TelemetryProcessor in the chain.
public SuccessfulDependencyFilter(ITelemetryProcessor next)
{
this.Next = next;
}
}

var requestTelemetry = telemetry as RequestTelemetry;

upvoted 11 times

Mev4953 1 year, 9 months ago

Yes, you are right. It should be;
public Filter (ITelemetryProcessor next)

upvoted 1 times

Mev4953 1 year, 9 months ago

ITelemetryInitializer:
Use telemetry initializers to define global properties that are sent with all telemetry; and to override selected behavior of the standard telemetry modules.
For example, the Application Insights for Web package collects telemetry about HTTP requests. By default, it flags as failed any request with a response code >= 400. But if you want to treat 400 as a success, you can provide a telemetry initializer that sets the Success property.

ITelemetryProcessor:
This technique gives you more direct control over what is included or excluded from the telemetry stream. You can use it in conjunction with Sampling, or separately. To filter telemetry, you write a telemetry processor and register it with the SDK

https://github.com/uglide/azure-content/blob/master/articles/application-insights/app-insights-api-filtering-sampling.md

upvoted 3 times

Saterial 2 years, 6 months ago

Looks Good

upvoted 8 times

Question #4 *Topic 19*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements -

Policy service -

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Policies -

Log policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events -

Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

PolicyLib -

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Other -

Anomaly detection service -

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring -

All web applications and services have health monitoring at the /health service endpoint.

Issues -

Policy loss -
When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue -
When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency -
Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code -

EventGridController.cs -
Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

EventGridController.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "inputl",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
EG55   }
EG56   private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57   {
EG58     . . .
EG59   }
EG60   private string GetLogData(string uri)
EG61   {
EG62     . . .
EG63   }
EG64   static string BlobStoreAccountSAS(string containerName)
EG65   {
EG66     . . .
EG67   }
EG68 }
```

LoginEvent.cs -

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

Question

DRAG DROP -

You need to ensure that PolicyLib requirements are met.

How should you complete the code segment? 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:

Code segments

- Process
- Initialize
- telemetry.Sequence
- ITelemetryProcessor
- ITelemetryInitializer
- telemetry.Context
- EventGridController.EventId.Value
- ((EventTelemetry)telemetry).Properties["EventId"]

Answer Area

```
public class IncludeEventId : [ code segment ]
{
   public void [ code segment ] (ITelemetry telemetry)
   {

      [ code segment ].Properties["EventId"] =
        [ code segment ];
   }
}
```

Correct Answer:

Code segments

- Process
- Initialize
- telemetry.Sequence
- ITelemetryProcessor
- ITelemetryInitializer
- telemetry.Context
- EventGridController.EventId.Value
- ((EventTelemetry)telemetry).Properties["EventId"]

Answer Area

```
public class IncludeEventId : ITelemetryInitializer
{
   public void Initialize (ITelemetry telemetry)
   {

      telemetry.Context .Properties["EventId"] =
        ((EventTelemetry)telemetry).Properties["EventId"] ;
   }
}
```

Scenario: You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications.

The PolicyLib library must:

- Exclude non-user actions from Application Insights telemetry.
- Provide methods that allow a web service to scale itself.
- Ensure that scaling actions do not disrupt application usage.

Box 1: ITelemetryInitializer -

Use telemetry initializers to define global properties that are sent with all telemetry; and to override selected behavior of the standard telemetry modules.

Box 2: Initialize -

Box 3: Telemetry.Context -

Box 4: ((EventTelemetry)telemetry).Properties["EventID"]

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-filtering-sampling

[Removed] Highly Voted 2 years, 6 months ago

* EventId is held by class EventGridController
* You can add properties to telemetry by implementing ITelemetryInitializer which defines the Initialize method.
* ITelemetry.Context.Properties is correct, but shouldnt be used any more as obsoloete

```
public class IncludeEventId : ITelemetryInitializer
{
public void Initialize (ITelemetry telemetry)
{
telemetry.Context.Properties["EventId"] = EventgridController.EventId.Value;
}
}
```

https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-custom-events-metrics#sampling-filtering-and-processing-telemetry
https://docs.microsoft.com/en-us/dotnet/api/microsoft.applicationinsights.datacontracts.telemetrycontext.properties?view=azure-dotnet#Microsoft_ApplicationInsights_DataContracts_TelemetryContext_Properties

upvoted 42 times

Knightie 10 months, 4 weeks ago

the default answer's last line doesn't make sense... it use the property value to put inside it's context.property value within itself...

a meaningful piece of code is to set the propertie value with another Global Object value (that the only sense making value calling inside a local method of a class)... so setting the Eventgrid's value to the pass-by-reference argument is the only sensible piece of code.

upvoted 1 times

Rakeshpro 1 year, 2 months ago

good job buddy!

upvoted 2 times

uffuchsi Highly Voted 8 months, 1 week ago

Box 1: ITelemetryInitializer
Box 2: Initialize
Box 3: Telemetry.Context
Box 4: EventGridController.EventId.Value

Scenario: The application reacts to events from Azure Event Grid and performs policy actions based on those events. The application must include the Event Grid Event ID field in all Application Insights telemetry.

\You can add properties to telemetry by implementing ITelemetryInitializer which defines the Initialize method.

https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-custom-events-metrics#sampling-filtering-and-processing-telemetry

upvoted 5 times

warchoon 6 months, 3 weeks ago

"Telemetry initializers are called before calling telemetry processors". So EventId.Value is not initialized by Process method yet.

upvoted 1 times

OPT_001122 Most Recent 10 months, 3 weeks ago

Box 1: ITelemetryInitializer
Box 2: Initialize
Box 3: Telemetry.Context
Box 4: ((EventTelemetry)telemetry).Properties["EventID"]

upvoted 3 times

OPT_001122 11 months, 2 weeks ago

anyone got this case study in exam?

upvoted 1 times

AzureDJ 1 year, 7 months ago

The given answer is correct. Box 4 is inside a function which takes in a parameter called: telemetry. The telemetry parameter, serves a purpose: to be used in the function in: ((EventTelemetry)telemetry).Properties["EventID"].

upvoted 2 times

AzureDJ 1 year, 7 months ago

The given answer is correct, except Box 4 should be: EventGridController.EventId.Value

upvoted 2 times

UnknowMan 2 years, 5 months ago

Not correct.
= > Use the EventId static prop of EventgridController

upvoted 4 times

Frakandel 2 years, 6 months ago

How static is the content of Event Grid Event ID... I agree with atomicicebreaker, bt doubt between "IProcess, Process, Context & EventHub" and "IInitialise, Initialise, Context & EventHub"

upvoted 1 times

Kuna_Lambo 2 years, 7 months ago

Box #4 is ok?

upvoted 4 times

rdemontis 2 years, 7 months ago

No, in my opinion it has no sense assign event id to telemetry object from the same object itself. you should take in from other objects, and in this case the only option possible seems to be EventGridController class.

upvoted 17 times

MrZoom 2 years, 7 months ago

Agreed. It states that "The application must include the Event Grid Event ID field in all Application Insights telemetry.".

upvoted 3 times

atomicicebreaker 2 years, 6 months ago

Yup, and the property is indeed set in a controller

upvoted 3 times

MiraA 2 years ago

Yes, it is set here.
Controller accepts JSON with the set of the events. The events are processed using foreach() and this is the place where EventGridController.EventId is set.
The trick is the EventId is defined as:
public static AsyncLocal<string> EventId = ...
and AsyncLocal "represents ambient data that is local to a given asynchronous control flow, such as an asynchronous method.".

https://docs.microsoft.com/en-us/dotnet/api/system.threading.asynclocal-1?view=net-5.0

upvoted 1 times

Topic 20 - Testlet 21

Question #1 *Topic 20*

Introductory Info

Case study -

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

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

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

To start the case study -

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

Background -

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web

API project that is deployed as an Azure Web App.

Overall architecture -

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing -

Employees may upload receipts in two ways:

Uploading using an Azure Files mounted folder

Uploading using the web application

Data Storage -

Receipt and employee information is stored in an Azure SQL database.

Documentation -

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for

Azure File upload, and instructions on how to configure the mounted folder.

Solution details -

Users table -

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application -

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing -

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob
Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging -
Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled.
Application Insights must always contain all log messages.

Requirements -

Receipt processing -
Concurrent processing of a receipt must be prevented.

Disaster recovery -
Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security -
User's SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
Receipt data must always be encrypted at rest.
All data must be protected in transit.
User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues -

Upload format issue -
Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File
Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal
Server error page.

Capacity issue -
During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

Log capacity issue -
Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Application code -

Processing.cs -

```
PC01 public static class Processing
PC02 {
PC03   public static class Function
PC04   {
PC05     [FunctionName("IssueWork")]
PC06     public static async Task Run([TimerTrigger("0 */5 * * * *")] TimerInfo timer, ILogger
log)
PC07     {
PC08       var container = await GetCloudBlobContainer();
PC09       foreach (var fileItem in await ListFiles())
PC10       {
PC11         var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12         var ms = new MemoryStream();
PC13         await file.DownloadToStreamAsync(ms);
PC14         var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15         await blob.UploadFromStreamAsync(ms);
PC16
PC17       }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob(CloudBlockBlob sourceBlob)
PC20     {
PC21        . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25       var cloudBlobClient = new CloudBlobClient(new Uri(". . ."), await GetCredentials());
PC26
PC27       await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28       return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32        . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36        . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(". . .");
PC38   }

PC39 }
```

Database.cs -

```
DB01 public class Database
DB02 {
DB03   private string ConnectionString =
DB04
DB05   public async Task<object> LoadUserDetails(string userId)
DB06   {
DB07
DB08     return await policy.ExecuteAsync(async () =>
DB09     {
DB10       using (var connection = new SqlConnection(ConnectionString))
DB11       {
DB12         await connection.OpenAsync();
DB13         using (var command = new SqlCommand("…", connection))
DB14         using (var reader = command.ExecuteReader())
DB15         {
DB16           …
DB17         }
DB18       }
DB19     });
DB20   }

DB21 }
```

ReceiptUploader.cs -

```
RU01 public class ReceiptUploader
RU02 {
RU03   public async Task UploadFile(string file, byte[] binary)
RU04   {
RU05     var httpClient = new HttpClient();
RU06     var response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU07     while (ShouldRetry(response))
RU08     {
RU09       response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU10     }
RU11   }
RU12   private bool ShouldRetry(HttpResponseMessage response)
RU13   {
RU14
RU15   }
RU16 }
```

ConfigureSSE.ps1 -

```
CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "..." -AccountName "..."
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "..."
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "..."
CS04  Set-AzureRmKeyVaultAccessPolicy `
CS05   -VaultName $keyVault.VaultName `
CS06   -ObjectId $storageAccount.Identity.PrincipalId `
CS07
CS08
CS09  Set-AzureRmStorageAccount `
CS10   -ResourceGroupName $storageAccount.ResourceGroupName `
CS11   -AccountName $storageAccount.StorageAccountName `
CS12   -EnableEncryptionService File `
CS13   -KeyvaultEncryption `
CS14   -KeyName $key.Name
CS15   -KeyVersion $key.Version `
CS16   -KeyVaultUri $keyVault.VaultUri
```

Question

You need to ensure receipt processing occurs correctly.

What should you do?

A. Use blob properties to prevent concurrency problems

B. Use blob SnapshotTime to prevent concurrency problems

C. Use blob metadata to prevent concurrency problems

D. Use blob leases to prevent concurrency problems

Correct Answer: *B*

You can create a snapshot of a blob. A snapshot is a read-only version of a blob that's taken at a point in time. Once a snapshot has been created, it can be read, copied, or deleted, but not modified. Snapshots provide a way to back up a blob as it appears at a moment in time.

Scenario: Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in

Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Reference:

https://docs.microsoft.com/en-us/rest/api/storageservices/creating-a-snapshot-of-a-blob

Community vote distribution

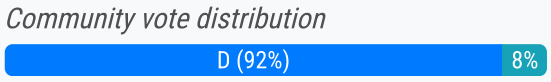

Arul1705 Highly Voted 2 years, 7 months ago

Answer is D: Use blob leases to prevent concurrency problems

upvoted 56 times

Shion2009 2 years, 7 months ago

I guess this could be a problem with the "Processing" scenario:
"Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user."
It seems a lease will only be kept alive for 60 seconds, so it shouldn't be an option.

upvoted 3 times

Bengkel 2 years, 7 months ago

A lease does not lock a file for read operations, just for write and delete operations. The lock duration can be 15 to 60 seconds, or can be infinite. So I guess the effect is the same?
https://docs.microsoft.com/en-us/rest/api/storageservices/lease-blob

upvoted 3 times

rdemontis 2 years, 7 months ago

Why should you avoid to use blob lease to guarantee access the blob file from the report link? Generally to give access to a blob from http you use SAS. In the SAS you can specify an expiration datetime according to your need. Blob lease is instead a way just to prevent concurrent access and puts a lock on blob only for write and delete operations. Any other can still view the content but can't modify or delete it. So in my opinion the correct answer is Blob lease. In addition I have found the same question on a udemy course test and the answer is just "Blob Lease".
https://docs.microsoft.com/en-us/rest/api/storageservices/lease-blob

upvoted 3 times

MrZoom 2 years, 7 months ago

Agreed. The question is about the receipt processing. The case states "Concurrent processing of a receipt must be prevented", which can't be done with snapshots and leases are made for this specifically. So answer should be D.

upvoted 5 times

Javierzgz Highly Voted 2 years, 7 months ago

I think it is lease (see this link)
https://docs.microsoft.com/en-us/azure/storage/blobs/concurrency-manage?tabs=dotnet

upvoted 11 times

kabbas Most Recent 4 months ago

Selected Answer: D

must be lease because if you use snapshot to do the processing another thread could write or update the original which could mean you are actually processing out of date data

upvoted 1 times

adilkhan 7 months, 2 weeks ago

D. Use blob leases to prevent concurrency problems.

To prevent concurrent processing of a receipt, we need to ensure that only one instance of the processor can access the receipt file at a given time. Blob leases can be used to implement this mechanism. A blob lease allows an application to acquire a lease on a blob for a specified period, during which no other application can modify the blob. When the lease expires or is released, another application can acquire the lease and modify the blob.

upvoted 2 times

AymanAkk 10 months ago

Selected Answer: D

The Lease Blob operation creates and manages a lock on a blob for write and delete operations. The lock duration can be 15 to 60 seconds, or can be infinite. In versions prior to 2012-02-12, the lock duration is 60 seconds
https://learn.microsoft.com/en-us/rest/api/storageservices/lease-blob

upvoted 1 times

AymanAkk 10 months ago

also this
Pessimistic concurrency for blobs
To lock a blob for exclusive use, you can acquire a lease on it. When you acquire the lease, you specify the duration of the lease. A finite lease may be valid from between 15 to 60 seconds. A lease can also be infinite, which amounts to an exclusive lock.
https://learn.microsoft.com/en-us/azure/storage/blobs/concurrency-manage?tabs=dotnet

upvoted 1 times

Knightie 10 months, 4 weeks ago

Selected Answer: B

B is correct, the issue is that one staff sent out the receipt but another staff deleted it. When the customer open, the receipt is gone. Extending the lease only give a limited time to access but it does not retain a snapshot of a copy for the customer to retrieve. So Snapshot is right.

upvoted 1 times

Knightie 10 months, 4 weeks ago

my bad, lease is right, it's a lock.
https://learn.microsoft.com/en-us/rest/api/storageservices/lease-blob

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: D

D. Use blob leases to prevent concurrency problem

upvoted 1 times

gmishra88 1 year ago

Sure, BlobLease is an option but that does not require me to read a huge scenario. I wonder what the scenario mean for this question. And I wonder why not A as well because Etag is a "property" !!

upvoted 1 times

OlivierPaudex 1 year, 2 months ago

I think this one will close the discussion

https://docs.microsoft.com/en-us/azure/storage/blobs/concurrency-manage?tabs=dotnet

To lock a blob for exclusive use, you can acquire a lease on it.
Answer is D

upvoted 1 times

Shotare 1 year, 4 months ago

Selected Answer: D

D is correct, although I have managed to solve such problem in real life using blobs metadata so I am not sure whether C could be correct aswell. But D is safer bet

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

Answer: D. Use blob leases to prevent concurrency problems
Reason: To lock a blob for exclusive use, you can acquire a lease on it. When you acquire the lease, you specify the duration of the lease. A finite lease may be valid from between 15 to 60 seconds.

Reference: https://docs.microsoft.com/en-us/azure/storage/blobs/concurrency-manage?tabs=dotnet#pessimistic-concurrency-for-blobs

upvoted 2 times

massnonn 1 year, 8 months ago

Selected Answer: D

Answer is D

upvoted 3 times

vilainchien 1 year, 8 months ago

Selected Answer: D

Use blob leases to prevent concurrency problems

upvoted 2 times

leonidn 1 year, 9 months ago

Selected Answer: D

For processing purposes "lease" is the simplest option.

upvoted 2 times

PhilLI 1 year, 9 months ago

It's quite vague. Receipts are uploaded to Azure Files. Some people place a file, wait for a while, then delete it, and upload it via the webinterface. Should we cater for this situation as well where duplicates may be introduced?
File will be copied to blob by the code, but it is not clear how long it will stay there, and how to prevent rereading the same file. I would either use blob storage events, or timestamps to filter the files to be processed. If there is a risk to read the same file twice (shortly after upload when filtering on timestamp) then a lease seems ok: if you can't acquire the lease then another instance of your code is already working on it.

upvoted 1 times

PhilLI 1 year, 9 months ago

I guess the concurrency is caused by the CRON job every 5 minutes. If processing of the available files takes longer than 5 minutes, concurrency issues may be caused.

upvoted 1 times

RajMasilamani 2 years, 1 month ago

https://docs.microsoft.com/en-us/rest/api/storageservices/lease-blob

The answer would be Blob lease.

Release, to free the lease if it is no longer needed so that another client may immediately acquire a lease against the blob.

upvoted 1 times

RajMasilamani 2 years, 1 month ago

Answer is D

upvoted 1 times

rhr 2 years, 6 months ago

Given answer is correct

upvoted 2 times

azurelearner666 2 years, 4 months ago

No, it's not. It's D, use blob leases.

upvoted 2 times

Question #2 *Topic 20*

Introductory Info

Case study -

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

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

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

To start the case study -

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

Background -

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web

API project that is deployed as an Azure Web App.

Overall architecture -

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing -

Employees may upload receipts in two ways:

Uploading using an Azure Files mounted folder

Uploading using the web application

Data Storage -

Receipt and employee information is stored in an Azure SQL database.

Documentation -

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for

Azure File upload, and instructions on how to configure the mounted folder.

Solution details -

Users table -

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application -

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing -

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob

Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging -
Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled.
Application Insights must always contain all log messages.

Requirements -

Receipt processing -
Concurrent processing of a receipt must be prevented.

Disaster recovery -
Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security -
User's SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
Receipt data must always be encrypted at rest.
All data must be protected in transit.
User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues -

Upload format issue -
Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File
Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal
Server error page.

Capacity issue -
During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

Log capacity issue -
Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Application code -

Processing.cs -

```
PC01 public static class Processing
PC02 {
PC03   public static class Function
PC04   {
PC05     [FunctionName("IssueWork")]
PC06     public static async Task Run([TimerTrigger("0 */5 * * * *")] TimerInfo timer, ILogger
log)
PC07     {
PC08       var container = await GetCloudBlobContainer();
PC09       foreach (var fileItem in await ListFiles())
PC10       {
PC11         var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12         var ms = new MemoryStream();
PC13         await file.DownloadToStreamAsync(ms);
PC14         var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15         await blob.UploadFromStreamAsync(ms);
PC16
PC17       }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob(CloudBlockBlob sourceBlob)
PC20     {
PC21        . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25       var cloudBlobClient = new CloudBlobClient(new Uri(". . ."), await GetCredentials());
PC26
PC27       await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28       return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32        . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36        . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(". . .");
PC38   }

PC39 }
```

Database.cs -

```
DB01 public class Database
DB02 {
DB03   private string ConnectionString =
DB04
DB05   public async Task<object> LoadUserDetails(string userId)
DB06   {
DB07
DB08     return await policy.ExecuteAsync(async () =>
DB09     {
DB10       using (var connection = new SqlConnection(ConnectionString))
DB11       {
DB12         await connection.OpenAsync();
DB13         using (var command = new SqlCommand("…", connection))
DB14         using (var reader = command.ExecuteReader())
DB15         {
DB16           …
DB17         }
DB18       }
DB19     });
DB20   }

DB21 }
```

ReceiptUploader.cs -

```
RU01 public class ReceiptUploader
RU02 {
RU03   public async Task UploadFile(string file, byte[] binary)
RU04   {
RU05     var httpClient = new HttpClient();
RU06     var response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU07     while (ShouldRetry(response))
RU08     {
RU09       response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU10     }
RU11   }
RU12   private bool ShouldRetry(HttpResponseMessage response)
RU13   {
RU14
RU15   }
RU16 }
```

ConfigureSSE.ps1 -

```
CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "..." -AccountName "..."
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "..."
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "..."
CS04  Set-AzureRmKeyVaultAccessPolicy `
CS05   -VaultName $keyVault.VaultName `
CS06   -ObjectId $storageAccount.Identity.PrincipalId `
CS07
CS08
CS09  Set-AzureRmStorageAccount `
CS10   -ResourceGroupName $storageAccount.ResourceGroupName `
CS11   -AccountName $storageAccount.StorageAccountName `
CS12   -EnableEncryptionService File `
CS13   -KeyvaultEncryption `
CS14   -KeyName $key.Name
CS15   -KeyVersion $key.Version `
CS16   -KeyVaultUri $keyVault.VaultUri
```

Question

You need to resolve the capacity issue.

What should you do?

A. Convert the trigger on the Azure Function to an Azure Blob storage trigger

B. Ensure that the consumption plan is configured correctly to allow scaling

C. Move the Azure Function to a dedicated App Service Plan

D. Update the loop starting on line PC09 to process items in parallel

Correct Answer: *D*

If you want to read the files in parallel, you cannot use forEach. Each of the async callback function calls does return a promise. You can await the array of promises that you'll get with Promise.all.

Scenario: Capacity issue: During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

```
PC08        var container = await GetCloudBlobContainer();
PC09        foreach (var fileItem in await ListFiles())
PC10        {
PC11          var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12          var ms = new MemoryStream();
PC13          await file.DownloadToStreamAsync(ms);
PC14          var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15          await blob.UploadFromStreamAsync(ms);
PC16
PC17        }
```

Reference:

https://stackoverflow.com/questions/37576685/using-async-await-with-a-foreach-loop

Community vote distribution

D (100%)

trance13 Highly Voted 2 years, 6 months ago

Receipts are uploaded to the File Storage (not Blob Storage) which does not support triggers.
Concurrent processing of a (SINGLE!) receipt must be prevented - so parallel processing is OK.
So answer D.

upvoted 28 times

ZodiaC 2 years, 3 months ago

1000% D !!!!!!!! CORRECT!

upvoted 3 times

PaulMD Highly Voted 2 years, 6 months ago

Cleared AZ-204 today, the question appeared, the option "D" was not there, but a "replace the solution with durable functions". I went for that.

upvoted 18 times

SnakePlissken 2 years, 5 months ago

Durable Functions is just an extension of Azure Functions that lets you write stateful functions in a serverless compute environment. It's not the solution for this problem. It proves that another answer is correct. I think that answer B, scaling, is the best option.

upvoted 7 times

HumbleYolo 1 year, 4 months ago

in consumption plan, scaling is automatic.

upvoted 1 times

ferut 2 years, 5 months ago

Durable functions will let the consumer get an immediate (async) response, but the processing remains. The duration till the file appears on the website doesn't change.
Doing the processing in parallel will make a change.

upvoted 3 times

ning 2 years, 2 months ago

Correct, if one instance of time trigger function is running, then there will not be a second instance starts, even when 5 minutes pass ... For a durable function, it can make sure immediate returns to allow second instance to start ...

upvoted 3 times

leonidn 1 year, 9 months ago

That makes sense. Running parallel tasks is not good practice for functions. Here we cannot predict the degree of parallelizm. But using durable function is the best choice.

upvoted 2 times

AndySmith Most Recent 1 week, 3 days ago

I believe it's "B" - since issue happens in busy period when CPU is over-utilized. Then only reasonable action will be to scale. And for that we should properly configure Consumption Plan.
"D" - could be an answer if Q was about slow speed of processing in normal situation, when CPU resources in enough. In this case, I/O operations are the bottleneck. But, if we try to spawn more thread when CPU is already super-busy, it would even worsen user experience.
And it's not "C" since Dedicated Plan is used in very specific situation. Exerpt:
"Consider a dedicated App Service plan in the following situations:
- You have existing, underutilized VMs that are already running other App Service instances.
- You want to provide a custom image on which to run your functions."

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: D

D. Update the loop starting on line PC09 to process items in parallel

upvoted 1 times

gmishra88 1 year ago

"Concurrent processing of a receipt must be prevented."
Microsoft has added this line as a red-herring to make the question taker not think parallelism as an option? What does "processing" mean here? What is "a receipt"? That in combination with the listFiles() method. Does "a receipt" contain multiple files? Does "processing of a receipt" in Microsoft dictionary mean uploading (processing, microsoft?) of the multiple files in "a receipt"
If the answer has durable functions then go for it without thinking deep. The requirements looks like a requirement for asynchronous processing because employees get an email (asynchronous) later. But any other answer is just not right and the question could send an intelligent developer (Microsoft excluded) into a loop of thoughts.

upvoted 2 times

gmishra88 1 year ago

Appreciate all the Microsoft-Technology-developers finding innovative reasons for the answers. But what is not clear is what that listFiles() method do. Which files are returned. That's a lot of assumptions to say you can do upload in parallel without knowing what files and their sizes. No wonder Microsoft-technologies are so buggy

upvoted 1 times

ReniRechner 1 year, 7 months ago

Selected Answer: D

A. Convert the trigger on the Azure Function to an Azure Blob storage trigger
=> won't help because we have Azure Fileshare
B. Ensure that the consumption plan is configured correctly to allow scaling
=> Trigger is time based. Multiple instances scanning the same folder => bad idea; also clearly stated in the requirements that parallel processing is not allowed
C. Move the Azure Function to a dedicated App Service Plan
=> the Trigger every 5 seconds should keep the function "alive". The function work is also not CPU bound so I cannot see a real benefit for ASP in this scenario
D. Update the loop starting on line PC09 to process items in parallel
=> might help.
D2 (alternative to D as by PaulMD) replace the solution with durable functions
=> looks even better than D

If D2 is an option I'd go for that.
Maybe they realized that the current "D" is not a really good solution and D2 is also way more "azure"
Otherwise D.

upvoted 11 times

kozchris 1 year, 8 months ago

The answer is C since this is a cold start problem.
"When using Azure Functions in the dedicated plan, the Functions host is always running, which means that cold start isn't really an issue."
https://azure.microsoft.com/en-us/blog/understanding-serverless-cold-start/

upvoted 1 times

coffecold 1 year ago

No, trigger is timed [Timertrigger....], so function execution never sleeps..

upvoted 1 times

eMax 1 year, 9 months ago

The answer reference is about JavaScript, not C# :))))))

upvoted 1 times

asdasdasg2 1 year, 9 months ago

D is not correct - while this would speed up performance, the prompt states that users report high delay during BUSY PERIODS. Clearly, the fact that it does not upload files in parallel would not solve that.

The problem must be that the consumption plan is not scaling the function app correctly to handle the load. C could theoretically help, but B is better.

Correct answer: B

upvoted 1 times

ning 2 years, 2 months ago

Only thing possible is D ...
File mount, is not blob storage, so cannot be trigger ...
This is a time trigger, so scale up will not help, only one instance will run ...
Only leave us with D

upvoted 4 times

Onuoa92 2 years, 5 months ago

Nobody is given us a correct answer

upvoted 2 times

ZodiaC 2 years, 3 months ago

D is 1000% correct

upvoted 1 times

Molte 1 year, 9 months ago

your 1000% comments under every single question does not help at all!

upvoted 16 times

[Removed] 2 years, 6 months ago

I vote for B. Reasoning:

A. Convert the trigger on the Azure Function to an Azure Blob storage trigger
> We are not dealing with a defect, but a performance degradation, so this would not help.

B. Ensure that the consumption plan is configured correctly to allow scaling
> It seems that "Maximum Scale Out Limit" is set to a value not appropriate for the usage pattern

C. Move the Azure Function to a dedicated App Service Plan
> Wont help.

D. Update the loop starting on line PC09 to process items in parallel
> I don't think it is a good idea to call an async method from within a foreach loop, also not from within Parallel.ForEach.

https://stackoverflow.com/questions/23137393/parallel-foreach-and-async-await

upvoted 5 times

warchoon 6 months, 3 weeks ago

Don't use parallel extensions in Azure. There are special Azure constructions for it.

upvoted 1 times

anirbanzeus 2 years, 5 months ago

well the function is started by a timer, meaning that the "event" that should trigger the scaling won't increase. Hence I do not think B is the correct choice (Ref: https://docs.microsoft.com/en-us/azure/azure-functions/event-driven-scaling).
Considering that we are uploading receipts to a Azure file storage A is also incorrect.
In the given scenario D is the one that makes the most sense.

upvoted 2 times

VR 2 years, 6 months ago

So what is the answer?

upvoted 4 times

kwaazaar 2 years, 6 months ago

D is the right answer, since the loop picks up all files in the container and scaling would make the files being processed more than once, potentially.
Change feed is not supported for file shares, so D is the only remaining option (though ugly as hell).

upvoted 2 times

jokergester 2 years, 7 months ago

A and C - converting to blob trigger with dedicated plan not consumption to avoid cold start and high availability of the function
D - is not enough since the the trigger is scheduled to every 5 mins - so users will still need to wait even it is already have been processed.

upvoted 1 times

nicolaus 2 years, 5 months ago

Answer is C. A is not possible as reports can also be uploaded using Azure Files. Consumption plan has a cold start (up to 10 minutes), so moving to dedicated plan will help

upvoted 6 times

kabbas 4 months ago

I agree, parallel processing is not going to help much here since the listfile() is already doing that. A dedicated plan will provide more resources

upvoted 1 times

PhilLI 1 year, 9 months ago

2 questions about C:
Will a cold start be an issue at all when it is triggered by a time trigger?
Could it be a dedicated App Service plan has stronger CPU allowing to process the files faster?
Besides that: if parallel processing is an option, I would go for that specially with the autoscaling options of a consumption plan (but where time trigger doesn't help?)

upvoted 2 times

aperez1979 2 years, 7 months ago

I think is better option change the trigger. A

upvoted 4 times

Beitran 2 years, 7 months ago

Indeed: "Concurrent processing of a receipt must be prevented."

upvoted 3 times

trance13 2 years, 6 months ago

No one wants to process a single receipt concurrently, each distinct file will be processed in parallel.

upvoted 6 times

kwaazaar 2 years, 6 months ago

Only blobs support change feed, not fileshare, which is used here.
https://docs.microsoft.com/nl-nl/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal

upvoted 3 times

Question #3 *Topic 20*

Introductory Info

Case study -

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

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

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

To start the case study -

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

Background -

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web

API project that is deployed as an Azure Web App.

Overall architecture -

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing -

Employees may upload receipts in two ways:

Uploading using an Azure Files mounted folder

Uploading using the web application

Data Storage -

Receipt and employee information is stored in an Azure SQL database.

Documentation -

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for

Azure File upload, and instructions on how to configure the mounted folder.

Solution details -

Users table -

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application -

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing -

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob

Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging -
Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled.
Application Insights must always contain all log messages.

Requirements -

Receipt processing -
Concurrent processing of a receipt must be prevented.

Disaster recovery -
Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security -
User's SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
Receipt data must always be encrypted at rest.
All data must be protected in transit.
User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues -

Upload format issue -
Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File
Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal
Server error page.

Capacity issue -
During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

Log capacity issue -
Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Application code -

Processing.cs -

```
PC01 public static class Processing
PC02 {
PC03   public static class Function
PC04   {
PC05     [FunctionName("IssueWork")]
PC06     public static async Task Run([TimerTrigger("0 */5 * * * *")] TimerInfo timer, ILogger
log)
PC07     {
PC08       var container = await GetCloudBlobContainer();
PC09       foreach (var fileItem in await ListFiles())
PC10       {
PC11         var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12         var ms = new MemoryStream();
PC13         await file.DownloadToStreamAsync(ms);
PC14         var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15         await blob.UploadFromStreamAsync(ms);
PC16
PC17       }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob(CloudBlockBlob sourceBlob)
PC20     {
PC21        . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25       var cloudBlobClient = new CloudBlobClient(new Uri(". . ."), await GetCredentials());
PC26
PC27       await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28       return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32        . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36        . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(". . .");
PC38   }

PC39 }
```

Database.cs -

```
DB01 public class Database
DB02 {
DB03   private string ConnectionString =
DB04
DB05   public async Task<object> LoadUserDetails(string userId)
DB06   {
DB07
DB08     return await policy.ExecuteAsync(async () =>
DB09     {
DB10       using (var connection = new SqlConnection(ConnectionString))
DB11       {
DB12         await connection.OpenAsync();
DB13         using (var command = new SqlCommand("…", connection))
DB14         using (var reader = command.ExecuteReader())
DB15         {
DB16           …
DB17         }
DB18       }
DB19     });
DB20   }

DB21 }
```

ReceiptUploader.cs -

```
RU01 public class ReceiptUploader
RU02 {
RU03   public async Task UploadFile(string file, byte[] binary)
RU04   {
RU05     var httpClient = new HttpClient();
RU06     var response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU07     while (ShouldRetry(response))
RU08     {
RU09       response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU10     }
RU11   }
RU12   private bool ShouldRetry(HttpResponseMessage response)
RU13   {
RU14
RU15   }
RU16 }
```

ConfigureSSE.ps1 -

```
CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "..." -AccountName "..."
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "..."
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "..."
CS04  Set-AzureRmKeyVaultAccessPolicy `
CS05   -VaultName $keyVault.VaultName `
CS06   -ObjectId $storageAccount.Identity.PrincipalId `
CS07
CS08
CS09  Set-AzureRmStorageAccount `
CS10   -ResourceGroupName $storageAccount.ResourceGroupName `
CS11   -AccountName $storageAccount.StorageAccountName `
CS12   -EnableEncryptionService File `
CS13   -KeyvaultEncryption `
CS14   -KeyName $key.Name
CS15   -KeyVersion $key.Version `
CS16   -KeyVaultUri $keyVault.VaultUri
```

Question

You need to resolve the log capacity issue.

What should you do?

A. Create an Application Insights Telemetry Filter

B. Change the minimum log level in the host.json file for the function

C. Implement Application Insights Sampling

D. Set a LogCategoryFilter during startup

Correct Answer: *C*

Scenario, the log capacity issue: Developers report that the number of log message in the trace output for the processor is too high, resulting in lost log messages.

Sampling is a feature in Azure Application Insights. It is the recommended way to reduce telemetry traffic and storage, while preserving a statistically correct analysis of application data. The filter selects items that are related, so that you can navigate between items when you are doing diagnostic investigations. When metric counts are presented to you in the portal, they are renormalized to take account of the sampling, to minimize any effect on the statistics.

Sampling reduces traffic and data costs, and helps you avoid throttling.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling

Community vote distribution

C (100%)

MikiStieger Highly Voted 2 years, 5 months ago

I had the exam last Friday and I notice that they change one sentence in case study from:
Application Insights must always contain all log messages.
to:
Application Insights currently contain all log messages.
So I chosen C. Implement Application Insights Sampling.

upvoted 29 times

OPT_001122 11 months, 2 weeks ago

so you mean
IF Application Insights must always contain all log messages. then A. Create an Application Insights Telemetry Filter

If Application Insights currently contain all log messages then C. Implement Application Insights Sampling.

upvoted 2 times

gmishra88 1 year ago

If that is the case then all three options are back on the table:
1. Sampling
2. Filtering
3. Adjusting the log level
4. Whatever the 4th mean

It totally depend on the situation. Practically you try the option and find the best solution. But I think Microsoft wants us to say Sampling.

upvoted 2 times

coffecold 1 year ago

I don't agree. Sampling does not meet "Application Insights must always contain ALL log messages" since it drops items.
https://learn.microsoft.com/en-us/azure/azure-monitor/app/sampling#how-sampling-works
I think it would be simply adjusting the log level in the app.

upvoted 3 times

coffecold 1 year ago

The fourth means : filtering on log level (in webjobs SDK 2.0).
https://learn.microsoft.com/en-us/azure/app-service/webjobs-sdk-how-to

upvoted 1 times

kozchris Highly Voted 1 year, 8 months ago

The answer is "B".

From the problem description:
"Developers report that the number of log messages in the trace output for the processor is too high"

The keywords are "too high".

The answer is to change the minimum logging level.

upvoted 6 times

chingdm 1 year, 8 months ago

agreed, it also said "trace output for the processor is too high, resulting in lost log messages" so it seems that sampling is already enabled and the logs are too many so resulting in lost messages since sampling only takes only portion of all the logs, so it helps if changing the minimum log level to reduce the logs or disable sampling but this is not in the answer selection.

upvoted 2 times

gmishra88 1 year ago

That's right, sherlock. But I don't think that is what Microsoft is looking for here. Those are red-herring options given to doom developers who never tried their flashy, state-of-the-art sampling feature. Yes, but these options depends on the developer and what exactly is shown in the logs and the level of logs.

upvoted 1 times

chingdm 1 year, 8 months ago

also this requirement "Application Insights must always contain all log messages." will eliminate enable sampling answer because app insights will filter some logs during sampling

upvoted 2 times

kabbas Most Recent 4 months ago

Selected Answer: C

Using sampling also allows you to preserve all your logs but reduces throttling

upvoted 1 times

JH81 4 months, 1 week ago

Selected Answer: C

"Application Insights must always contain all log messages." is a stupid requirement. What does that even mean? It would help if it indicated to what level. If some jr programmer set the logging level to Verbose then that's the problem and the answer should be B change the minimum logging level. A and C would indicate you are not meeting the requirement to log all messages. D is not a real answer. But I guess C sampling is what MS is looking for since ""Sampling also helps you avoid Application Insights throttling your telemetry." which is the issue.
https://learn.microsoft.com/en-us/azure/azure-monitor/app/sampling

upvoted 2 times

halfway 6 months, 4 weeks ago

Selected Answer: C

According to the reference below: "Sampling is a feature in Application Insights. It's the recommended way to reduce telemetry traffic, data costs, and storage costs, while preserving a statistically correct analysis of application data. Sampling also helps you avoid Application Insights throttling your telemetry. "
Ref: https://learn.microsoft.com/en-us/azure/azure-monitor/app/sampling?tabs=net-core-new

upvoted 1 times

gmishra88 1 year ago

Filtering or Sampling cannot be the answers considering, as far as I can guess from Microsoft documentation, that it is about NOT sending the telemetry (log entries?) to Application Insights. So, how can the argument that the requirement says "must contain all log message" (whatever that means in the real world). Sampling/Filtering and Chaning-log-levels have all the same effect. From the Microsoft documentation I could not really be clear of what exactly "sampling" is but it looks like it is not about "trace" but about metrics.

upvoted 3 times

coffecold 1 year ago

Agree, exactly because of that:
Filtering and Sampling drop messages. So if you must not drop messages between app and app Insights, the only thing that is left is to reduce logged messages (C)

upvoted 1 times

coffecold 1 year ago

I mean option B.

upvoted 1 times

mandusya 1 year, 10 months ago

We had similar a few weeks ago
sampling worked fine

upvoted 1 times

phvogel 1 year, 12 months ago

From https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-filtering-sampling
"Sampling reduces the volume of telemetry without affecting your statistics. "
"Filtering with telemetry processors lets you filter out telemetry in the SDK before it's sent to the server" (and the other options would also eliminate the log traces)
So only sampling will meet the requirement of containing all log messages.

upvoted 2 times

ning 2 years, 2 months ago

I am thinking of A to exclude TraceWrite logging ...
Since all log messages are required, so you cannot sampling or change log levels, those are leading to lose log entries ...
LogCategoryFilter I cannot find anything with that from documentation ...

upvoted 3 times

[Removed] 2 years, 6 months ago

A. Create an Application Insights Telemetry Filter

> A filter can be created by either implementing ITelemetryProcessor or by implementing ITelemetryInitializer. However MS recommends to use sampling

https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-filtering-sampling#filtering

B. Change the minimum log level in the host.json file for the function

> Can be ruled out because req. says app insights must always contain all log messages

C. Implement Application Insights Sampling

> Is the recommended way to reduce telemetry traffic, data costs and storage costs
> https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling
> https://blog.ramondeklein.nl/2017/05/05/filtering-application-insights/

D. Set a LogCategoryFilter during startup

> No idea what that is, but feels like pointint in the same direction as the B does

upvoted 5 times

gmishra88 1 year ago

Bingo. It does say "Application Insights must always contain all log messages." I did miss that sentence. This is an eyesight test than a developer exam.

upvoted 3 times

BrettusMaximus 2 years, 6 months ago

It can't be C. - AI Sampling.
Rule 1: Application Insights must always contain all log messages
If you read the Sampling doco carefully. It does filter logs and does not record all the transaction records (but yes it keeps a count for statistics only

(See adaptive and fixed sampling)) https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling

Fact 2: The processor also has TraceWriter logging enabled.
TraceWriter logs are generally used for debugging and are not "Official" transactional logs. https://stackify.com/logging-azure-functions/
Options A, B or D are candidates to filter these logs.
Option B and D would stop the actual logs being generated but may also remove some transactional logs. It would also not let the developers do their debugging (the purpose of TraceWriter in the first place).

Thus this leaves the only option A.
A. Create an Application Insights Telemetry Filter (to filter the trace writer logging)

upvoted 3 times

aperez1979 2 years, 7 months ago

I think it could be the b

upvoted 1 times

MrZoom 2 years, 7 months ago

The case states "The processor also has TraceWriter logging enabled.
Application Insights must always contain all log messages."

For this reason B isn't an option, and neither is A or D. These would all change what log messages are sent to AI, which isn't according to reqs.
So this leaves C. Sampling just groups messages together with a count, causing less traffic to AI but the same results.

upvoted 9 times

kwaazaar 2 years, 6 months ago

Sampling is enabled by default and can be turned off (althoiugh AI ingress may still drop entries when overloaded). For metrics sampling does exactly as MrZoom describes.

As I see it, irrelevant logs must be prevented. My first step would be to adjust the minimum loglevel,if possible. A telemetry processor to filter telemetry/logs technically could work to, but is intended for filtering specific entries.

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

B is correct.
Because it says "Application Insights must always contain all log messages.", sampling is not a valid answer. In fact, sampling could be the very reason that log messages are lost.
See: https://docs.microsoft.com/en-us/azure/azure-monitor/app/asp-net-trace-logs
"I don't see some log entries that I expected".

upvoted 2 times

gmishra88 1 year ago

That's an excellent point. Kudos. I will need three days to prepare a legal answer for this question. How much did you say the exam allows? 180 minutes?
So, according to you the filter is also not an option? Nice, but I do not know how can I not see the vague requirement that says "must contain all log messages". Hope I do not read that line in the certification exam and so, I can say B is the right answer like a good developer's natural instinct.

upvoted 2 times

oceane0316 10 months, 1 week ago

Thank you gmishra88 for the funny comments you left along the way throughout this painful study session.

upvoted 3 times

naivecoder786 8 months ago

Haha so true ! I think he is the same guy who was from start but is now showin as "Removed User" xD

upvoted 2 times

Topic 21 - Testlet 22

Question #1 *Topic 21*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -
The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
HOTSPOT -
You need to implement event routing for retail store location data.
Which configurations should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Event data	Configuration
Source	Azure Blob Storage Azure Event Grid Azure Service Bus Azure Event Hub
Receiver	Azure Event Grid Azure Event Hub Azure Service Bus Azure Blob Storage
Handler	Azure Function App Azure Logic App Azure Event Grid Azure Blob Storage

Correct Answer:

Answer Area

Event data	Configuration
Source	Azure Blob Storage
Receiver	Azure Event Grid
Handler	Azure Logic App

Box 1: Azure Blob Storage -
Azure event publishers and event handlers are at the core of the Event Grid routing-service. Event Grid listens to Azure event publishers, such as Blog Storage, then reacts by routing specific events to Azure event handlers, such as WebHooks. You can easily control this entire process at a granular level through event subscriptions and event filters.

Box 2: Azure Event Grid -
Azure Event Grid is a highly scalable event-routing service that listens for specific system events, then reacts to them according to your precise specifications. In the past, event handling has relied largely on polling ג€" a high latency, low-efficiency approach that can prove prohibitively expensive at scale.

Box 3: Azure Logic App -
Event Grid's supported event handlers currently include Event Hubs, WebHooks, Logic Apps, Azure Functions, Azure Automation and Microsoft Flow.
Reference:
https://www.appliedi.net/blog/using-azure-event-grid-for-highly-scalable-event-routing

le129 Highly Voted 1 year, 1 month ago
Shouldn't it be Azure Function App for the handler?
upvoted 29 times

Dani_ac7 1 year, 1 month ago
Yes, logic app is not mentioned
upvoted 5 times

willchenxa 1 year, 1 month ago
Agree, there is no logic app in the solution.
upvoted 6 times

OPT_001122 Highly Voted 11 months, 2 weeks ago
Retail store locations -
Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.
Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Logic app is not mentioned anywhere in the case study.
So #3 should be Function app?
upvoted 7 times

juanckar Most Recent 3 months, 3 weeks ago
This was on the exam (July 2023). Went with blob/grid/function. Scored 917

upvoted 5 times

tawanda_belkis 6 months, 3 weeks ago

Is this microsoft question or Azure Function app

upvoted 1 times

warchoon 6 months, 3 weeks ago

"process sales financials" needs messages not events, "immediately" => Box 2: Azure Service Bus
"Azure Functions must" => Box 3: Azure Function App

upvoted 2 times

surprise0011 6 months, 2 weeks ago

exactly. Event should not carry much information, so Azure Service Bus should be utilized

upvoted 1 times

warchoon 6 months, 3 weeks ago

ok, this is wrong. Because 21/1 has no Service Bus option and "financial" is probably for the audit trigger only.
So Box:2 must be Event Grid

upvoted 1 times

warchoon 6 months, 3 weeks ago

... I mean 29/1 has no Service Bus

upvoted 2 times

Topic 22 - Testlet 23

Question #1 *Topic 22*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

LabelMaker app -

Coho Winery produces, bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends data to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the LabelMaker application to include artwork and text for custom label designs.

Requirements. Data -

You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using SQL.

Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Requirements. Security -

You have the following security requirements:

Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.

▪

External partners must use their own credentials and authenticate with their organization's identity management solution.

External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.

Storage of e-commerce application settings must be maintained in Azure Key Vault.

E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).

Conditional access policies must be applied at the application level to protect company content.

The LabelMaker application must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

Requirements. LabelMaker app -

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure

Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture -

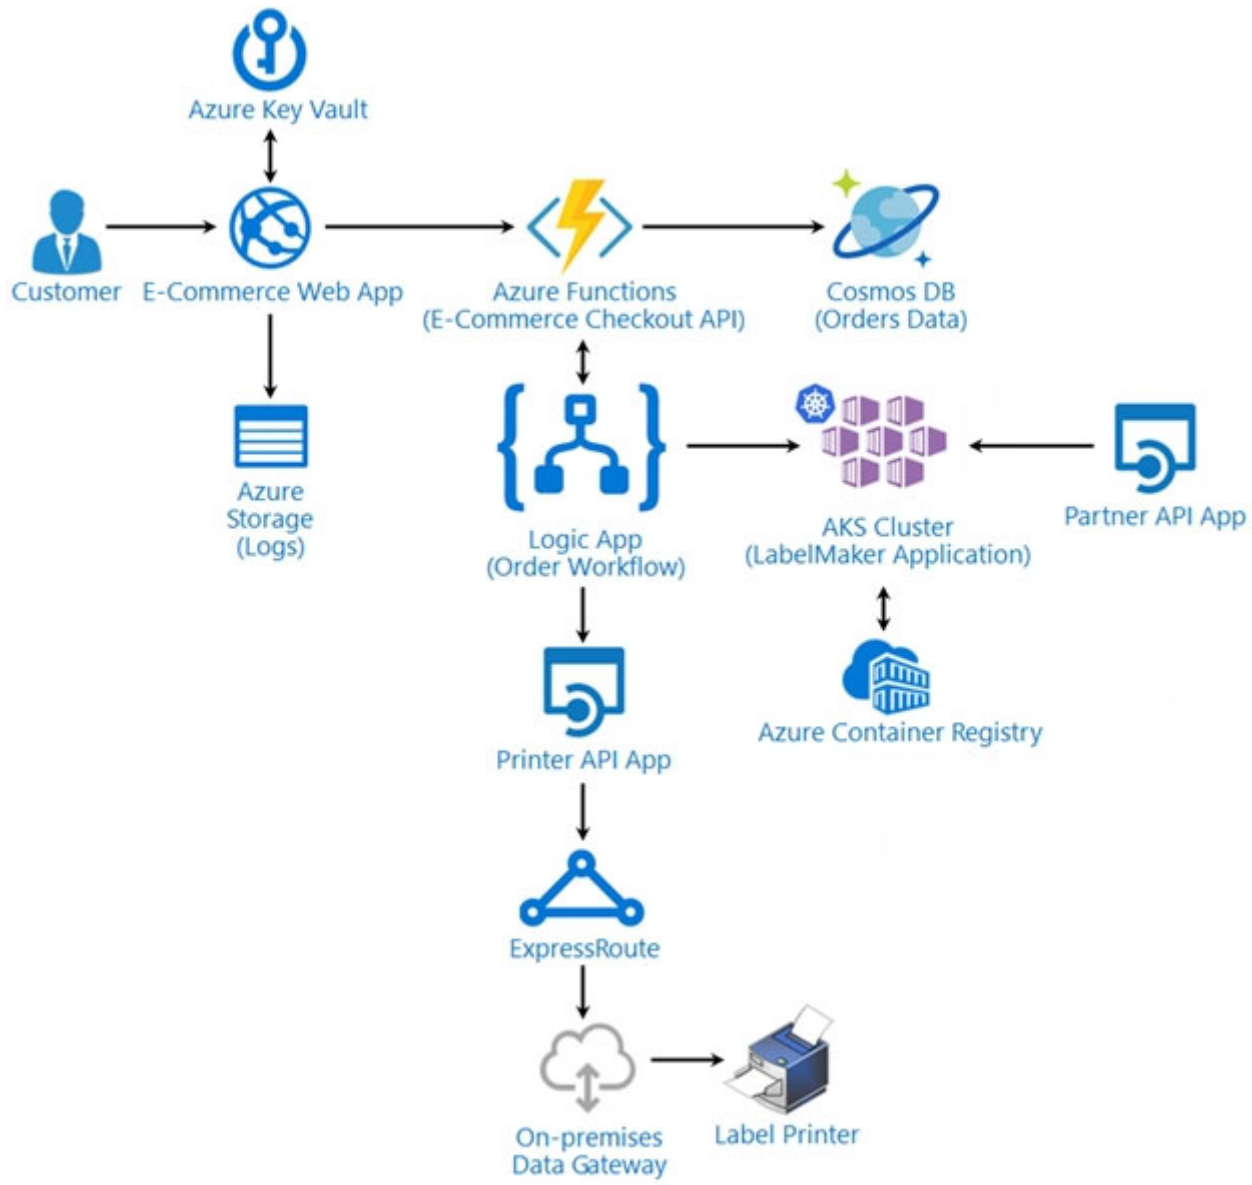

Issues -

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order.json -

Relevant portions of the app files are shown below. Line numbers are included for reference only.

This JSON file contains a representation of the data for an order that includes a single item.

Order.json

```
01 {
02   "id" : 1,
03   "customers" : [
04   {
05      "familyName" : "Doe",
06      "givenName" : "John",
07      "customerid" : 5
08    }
09   ],
10   "line_items" : [
11   {
12     "fulfillable_quantity" : 1,
13     "id": 6,
14     "price" : "199.99" ,
15     "product_id" : 7513594,
16     "quantity": 1,
17     "requires_shipping" : true ,
18     "sku": "SFC-342-N" ,
19     "title" : "Surface Go" ,
20     "vendor" : "Microsoft" ,
21     "name" : "Surface Go - 8GB" ,
22     "taxable" : true ,
23     "tax_lines" : [
24     {
25       "title" : "State Tax" ,
26       "price" : "3.98" ,
27       "rate" : 0.06
28      }
29     ],
```

Question

You need to troubleshoot the order workflow.

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

NOTE: Each correct selection is worth one point.

A. Review the API connections.

B. Review the activity log.

C. Review the run history.

D. Review the trigger history.

Correct Answer: *CD*

Scenario: The order workflow fails to run upon initial deployment to Azure.

Check runs history: Each time that the trigger fires for an item or event, the Logic Apps engine creates and runs a separate workflow instance for each item or event. If a run fails, follow these steps to review what happened during that run, including the status for each step in the workflow plus the inputs and outputs for each step.

Check the workflow's run status by checking the runs history. To view more information about a failed run, including all the steps in that run in their status, select the failed run.

Example:

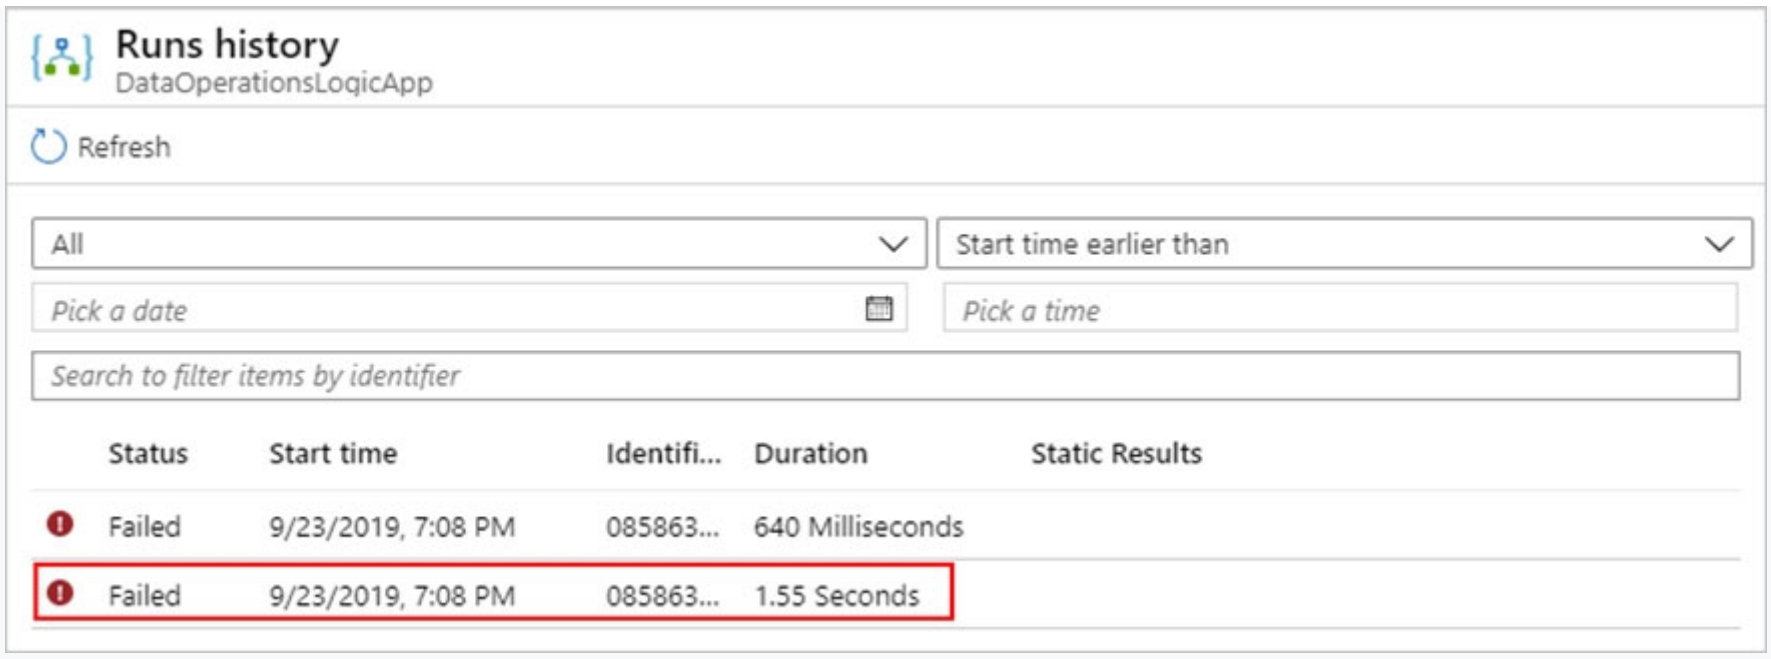

Check the trigger's status by checking the trigger history

To view more information about the trigger attempt, select that trigger event, for example:

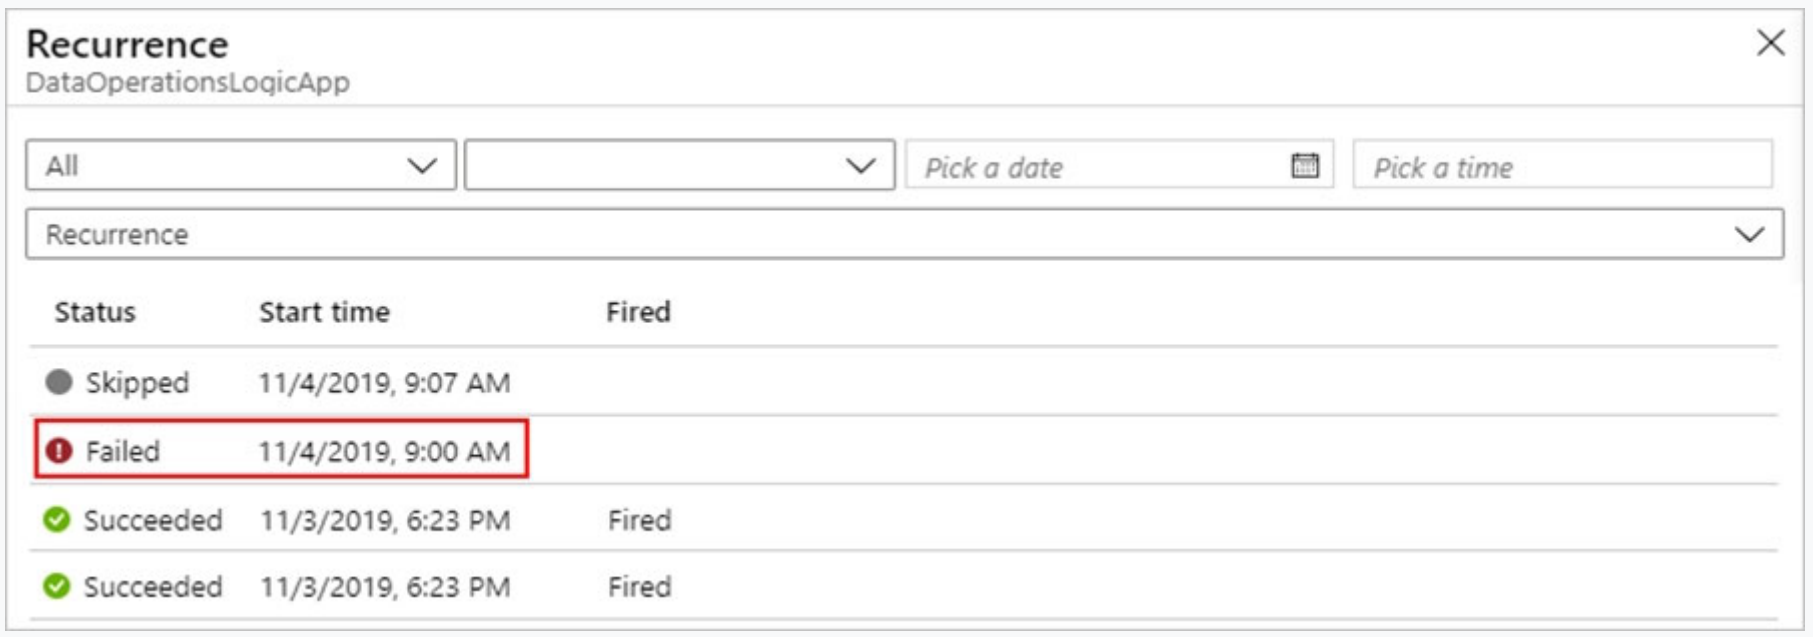

Reference:

https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-diagnosing-failures

Community vote distribution

CD (80%) BD (20%)

7ack Highly Voted 2 years, 4 months ago

Answer seems correct

upvoted 8 times

mkqwerty 1 year, 11 months ago

very convenience answer

upvoted 3 times

Net_IT Most Recent 1 day, 18 hours ago

Selected Answer: CD

C and D are the correct answers.
https://learn.microsoft.com/en-us/azure/logic-apps/logic-apps-diagnosing-failures?tabs=consumption

upvoted 1 times

macobuzi 1 month, 3 weeks ago

Selected Answer: BD

Correct answer should be B and D. As for why not C, the specific issue is that the "order workflow fails to run upon initial deployment to Azure", so the initial deployment itself might not have generated any run history if the workflow didn't run successfully. Run history typically records instances of workflow executions.

upvoted 1 times

macobuzi 1 month, 3 weeks ago

Please Ignore my answer, the default answer was right because we have the Logic App in the architecture!

upvoted 1 times

VarunDashora26 2 months, 1 week ago

Correct answer in my opinion should be B and D.
"The order workflow fails to run upon initial deployment to Azure."
This means you will not get any info in C. Run History. And A. Checking the API connection will be helpful only when the workflow is up and running and API calls are made.

upvoted 1 times

macobuzi 1 month, 3 weeks ago

I agree, even chatGPT chooses B and D.
upvoted 1 times

macobuzi 1 month, 3 weeks ago
Sorry, C and D is correct because the app uses Logic App, It can still check the Run History: "Each time that the trigger fires for an item or event, the Logic Apps engine creates and runs a separate workflow instance for each item or event. If a run fails, follow these steps to review what happened during that run, including the status for each step in the workflow plus the inputs and outputs for each step"
upvoted 1 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: CD
C. Review the run history.
D. Review the trigger history.
upvoted 3 times

adilkhan 8 months, 3 weeks ago
why the answer is C and D?
upvoted 1 times

deborahQ 1 year ago
It says: The order workflow fails to run upon initial deployment to Azure.
It seems to me, it never runs.
upvoted 3 times

le129 1 year, 1 month ago
what is the difference between run history and trigger history
upvoted 2 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

LabelMaker app -

Coho Winery produces, bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends data to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the LabelMaker application to include artwork and text for custom label designs.

Requirements. Data -

You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using SQL.

Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Requirements. Security -

You have the following security requirements:

Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.

▪

External partners must use their own credentials and authenticate with their organization's identity management solution.

External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.

Storage of e-commerce application settings must be maintained in Azure Key Vault.

E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).

Conditional access policies must be applied at the application level to protect company content.

The LabelMaker application must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

Requirements. LabelMaker app -

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure

Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture -

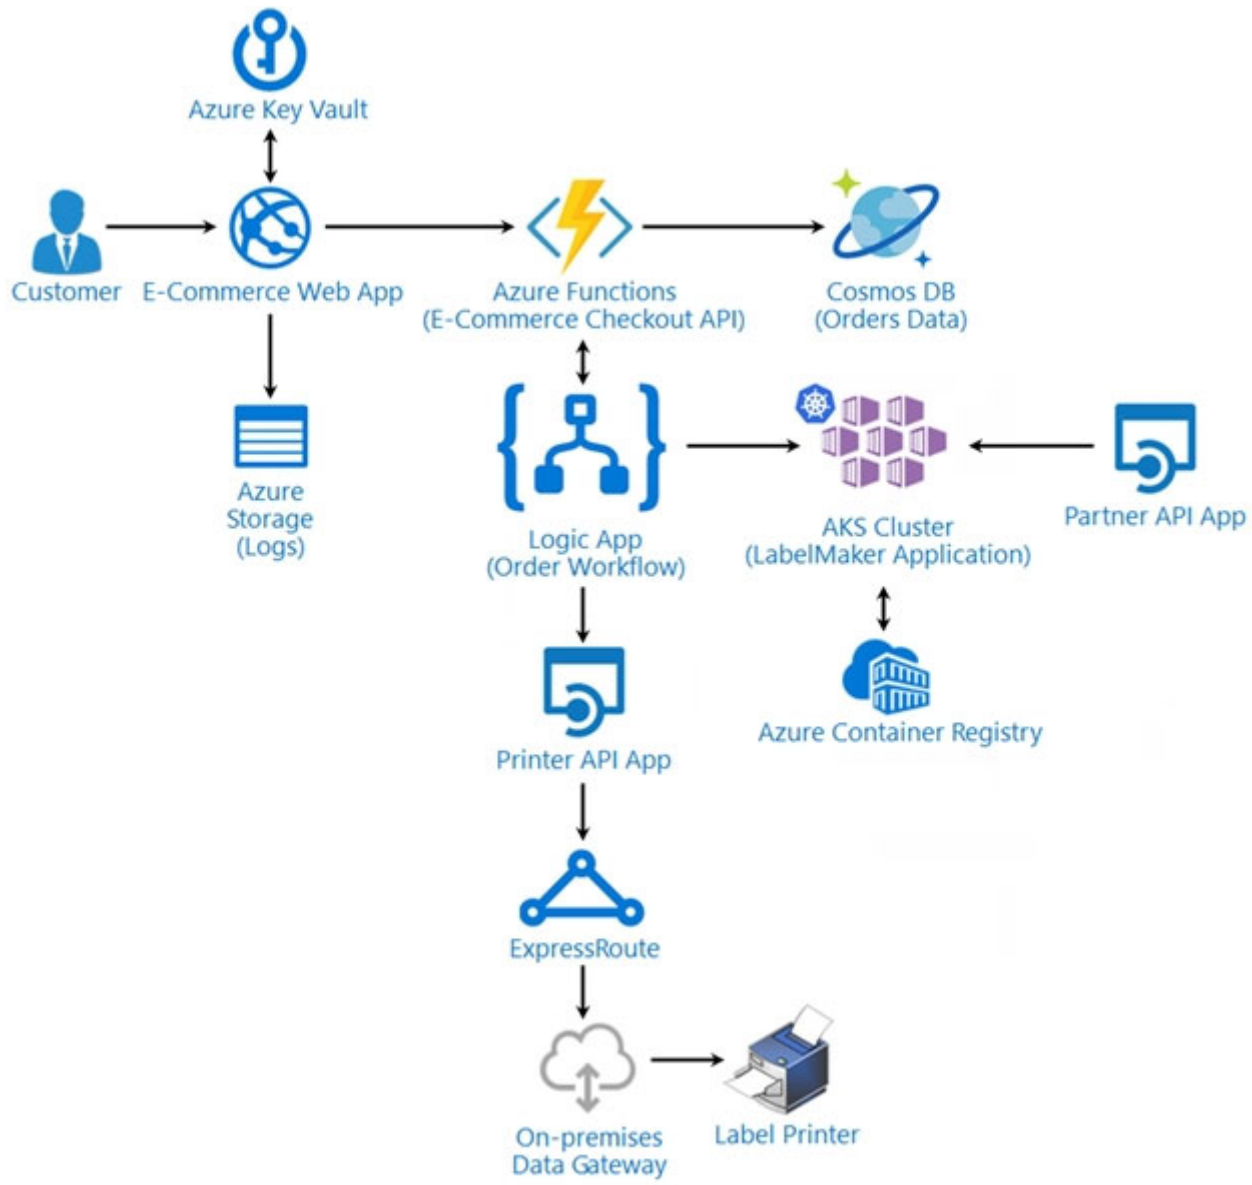

Issues -

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order.json -

Relevant portions of the app files are shown below. Line numbers are included for reference only.

This JSON file contains a representation of the data for an order that includes a single item.

Order.json

```
01 {
02   "id" : 1,
03   "customers" : [
04   {
05      "familyName" : "Doe",
06      "givenName" : "John",
07      "customerid" : 5
08    }
09   ],
10   "line_items" : [
11   {
12     "fulfillable_quantity" : 1,
13     "id": 6,
14     "price" : "199.99" ,
15     "product_id" : 7513594,
16     "quantity": 1,
17     "requires_shipping" : true ,
18     "sku": "SFC-342-N" ,
19     "title" : "Surface Go" ,
20     "vendor" : "Microsoft" ,
21     "name" : "Surface Go - 8GB" ,
22     "taxable" : true ,
23     "tax_lines" : [
24     {
25       "title" : "State Tax" ,
26       "price" : "3.98" ,
27       "rate" : 0.06
28      }
29     ],
```

Question

HOTSPOT -

You need to update the order workflow to address the issue when calling the Printer API App.

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:

Answer Area

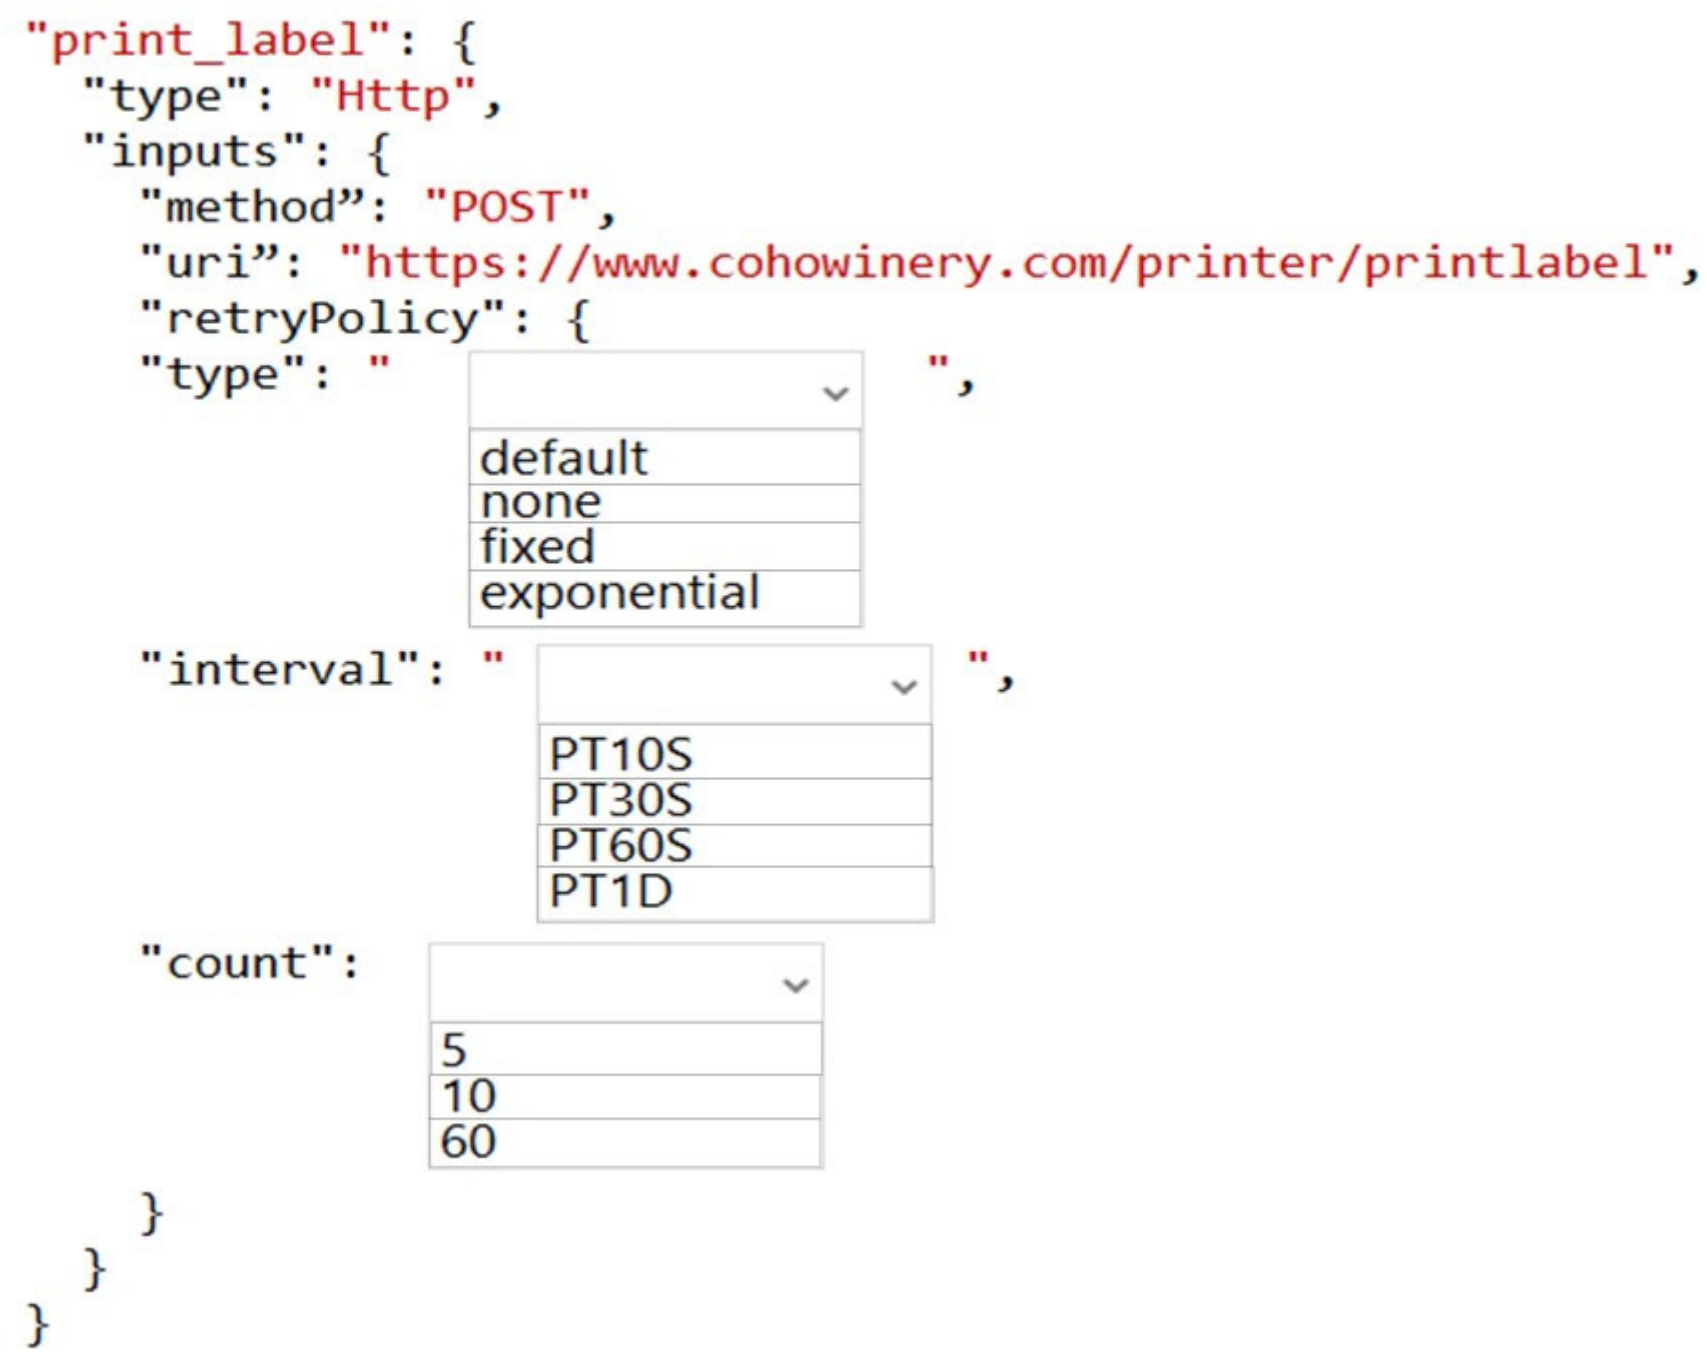

Correct Answer:

Answer Area

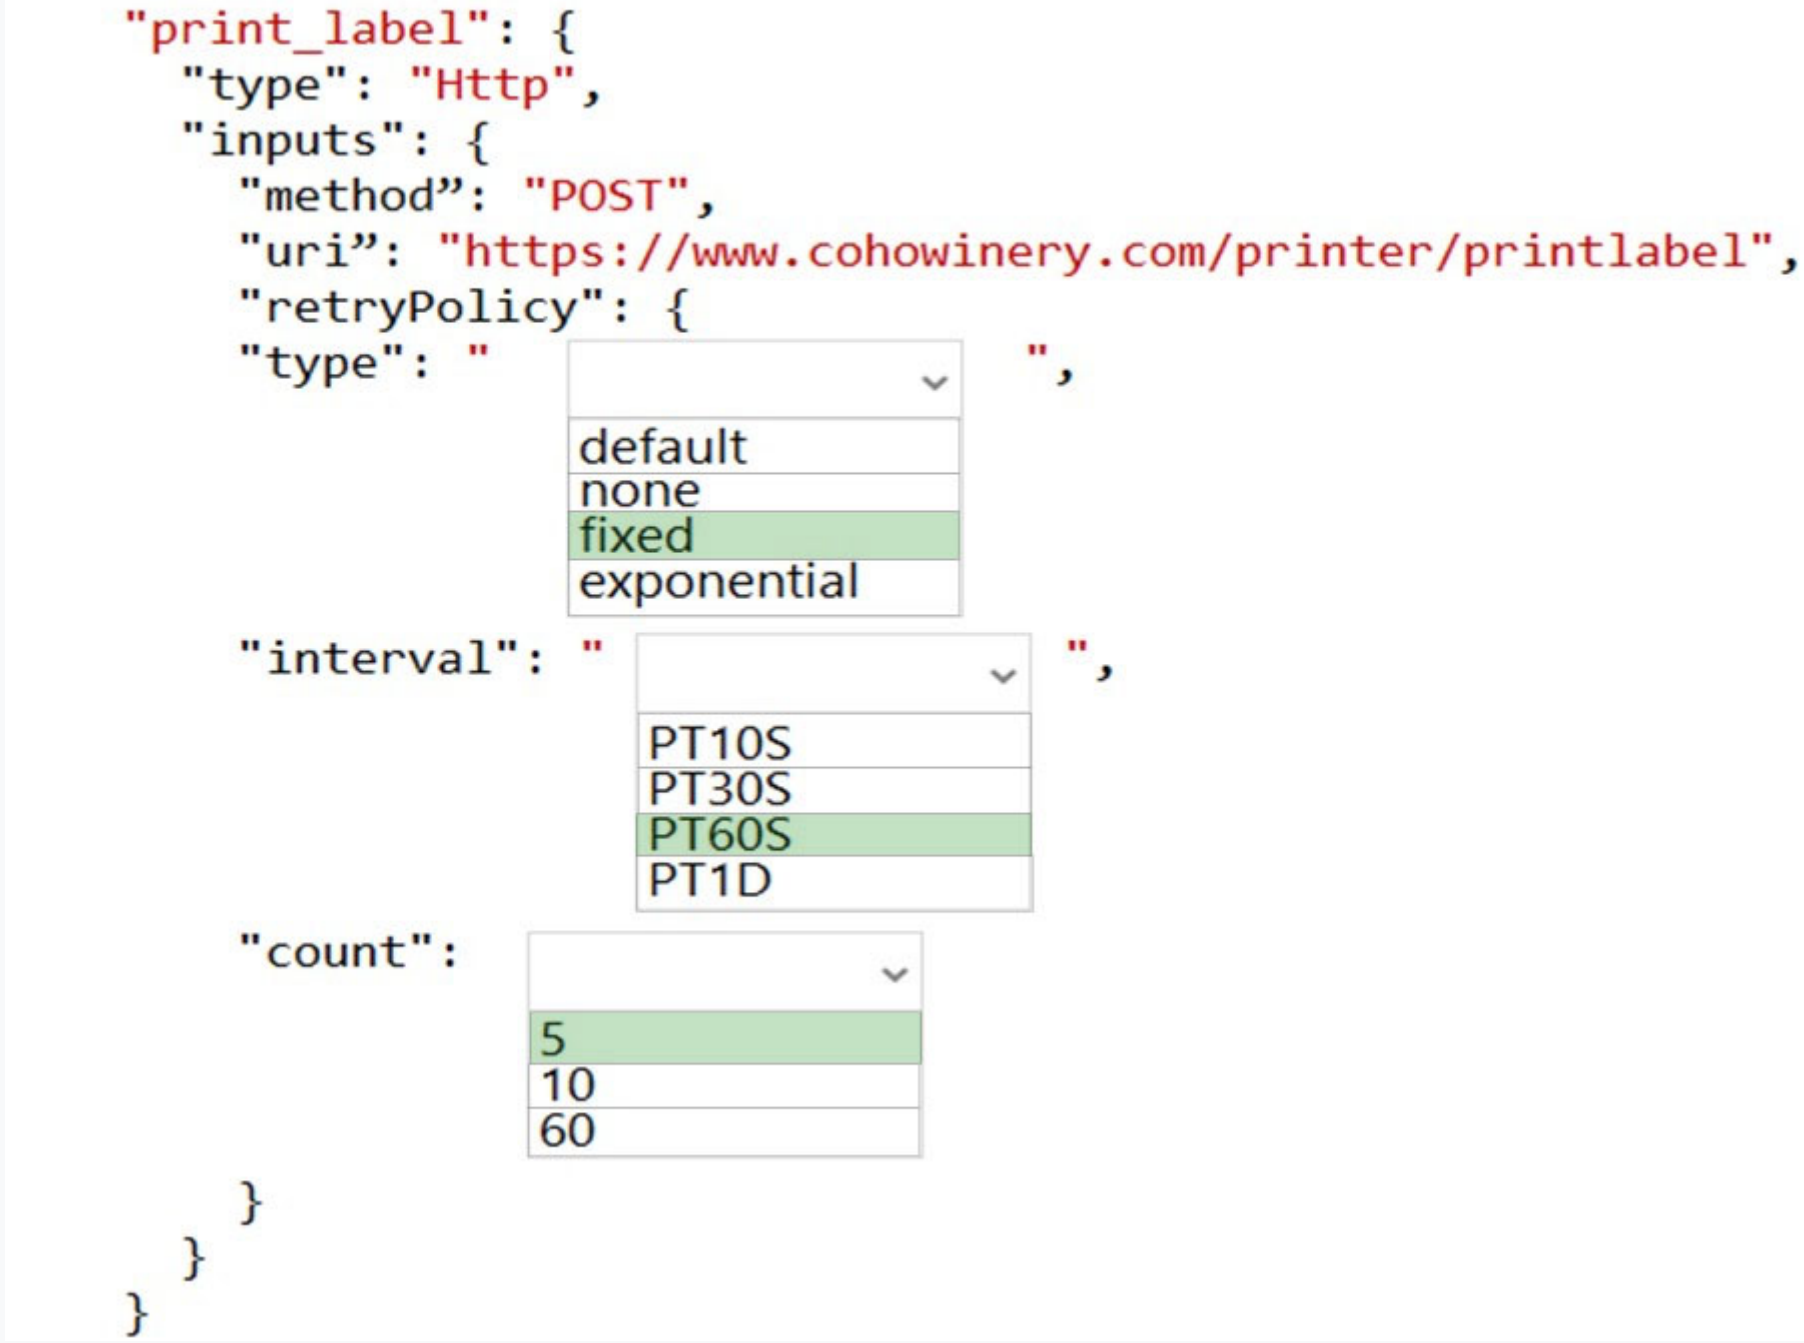

Box 1: fixed -

The 'Default' policy does 4 exponential retries and from my experience the interval times are often too short in situations.

Box 2: PT60S -

We could set a fixed interval, e.g. 5 retries every 60 seconds (PT60S).

PT60S is 60 seconds.

Scenario: Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

Box 3: 5 -

Reference:

https://michalsacewicz.com/error-handling-in-power-automate/

jay158 Highly Voted 2 years, 3 months ago

1. Fixed
2. PT10S
3. 5

PT10 means retry after 10sec ,
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-exception-handling

upvoted 51 times

ensa 2 years ago

Currently it is 10, because fail after 10 min. we should increase the time interval and I am not sure to choose PT30 or PT60

upvoted 2 times

ensa 2 years ago

Sorry, Jay158 is right. There is requirement in the question

upvoted 2 times

Branners 2 years, 3 months ago

Agreed. PT10S means a 10 second delay between each try defined as the retry value.

upvoted 4 times

vilainchien 1 year, 8 months ago

Hi, I don't understand why PT10S is the correct answer. The function could fail immediately and the printer only accept 5 requests in 1 minute. If you set 5 retries with an interval of 10 seconds, you will have 6 requests in 1 minute?

upvoted 1 times

Knightie 10 months, 4 weeks ago

it will only retry 5 times and then fail and not going to retry anymore. So no 6 times in a min.

upvoted 1 times

chingdm 1 year, 8 months ago

you have to also consider the time of executing the printer request, say printer timed-out in 0 seconds and first attempt issued, say the printer timed-out in 10 seconds, then request would wait 10 seconds before trying again, so the second retry attempt would be in the 20th seconds, so it will be less than 5 attempts in 1 minute.

upvoted 1 times

Gwak Highly Voted 2 years, 3 months ago

I think Given ans is right "PT60 and count 5" menas try 5 times in 60sec.

upvoted 8 times

j888 2 years, 2 months ago

If you read the following:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-exception-handling#fixed-interval
pt60s = 60s delay between each attempt.
PT10s will be the better answer, 5 attempts with each 10 seconds delay satisfy the requirement “The label printer must only receive up to 5 attempts within one minute”.

upvoted 7 times

yosri_c_sharp Most Recent 4 months, 2 weeks ago

Logical App no more included in az-204

upvoted 3 times

adilkhan 8 months, 3 weeks ago

Ans is correct https://learn.microsoft.com/en-us/azure/logic-apps/logic-apps-exception-handling?tabs=consumption Fixed, PT60S, 5

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer:
1. Fixed
2. PT10S
3. 5

https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-exception-handling

upvoted 2 times

leonidn 1 year, 9 months ago

Fixed + PT10S, 5.
None and default do not require "interval" parameter. Exponential requires additional parameters. "Fixed" is the only option that fits the schema. PT10S means that after the timeout occurs, the logic app waits 10 more seconds and tries again. It means that there will be up to 3 attempts in a minute. That fits the requirement. Extended interval does not make sense here from the user experience perspective. It is more desirable to reduce the time interval. As well as setting more than 5 attempts does not make sense. If printer cannot recover after 5 attempts it's not make sense to continue trying.

upvoted 4 times

RajMasilamani 2 years, 1 month ago

Answer to address or fix the issue.So the time interval increased to 60 seconds and for 5 attempts.As problem stated for the given fixed values its giving an error.To overcome this we need to increase the retry time .

upvoted 1 times

koreshulya 2 years, 4 months ago

I guess interval should be PT10S.
“The label printer must only receive up to 5 attempts within one minute”.

upvoted 7 times

vokep77043 2 years, 2 months ago

One minute is 60 seconds - so it won't be 5 attempts... lol

upvoted 1 times

lexowe9241 2 years, 1 month ago

https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-exception-handling#fixed-interval "This retry policy attempts to get the latest news two more times after the first failed request with a 30-second delay between each attempt:"
"retryPolicy": {
"type": "fixed",
"interval": "PT30S",
"count": 2
}

Back to our case: PT10S and 5 attempts means : 1st attempt plus 5 more. That is 60seconds in total.

upvoted 6 times

Topic 23 - Testlet 24

Question #1 *Topic 23*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment -

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk Server 2016. The VM runs the following workflows:

Ocean Transport `" This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.

Inland Transport `" This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

Container API `" This API provides container information including weight, contents, and other attributes.

Location API `" This API provides location information regarding shipping ports of call and trucking stops.

Shipping REST API `" This API provides shipping information for use and display on the shipping website.

Shipping Data -

The application uses MongoDB JSON document storage database for all container and transport information.

Shipping Web Site -

The site displays shipping container tracking information and container contents. The site is located at http://shipping.wideworldimporters.com/

Proposed solution -

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:

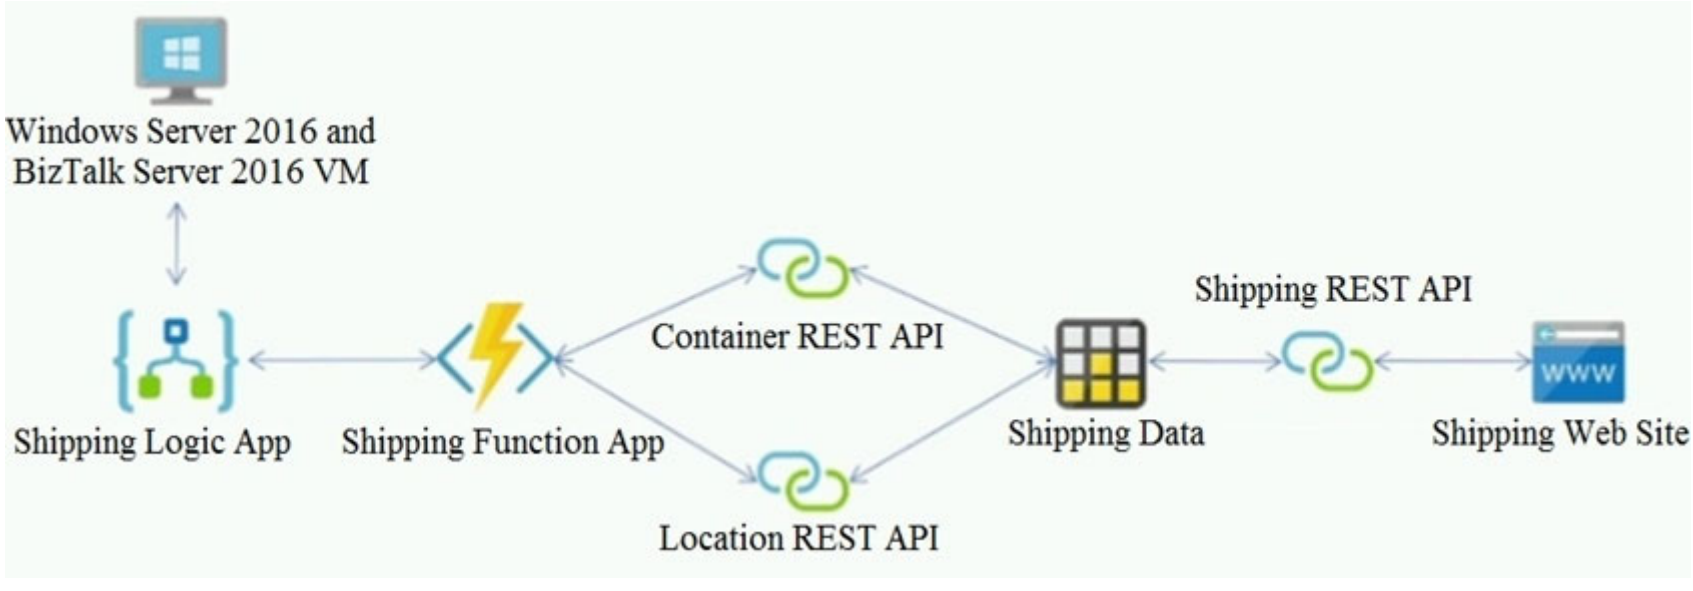

Requirements -

Shipping Logic app -
The Shipping Logic app must meet the following requirements:
Support the ocean transport and inland transport workflows by using a Logic App.
Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app -
Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

REST APIs -
The REST API's that support the solution must meet the following requirements:
Secure resources to the corporate VNet.
Allow deployment to a testing location within Azure while not incurring additional costs.
Automatically scale to double capacity during peak shipping times while not causing application downtime.
Minimize costs when selecting an Azure payment model.

Shipping data -
Data migration from on-premises to Azure must minimize costs and downtime.

Shipping website -
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Issues -

Windows Server 2016 VM -
The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

Shipping website and REST APIs -
The following error message displays while you are testing the website:
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource.
Origin 'http://test.wideworldimporters.com/' is therefore not allowed access.

Question
DRAG DROP -
You need to support the message processing for the ocean transport workflow.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.
Select and Place:

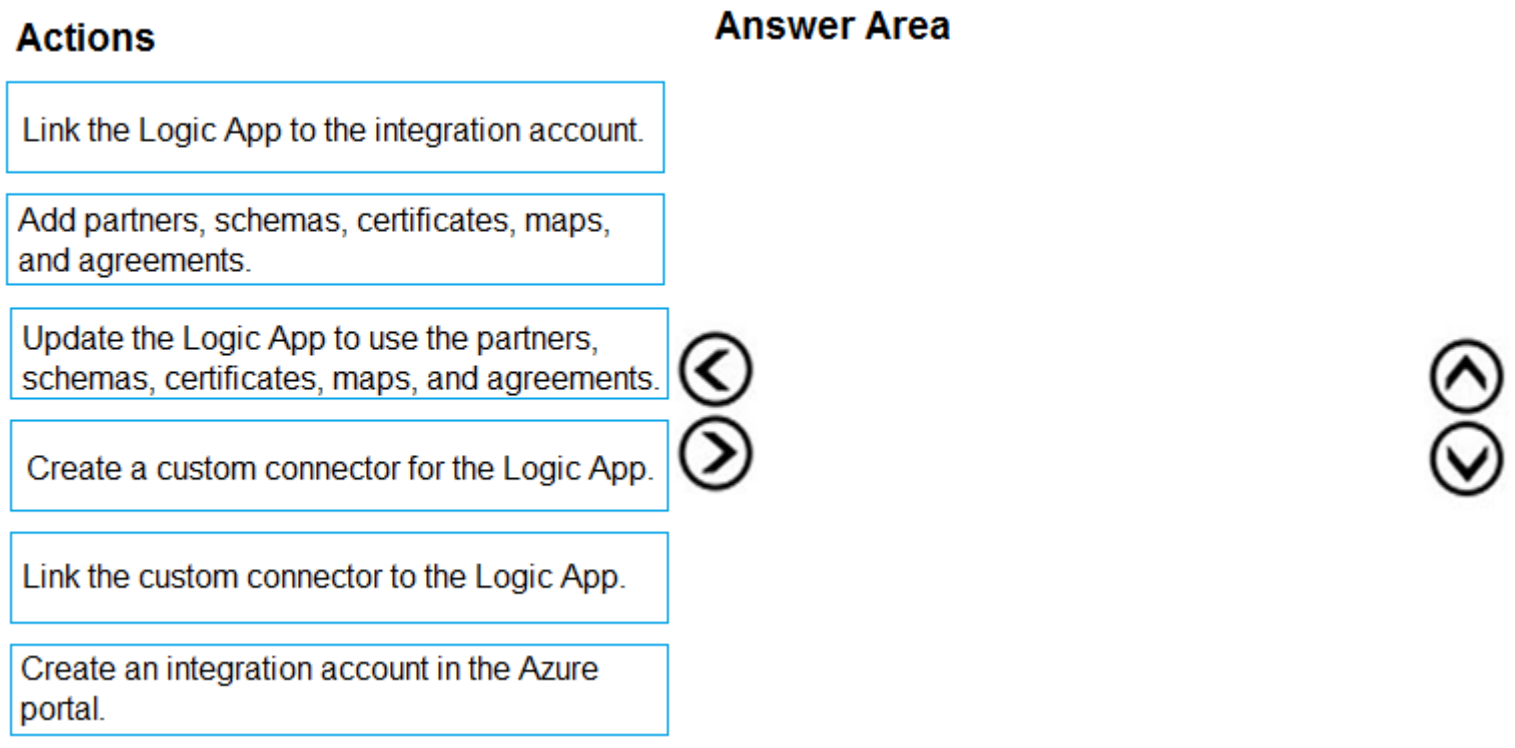

Correct Answer:

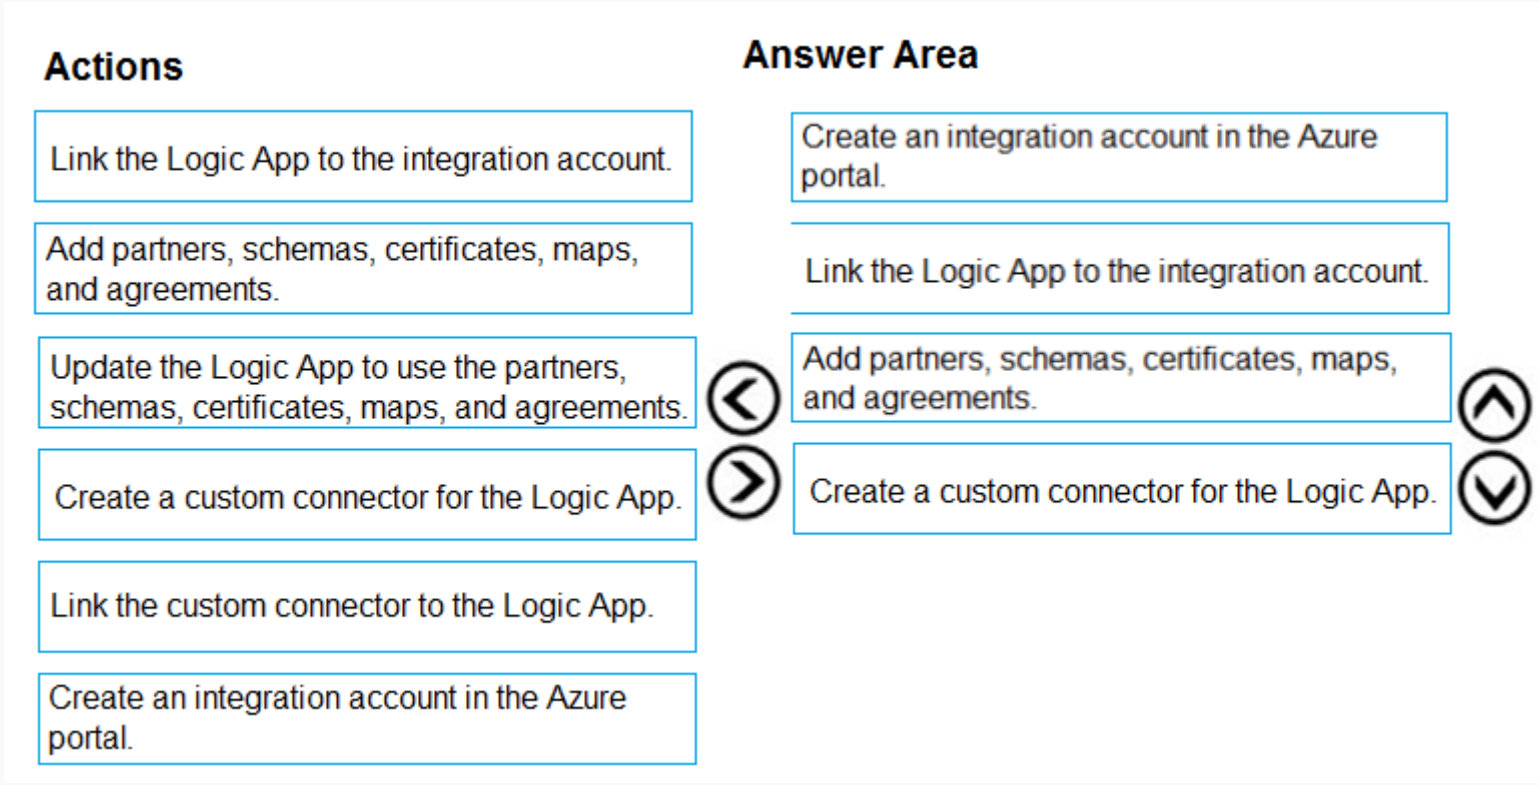

Step 1: Create an integration account in the Azure portal

You can define custom metadata for artifacts in integration accounts and get that metadata during runtime for your logic app to use. For example, you can provide metadata for artifacts, such as partners, agreements, schemas, and maps - all store metadata using key-value pairs.

Step 2: Link the Logic App to the integration account

A logic app that's linked to the integration account and artifact metadata you want to use.

Step 3: Add partners, schemas, certificates, maps, and agreements

Step 4: Create a custom connector for the Logic App.

Reference:

https://docs.microsoft.com/bs-latn-ba/azure/logic-apps/logic-apps-enterprise-integration-metadata

coffecold Highly Voted 1 year ago

This is about logic app, which is not part anymore of the exam

However :

1.Create integration account
2.Link Logic app
3.Add partners...
4.Update logic app

https://learn.microsoft.com/en-us/azure/logic-apps/add-artifacts-integration-service-environment-ise#create-integration-accounts

Why not custom connector? X12 is Managed connector, as it is opposed to a custom connector
https://learn.microsoft.com/en-us/azure/logic-apps/add-artifacts-integration-service-environment-ise#add-ise-connectors

upvoted 14 times

adilkhan Most Recent 8 months, 3 weeks ago

1. create integration accountWhen you're done, select Create.

Link your logic app to your integration account in the usual way.

Continue by adding resources to your integration account, such as trading partners and agreements.
https://learn.microsoft.com/en-us/azure/logic-apps/add-artifacts-integration-service-environment-ise#create-integration-accounts. Answers are correct!
To manage integration accounts in your ISE, see Manage your integration service environment.

upvoted 3 times

tester2023 8 months, 1 week ago

The reference article above supports the order of events.
1. Create logic app (this is assumed to be done in our case study)
2. Create Integration account
3. Link logic app
4. Add partners, schemas, certificates...
5. Create (add) a custom connector

upvoted 2 times

OPT_001122 10 months, 3 weeks ago

1. add parameters
2. create integration accounts
3. link logic app to integration accounts
4. update param in logic app

upvoted 1 times

OPT_001122 10 months, 2 weeks ago

disregard above ans

1.Create integration account
2.Link Logic app
3.Add partners...
4.Update logic app

upvoted 4 times

willchenxa 1 year, 1 month ago

it should be:
1. add parameters
2. create integration accounts
3. link logic app to integration accounts
4. update param in logic app

upvoted 4 times

ArturKon 1 year, 1 month ago

Source?

upvoted 6 times

sghaha 1 year, 5 months ago

Korean page
https://docs.microsoft.com/ko-kr/azure/logic-apps/logic-apps-overview

upvoted 1 times

Question #2 *Topic 23*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment -

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk Server 2016. The VM runs the following workflows:

Ocean Transport `" This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.

Inland Transport `" This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

Container API `" This API provides container information including weight, contents, and other attributes.

Location API `" This API provides location information regarding shipping ports of call and trucking stops.

Shipping REST API `" This API provides shipping information for use and display on the shipping website.

Shipping Data -

The application uses MongoDB JSON document storage database for all container and transport information.

Shipping Web Site -

The site displays shipping container tracking information and container contents. The site is located at http://shipping.wideworldimporters.com/

Proposed solution -

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:

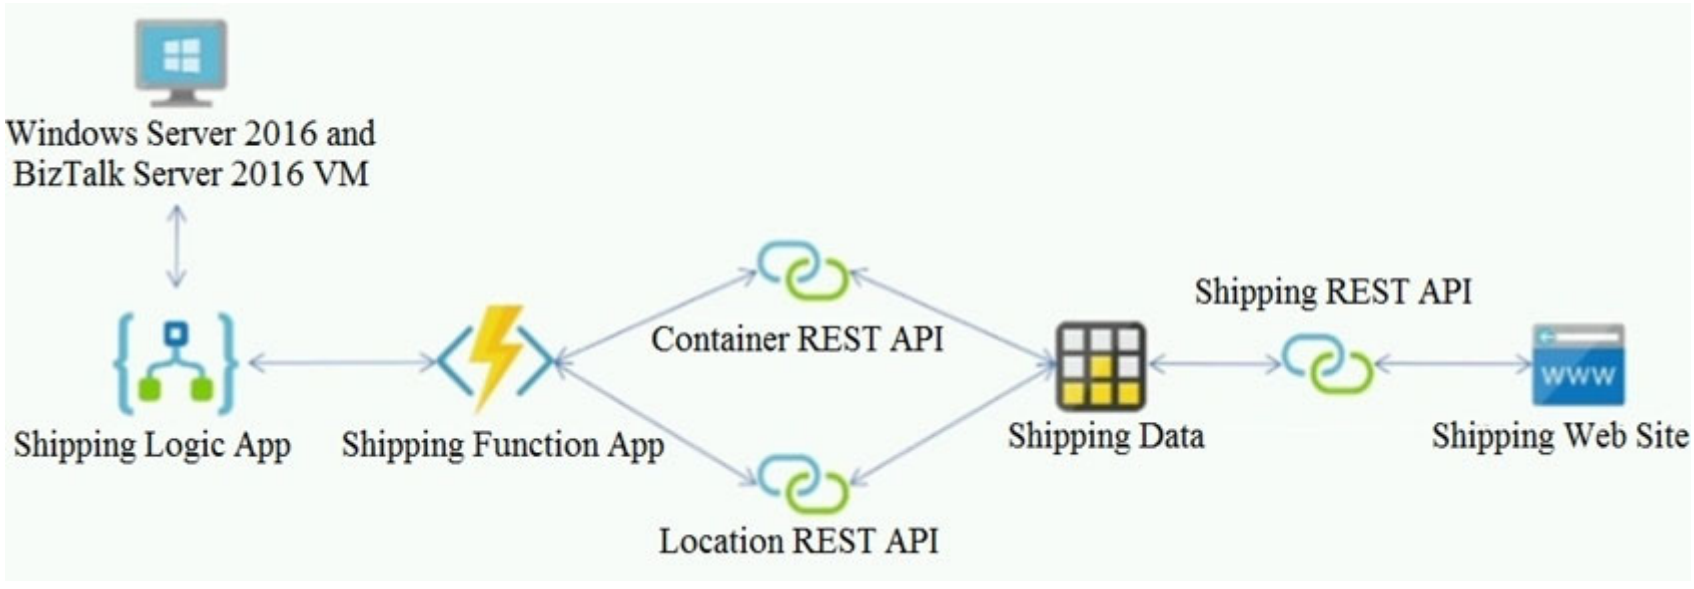

Requirements -

Shipping Logic app -
The Shipping Logic app must meet the following requirements:
Support the ocean transport and inland transport workflows by using a Logic App.
Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app -
Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

REST APIs -
The REST API's that support the solution must meet the following requirements:
Secure resources to the corporate VNet.
Allow deployment to a testing location within Azure while not incurring additional costs.
Automatically scale to double capacity during peak shipping times while not causing application downtime.
Minimize costs when selecting an Azure payment model.

Shipping data -
Data migration from on-premises to Azure must minimize costs and downtime.

Shipping website -
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Issues -

Windows Server 2016 VM -
The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

Shipping website and REST APIs -
The following error message displays while you are testing the website:
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource.
Origin 'http://test.wideworldimporters.com/' is therefore not allowed access.

Question
You need to support the requirements for the Shipping Logic App.
What should you use?

A. Azure Active Directory Application Proxy

B. Site-to-Site (S2S) VPN connection

C. On-premises Data Gateway

D. Point-to-Site (P2S) VPN connection

Correct Answer: *C*
Before you can connect to on-premises data sources from Azure Logic Apps, download and install the on-premises data gateway on a local computer. The gateway works as a bridge that provides quick data transfer and encryption between data sources on premises (not in the cloud) and your logic apps.
The gateway supports BizTalk Server 2016.
Note: Microsoft have now fully incorporated the Azure BizTalk Services capabilities into Logic Apps and Azure App Service Hybrid Connections.
Logic Apps Enterprise Integration pack bring some of the enterprise B2B capabilities like AS2 and X12, EDI standards support
Scenario: The Shipping Logic app must meet the following requirements:
☞ Support the ocean transport and inland transport workflows by using a Logic App.
☞ Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
☞ Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
☞ Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.
Reference:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-gateway-install

Community vote distribution

C (100%)

programmingbot Highly Voted 3 years, 3 months ago

Agreed: On-premises Data Gateway

upvoted 34 times

kakajos Most Recent 1 month ago

I was about to settle on S2S VPN connection until I read this: The On-premises data gateway acts as a bridge to provide quick and secure data transfer between on-premises data (data that isn't in the cloud) and several Microsoft cloud services. These cloud services include Power BI, PowerApps, Power Automate, Azure Analysis Services, and Azure Logic Apps

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: C

On-premises data Gateway

upvoted 2 times

BrettusMaximus 2 years, 6 months ago

B. Site to Site s2s VPN
Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.
I read this as legacy apps need to connect to the BizTalk server in Azure. This is how the work flow starts.

upvoted 1 times

Sylph 2 years, 8 months ago

The provided answer is probably correct because you can connect to REST/SOAP services over the Data Gateway using a custom connector.

Snip from documentation:
"You can also create custom connectors that connect to data sources over HTTP or HTTPS by using REST or SOAP. Although the gateway itself doesn't incur extra costs, the Logic Apps pricing model applies to these connectors and other operations in Azure Logic Apps."
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-gateway-connection#supported-data-sources

upvoted 1 times

paru123456789 2 years, 8 months ago

Answer: On premises data gateway

upvoted 1 times

fadikh 2 years, 8 months ago

Hasn't the VM been already migrated to Azure?

upvoted 3 times

coffecold 1 year ago

Exactly my first thoughts.. that VM is already in Azure.
"The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations"

upvoted 1 times

coffecold 1 year ago

But as none of the answers suit then, C is the best option. Let's pretend that the VM is on premises.

upvoted 1 times

khoant 2 years, 9 months ago

Answer is correct. The on-premises shipping application must be moved to Azure.

upvoted 1 times

amanp 2 years, 11 months ago

S2S uses a connection to connect the Vnet to on-premises private network. This does not help in logic app creation. Hence, the given answer is correct

upvoted 1 times

Brak 2 years, 9 months ago

The question does NOT ask if S2S VPN would "help in logic app creation."

The question does ask for on-prem connectivity for legacy apps. Data Gateway provides access to on-prem data sources for tools like PowerBI or Analysis Services (which are not needed by the workflow), not legacy apps or support final BizTalk migration. These need general on-prem to VNet connectivity.

The answer is S2S VPN.

upvoted 7 times

pac1311 2 years, 9 months ago

Your answer is incorrect, the question does ask for on-prem connection to legacy app, see the following text from the question in the shipping logic app part.
- Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

upvoted 2 times

Brak 2 years, 9 months ago

Exactly. Data Gateway does not provide "connectivity to support legacy apps." So it's not the answer. That leaves S2S.

upvoted 2 times

pieronegri 2 years, 8 months ago

you are right Brak, that would be used in conjunction with Integration Service Environment for Logic apps so to secure them with the on premise VNet

upvoted 2 times

dineshkm06tnj 2 years, 12 months ago

Agreed: On-premises Data Gateway

upvoted 2 times

sebainones 3 years ago

Agreed with "On-premises data Gateway".
But what piece of information should I use/consider to discard "S2S VPN" as valid answer?

upvoted 4 times

MrZoom 2 years, 7 months ago

You shouldn't discard it. See comments from Brak and pieronegri, above. For info on ISE, see https://docs.microsoft.com/en-us/azure/logic-apps/connect-virtual-network-vnet-isolated-environment

upvoted 2 times

Topic 24 - Testlet 25

Question #1 *Topic 24*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

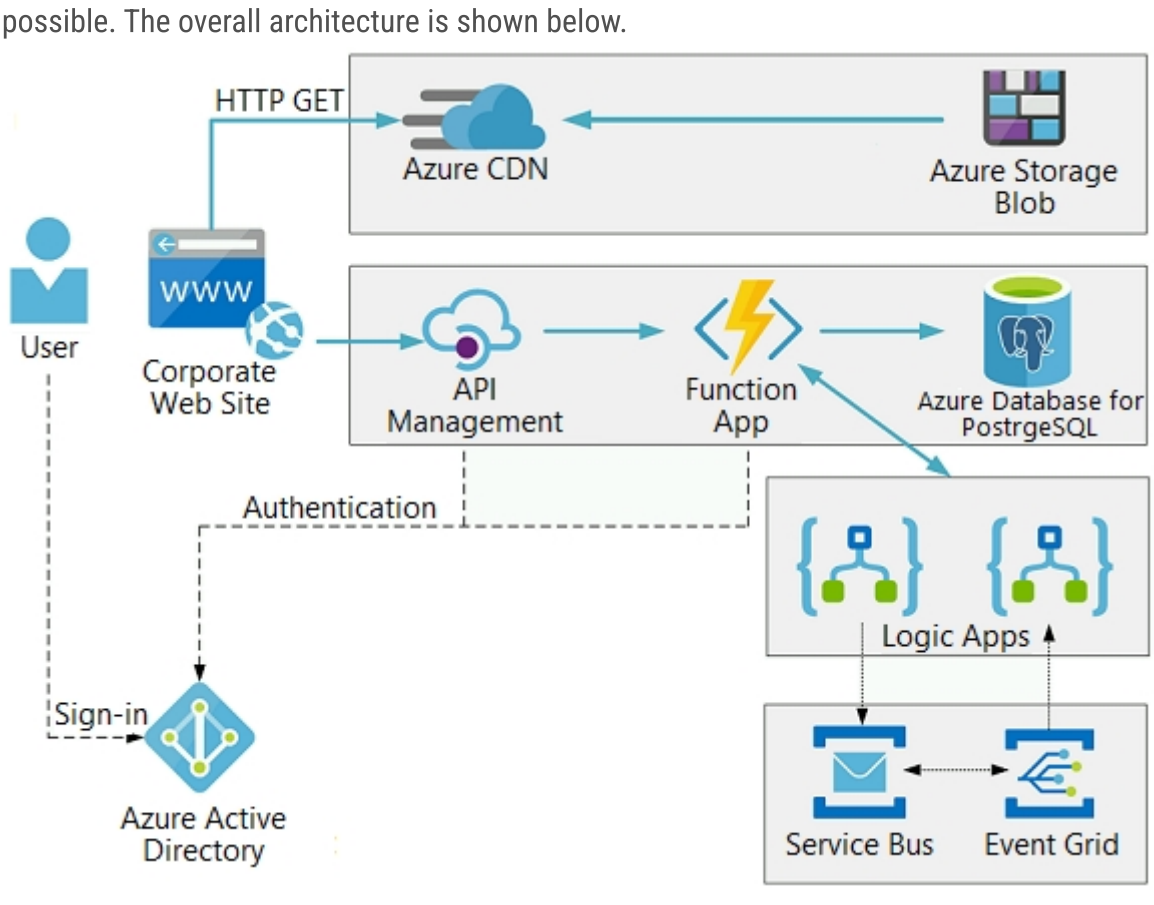

User authentication -

The following steps detail the user authentication process:
1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -

Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

HOTSPOT -

You need to configure the integration for Azure Service Bus and Azure Event Grid.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

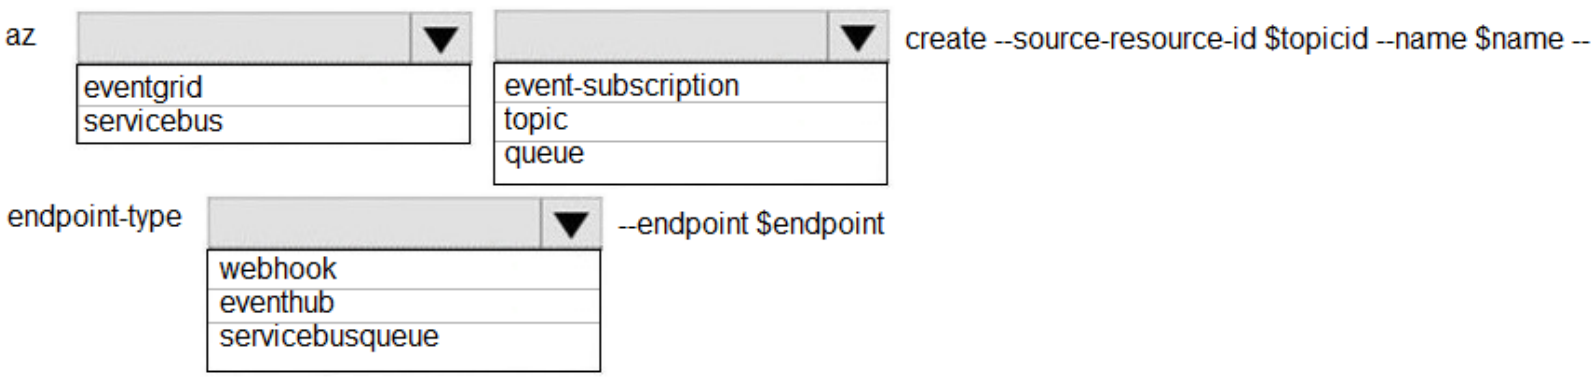

Correct Answer:

Answer Area

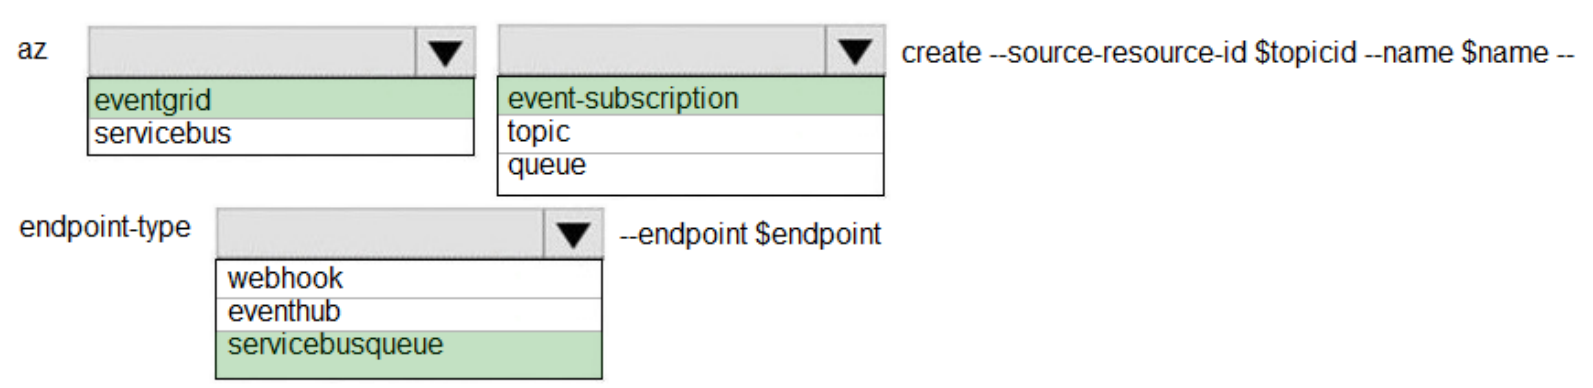

Box 1: eventgrid -

To create event subscription use: az eventgrid event-subscription create

Box 2: event-subscription -

Box 3: servicebusqueue -
Scenario: Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.
Reference:
https://docs.microsoft.com/en-us/cli/azure/eventgrid/event-subscription?view=azure-cli-latest#az_eventgrid_event_subscription_create

malay1232489 Highly Voted 2 years, 7 months ago
correcto...

https://docs.microsoft.com/en-us/azure/event-grid/handler-service-bus
upvoted 39 times

MrZoom 2 years, 7 months ago
Agreed. Note that the casus also denotes the opposite integration: "Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing". But none of the given options seem to do just this, so the given answer is correct.
upvoted 6 times

coffecold 1 year ago
Exact, this is only one piece of the integration. That for grid to service bus.
upvoted 1 times

surprise0011 6 months, 2 weeks ago
received 2023-04-17 went with given answer, score 926
upvoted 2 times

lugospod Highly Voted 1 year, 9 months ago
Got this one 01/2022. Went with most voted (to avoid writing answers again)
upvoted 8 times

Vmwarevirtual Most Recent 5 months ago
Appeared in exam I took on 27 May 2023
1 eventgrid
2 event-subscription
3 servicebusqueue
Following is the example from Azcli command of az eventgrid

Create a new event subscription for an Azure subscription, using default filters, and an Azure ServiceBusQueue as a destination.
az eventgrid event-subscription create --name es2 \
--source-resource-id /subscriptions/{SubID} \
--endpoint-type servicebusqueue \
--endpoint /subscriptions/{SubID}/resourceGroups/TestRG/providers/Microsoft.ServiceBus/namespaces/ns1/queues/queue1
upvoted 2 times

OPT_001122 11 months, 2 weeks ago
correct
upvoted 1 times

BishopeL 2 years, 2 months ago
Absolutely correct
upvoted 3 times

rustycables 2 years, 2 months ago
Given is valid... but DYOR.
https://docs.microsoft.com/en-us/cli/azure/eventgrid/event-subscription?view=azure-cli-latest#az_eventgrid_event_subscription_create
upvoted 2 times

MrXBasit 2 years, 3 months ago
Correcto
upvoted 2 times

Question #2 *Topic 24*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

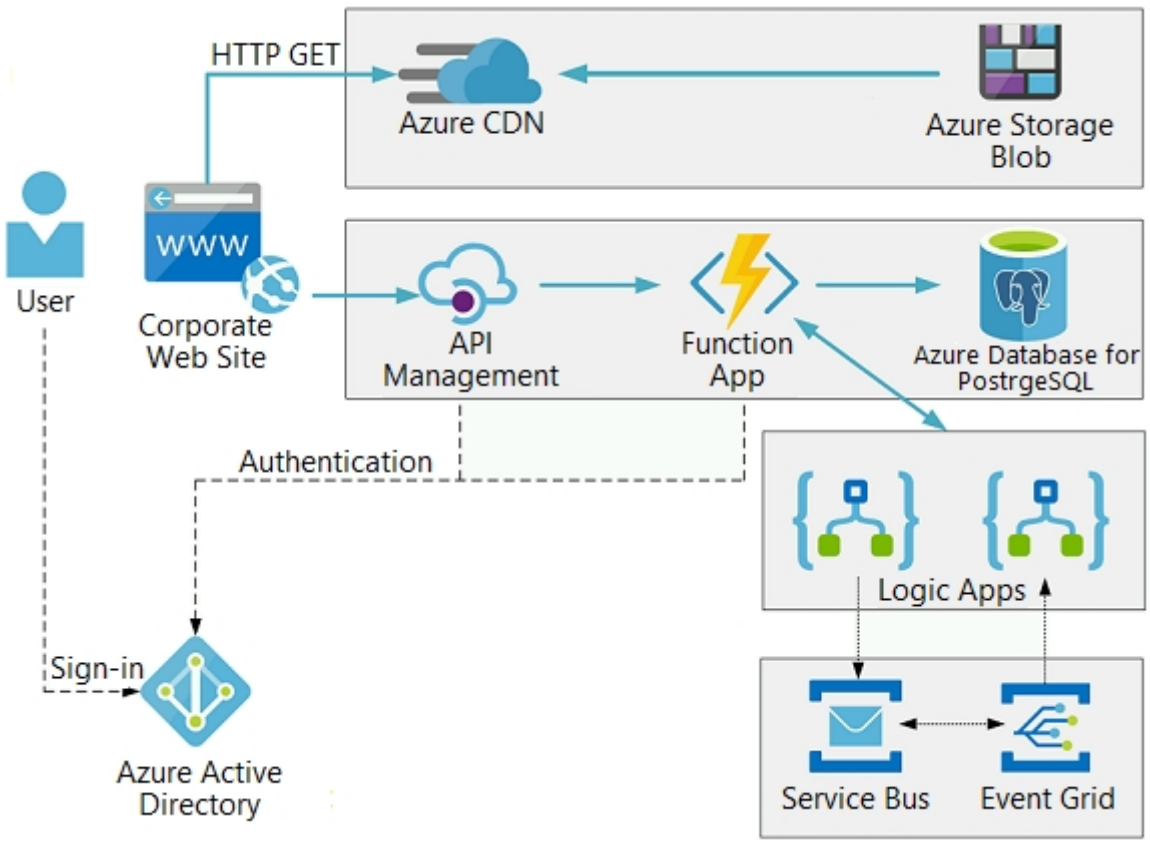

User authentication -

The following steps detail the user authentication process:

1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -
Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

You need to ensure that all messages from Azure Event Grid are processed.

What should you use?

A. Azure Event Grid topic

B. Azure Service Bus topic

C. Azure Service Bus queue

D. Azure Storage queue

E. Azure Logic App custom connector

Correct Answer: *C*

As a solution architect/developer, you should consider using Service Bus queues when:

☞ Your solution needs to receive messages without having to poll the queue. With Service Bus, you can achieve it by using a long-polling receive operation using the TCP-based protocols that Service Bus supports.

Reference:

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted

Community vote distribution

C (100%)

BrettusMaximus Highly Voted 2 years, 6 months ago

Answer is Correct C. Service Bus
Only C. ServiceBus Queue and B. Service Bus Topic give a guarantee of delivery.
Since Topics can only receive messages from a Queue that leaves only C.

upvoted 24 times

surprise0011 6 months, 2 weeks ago

received 2023-04-17 went with given answer, score 926

upvoted 2 times

agcertif Highly Voted 2 years, 8 months ago

Azure Event Grid must use Azure Service Bus for queue-based load leveling.

upvoted 10 times

ZodiaC 2 years, 3 months ago

YEP CORRECT IS C ! 100% >

I HATE THIS QUESTION,, Who create this one LOL
upvoted 6 times

Vmwarevirtual Most Recent 5 months ago
Appear in exam I took at 27-May-2023

I chose C - Azure Service Bus queue
https://learn.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept?source=recommendations&tabs=event-grid-event-schema#how-many-events-are-emitted-and-how-often
upvoted 1 times

szumik 8 months, 1 week ago
had it on 22/02/2023. Went with SB queue and topic, score 846
upvoted 1 times

OPT_001122 11 months, 2 weeks ago
Selected Answer: C
C. Azure Service Bus queue
upvoted 2 times

Eltooth 1 year, 4 months ago
Selected Answer: C
C is correct answer.
Service Bus Queue.
upvoted 2 times

Bogdan75 1 year, 7 months ago
I agree the provided answer works – but why not consider Azure Service Bus topics (+subscriptions) instead?
upvoted 1 times

kozchris 1 year, 8 months ago
Easy, C - Service Bus Queue.
From the problem description:
"Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering."
upvoted 4 times

lugospod 1 year, 9 months ago
Got this one 01/2022. Went with service bus queue since it war requested.
upvoted 4 times

phvogel 1 year, 12 months ago
I think Service Bus is the right answer but not because of the guarantee of delivery (Event Grid already promises "at least once" delivery). I'd pick it because of the dead letter queue that ensures that every message could be processed (or, at least, looked at).

In fact, without the dead letter queue, I'd pick Storage Queue because of its larger capacity and longer hold times (you could image the Logic App -- already using the Service Bus for load leveling -- falling so far behind that the number of messages hits the capacity limit on the Service Bus.

It's just a terrible, terrible question.
upvoted 3 times

yoloswag 2 years, 4 months ago
Horrible question, just why Microsoft, why...
upvoted 7 times

shoguns6 2 years, 4 months ago
Bit confusing, as per arch diagram, from azure event grid, arrows ends at LogicApp, meaning, logic app is subscriber to event grid... will it not be 'logic app connector'?
upvoted 2 times

Spooky7 2 years, 4 months ago
Question doesn't make sense. Event Grid events are the consequence of Service Bus messages (when new message appears is Service Bus than Event Grid emits an event to subscribers). So how Service Bus can be an answer which ensures that all EventGrid events are processed?
upvoted 1 times

pac1311 2 years, 8 months ago
Weird question, not sure about the answer probably service bus queue or storage queue
upvoted 1 times

danielcr 2 years, 8 months ago
Check the Case Study: "Azure Event Grid must use Azure Service Bus for queue-based load leveling." so.... C, Azure Server Bus queue (IMHO)
upvoted 5 times

cbn 2 years, 8 months ago
Question is bit confusing. Not sure on the answer.
upvoted 2 times

cbn 2 years, 8 months ago

Looking more into the question, the given answer looks correct.

Event grid sends all events to the service bus queue, and it is the best place to check if all events received by event grid are processed correctly.

upvoted 7 times

Topic 25 - Testlet 26

Question #1 *Topic 25*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements -

Policy service -

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Policies -

Log policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events -

Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

PolicyLib -

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Other -

Anomaly detection service -

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring -

All web applications and services have health monitoring at the /health service endpoint.

Issues -

Policy loss -
When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue -
When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency -
Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code -

EventGridController.cs -
Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

EventGridController.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "input1",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
EG55   }
EG56   private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57   {
EG58     . . .
EG59   }
EG60   private string GetLogData(string uri)
EG61   {
EG62     . . .
EG63   }
EG64   static string BlobStoreAccountSAS(string containerName)
EG65   {
EG66     . . .
EG67   }
EG68 }
```

LoginEvent.cs -

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

Question

DRAG DROP -

You need to add code at line EG15 in EventGridController.cs to ensure that the Log policy applies to all services.

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:

Code segments

topic

status

eventType

Succeeded

operationName

resourceProvider

Answer Area

```
if {
  @event[ "data" ][ "   code segment   " ].ToString() == "   code segment   "
  &&
  @event[ "data" ][ "   code segment   " ].ToString() == "Microsoft.Web/sites/write"
)
```

Correct Answer:

Code segments

topic

status

eventType

Succeeded

operationName

resourceProvider

Answer Area

```
if {
  @event[ "data" ][ "   status   " ].ToString() == "   Succeeded   "
  &&
  @event[ "data" ][ "   operationName   " ].ToString() == "Microsoft.Web/sites/write"
)
```

Scenario, Log policy: All Azure App Service Web Apps must write logs to Azure Blob storage.

Box 1: Status -

Box 2: Succeeded -

Box 3: operationName -

Microsoft.Web/sites/write is resource provider operation. It creates a new Web App or updates an existing one.

Reference:

https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations

MrZoom Highly Voted 2 years, 7 months ago

Answer seems to be correct.

Once a new web app is created, an event is triggered from the resource group...
https://docs.microsoft.com/en-us/azure/event-grid/event-schema-resource-groups?tabs=event-grid-event-schema

That event contains the status, which must be Succeeded in order for it to make sense to apply the policy. The above site also mentions that the event["data"] contains an operationName property for resource group events, which in this case is an "Microsoft.web/sites/write" operation.

upvoted 26 times

MiraA 2 years ago

operationName ... Name of the operation.
status ... String describing the status of the operation. Some common values are: Started, In Progress, Succeeded, Failed, Active, Resolved.
https://docs.microsoft.com/en-us/azure/azure-monitor/essentials/activity-log-schema#administrative-category

Microsoft.Web/sites/Write ... Create a new Web App or update an existing one
https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations

But why do we need to call EnsureLogging() upon creating/updating the Web App?
I cannot find such requirement in the assignment. Or am I wrong?

upvoted 4 times

OPT_001122 Most Recent 10 months, 3 weeks ago

1: Status

2: Succeeded

3: operationName

upvoted 2 times

ning 2 years, 2 months ago

The given answer seems logical, but the wording is so confusing ...
This operation is for when a new web site is created or updated, or a web log is generated ...

upvoted 1 times

MiraA 2 years ago

Confusing is the opening brace - it should be "(" instead of "{" as it is a logical expression for the "if" statement. ;-)

upvoted 3 times

Faizs 2 years, 5 months ago

Correct

upvoted 1 times

Question #2 *Topic 25*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements -

Policy service -

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Policies -

Log policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events -

Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

PolicyLib -

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Other -

Anomaly detection service -

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring -

All web applications and services have health monitoring at the /health service endpoint.

Issues -

Policy loss -
When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue -
When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency -
Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code -

EventGridController.cs -
Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

EventGridController.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "input1",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
EG55   }
EG56   private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57   {
EG58     . . .
EG59   }
EG60   private string GetLogData(string uri)
EG61   {
EG62     . . .
EG63   }
EG64   static string BlobStoreAccountSAS(string containerName)
EG65   {
EG66     . . .
EG67   }
EG68 }
```

LoginEvent.cs -

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

Question

HOTSPOT -

You need to insert code at line LE03 of LoginEvent.cs to ensure that all authentication events are processed correctly.

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:

Answer Area

public string [▾] (get; set;)
- id
- eventType
- dataVersion
- metadataVersion

public string [▾] (get; set;)
- id
- eventType
- dataVersion
- metadataVersion

public string [▾] (get; set;)
- id
- eventType
- dataVersion
- metadataVersion

Correct Answer:

Answer Area

public string [▾] (get; set;)
- id
- eventType
- dataVersion
- metadataVersion

public string [▾] (get; set;)
- id
- eventType
- dataVersion
- metadataVersion

public string [▾] (get; set;)
- id
- eventType
- dataVersion
- metadataVersion

Box 1: id -

id is a unique identifier for the event.

Box 2: eventType -
eventType is one of the registered event types for this event source.

Box 3: dataVersion -
dataVersion is the schema version of the data object. The publisher defines the schema version.
Scenario: Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.
The following example shows the properties that are used by all event publishers:

```
[
{
"topic": string,
"subject": string,
"id": string,
"eventType": string,
"eventTime": string,
"data":{
object-unique-to-each-publisher
},
"dataVersion": string,
"metadataVersion": string
}
]
```

Reference:
https://docs.microsoft.com/en-us/azure/event-grid/event-schema

somenkr Highly Voted 2 years, 3 months ago
Answer is correct : metadataVersion string Not required, but if included, must match the Event Grid Schema
upvoted 13 times

JH81 Most Recent 4 months, 1 week ago
Wow MS, id and eventType I get but we then have to decide between dataVersion and metadataVersion and somehow know that metadataVersion only has a getter??? Another wonderful question.
upvoted 3 times

plusJoyed 5 months, 4 weeks ago
Why does the order matter for the 3 values? Isn't this just a class model/interface?
upvoted 4 times

devex 1 month, 3 weeks ago
Order does not matter and as long as you do not select duplicates, you get at least 2/3 points. Strange question.
upvoted 1 times

UnknowMan 2 years, 5 months ago
Correct (Metadataversion have not a Setter so is Dataversion), Id and eventType is required in Event Grid Schema
upvoted 4 times

faizalzain 2 years, 6 months ago
the last one i think should be metadataversion
upvoted 2 times

jay158 2 years, 3 months ago
NO , the only allowed value at present for metadataversion is 1
https://docs.microsoft.com/en-us/azure/event-grid/event-schema
upvoted 2 times

ning 2 years, 2 months ago
No, metadataversion is determined by event grid, not sender
upvoted 2 times

Mike_St 2 years, 7 months ago
As shown in the example in the explanation of the answer, all 4 values are needed BUT from programmer stand point... ID should not be public and have getter and setter... it should be a private immutable value... so in my opionion it should be eventType dataVersion and metadataVersion (i am not 100% sertain of that answer)
upvoted 4 times

rdemontis 2 years, 7 months ago

I think answer is correct because id is a required field in EventGrid Schema and if you loog at the EventGridEvent class from Microsoft libraries the property has both getter and setter methods defined.
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.eventgrid.models.eventgridevent.id?view=azure-dotnet#Microsoft_Azure_EventGrid_Models_EventGridEvent_Id

upvoted 8 times

aperez1979 2 years, 7 months ago

I agree, metadaversion only get https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.eventgrid.models.eventgridevent.metadataversion?view=azure-dotnet

upvoted 10 times

clarionprogrammer 2 years, 6 months ago

Good point. I can't believe their question choice is so nit-picky.

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

This question proves how ridiculous this test is.

upvoted 21 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements -

Policy service -

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Policies -

Log policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events -

Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

PolicyLib -

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Other -

Anomaly detection service -

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring -

All web applications and services have health monitoring at the /health service endpoint.

Issues -

Policy loss -
When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue -
When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency -
Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code -

EventGridController.cs -
Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

EventGridController.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "input1",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
EG55   }
EG56   private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57   {
EG58     . . .
EG59   }
EG60   private string GetLogData(string uri)
EG61   {
EG62     . . .
EG63   }
EG64   static string BlobStoreAccountSAS(string containerName)
EG65   {
EG66     . . .
EG67   }
EG68 }
```

LoginEvent.cs -

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

Question

HOTSPOT -

You need to implement the Log policy.

How should you complete the EnsureLogging method in EventGridController.cs? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
var client = new WebSiteManagementClient(. . .);
var id = ParseResourceID(resource);
var appSettings = new StringDictionary(name: "properties",
  properties: new Dictionary<string, string> {
  {"DIAGNOSTICS_AZUREBLOBCONTAINERSASURL", BlobStoreAccountSAS(" ")},

  {"DIAGNOSTICS_AZUREBLOBRETENTIONINDAYS", " "}

});
client.WebApps. (
  id.resourceGroup,
  id.name, appSettings);
```

(" " options: logs, logdrop)

(" " options: 15, 30)

(options: UploadLoggingSettings, UpdateApplicationSetting)

Correct Answer:

Answer Area

```
var client = new WebSiteManagementClient(. . .);
var id = ParseResourceID(resource);
var appSettings = new StringDictionary(name: "properties",
  properties: new Dictionary<string, string> {
  {"DIAGNOSTICS_AZUREBLOBCONTAINERSASURL", BlobStoreAccountSAS(" ")},

  {"DIAGNOSTICS_AZUREBLOBRETENTIONINDAYS", " "}

});
client.WebApps. (
  id.resourceGroup,
  id.name, appSettings);
```

(" " options: logs, logdrop)

(" " options: 15, 30)

(options: UploadLoggingSettings, UpdateApplicationSetting)

Box 1: logdrop -

All log files should be saved to a container named logdrop.

Box 2: 15 -
Logs must remain in the container for 15 days.
Box 3: UpdateApplicationSettings
All Azure App Service Web Apps must write logs to Azure Blob storage.
Reference:
https://blog.hompus.nl/2017/05/29/adding-application-logging-blob-to-a-azure-web-app-service-using-powershell/

rdemontis Highly Voted 2 years, 7 months ago
answer is correct
upvoted 22 times

JH81 Most Recent 4 months, 1 week ago
First 2 are simple and straight from the requirements but why do you have to know a property name for a .NET class for the 3rd??
upvoted 2 times

OPT_001122 11 months, 2 weeks ago
Correct
upvoted 2 times

gmishra88 1 year ago
What is BlobStoreAccountSAS? I cannot find anything with that reference. Microsoft just randomly expects me to remember these class names. Google, sorry Bing, doesn't Bing it for them
upvoted 2 times

coffecold 1 year ago
BlobStoreAccountSAS() is just a method listed in the case on line EG64
upvoted 3 times

SivajiTheBoss 1 year, 7 months ago
Correct Answer provided
upvoted 2 times

AzureDJ 1 year, 7 months ago
Given answer is correct
upvoted 1 times

mc0re 2 years, 2 months ago
This is too simple. One doesn't need to know anything except for the requirements...
upvoted 3 times

xRiot007 1 year, 4 months ago
Right, maybe MS needs even more complicated tests, so we can memorize even more info that we can find using Google in 10 minutes or less.
upvoted 8 times

BroGood 2 years, 4 months ago
given answer is correct
upvoted 2 times

UnknowMan 2 years, 5 months ago
Correct ;)
upvoted 2 times

Topic 26 - Testlet 3

Question #1 *Topic 26*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements -

Policy service -

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Policies -

Log policy -

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events -

Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

PolicyLib -

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Other -

Anomaly detection service -

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring -

All web applications and services have health monitoring at the /health service endpoint.

Issues -

Policy loss -
When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue -
When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency -
Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code -

EventGridController.cs -
Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

EventGridController.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "inputl",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54       }
EG55     }
EG56     private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57     {
EG58       . . .
EG59     }
EG60     private string GetLogData(string uri)
EG61     {
EG62       . . .
EG63     }
EG64     static string BlobStoreAccountSAS(string containerName)
EG65     {
EG66       . . .
EG67     }
EG68 }
```

LoginEvent.cs -

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

Question

You need to resolve a notification latency issue.

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

NOTE: Each correct selection is worth one point.

A. Set Always On to true.

B. Ensure that the Azure Function is using an App Service plan.

C. Set Always On to false.

D. Ensure that the Azure Function is set to use a consumption plan.

Correct Answer: *AB*

Azure Functions can run on either a Consumption Plan or a dedicated App Service Plan. If you run in a dedicated mode, you need to turn on the Always On setting for your Function App to run properly. The Function runtime will go idle after a few minutes of inactivity, so only HTTP triggers will actually "wake up" your functions. This is similar to how WebJobs must have Always On enabled.

Scenario: Notification latency: Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

Anomaly detection service: You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service.

If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Reference:

https://github.com/Azure/Azure-Functions/wiki/Enable-Always-On-when-running-on-dedicated-App-Service-Plan

Community vote distribution

AB (100%)

andsol Highly Voted 2 years, 7 months ago

Correct

upvoted 29 times

mlantonis Highly Voted 2 years, 5 months ago

Correct Answer: A and B

Always On enables waking up on HTTP trigger, but does not prevent the exceeding the max time out time of 230 seconds.
If your function app is on the Consumption plan, there can be up to a 10-minute delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled. You can also use an Event Grid trigger with your Blob storage account.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/dedicated-plan#always-on

https://github.com/Azure/Azure-Functions/wiki/Enable-Always-On-when-running-on-dedicated-App-Service-Plan

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview

https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp#event-grid-trigger

upvoted 17 times

uffuchsi Most Recent 8 months, 1 week ago

Selected Answer: AB

Always On is required and you need App Servie Plan for that feature

upvoted 2 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: AB

Correct

upvoted 1 times

DP_Bhatt 11 months, 3 weeks ago

Selected Answer: AB

Correct Answer: A and B

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

A and B is correct

upvoted 1 times

Naval708 2 years, 5 months ago

A and B is correct

upvoted 2 times

UnknowMan 2 years, 5 months ago

Correct ;)

upvoted 3 times

Topic 27 - Testlet 4

Question #1 *Topic 27*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

Messages are sent to ContentUploadService.

Content is processed by ContentAnalysisService.

After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -
You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
apiVersion: '2018-10-01'
type: Microsoft.ContainerInstance/containerGroups
location: westus
name: contentUploadService
properties:
    containers:
    - name: service
        properties:
        image: contoso/contentUploadService:latest
        ports:
        -  port: 80
           protocol: TCP
        resources:
          requests:
             cpu: 1.0
             memoryInGB: 1.5

ipAddress:
    ip: 10.23.121.112
    ports:
      - port: 80
        protocol: TCP

networkProfile:

id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14        "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

HOTSPOT -

You need to ensure that validation testing is triggered per the requirements.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
var event = getEvent();
if (event.eventType === '                    '
```

ImagePushed
RepositoryItem
ImageDeployed
RepositoryUpdated

```
 && event.data.target.          === 'contentanalysisservice'
```

aci
image
service
repository

```
 && event.                    .contains('contosoimages'))
```

topic
service
repository
imageCollection

```
{
  startValidationTesting();
}
```

Correct Answer:

Answer Area

```
var event = getEvent();
if (event.eventType === '                    ')
```

- ImagePushed
- RepositoryItem
- ImageDeployed
- RepositoryUpdated

```
 && event.data.target.          === 'contentanalysisservice'
```

- aci
- image
- service
- repository

```
 && event.                    .contains('contosoimages'))
```

- topic
- service
- repository
- imageCollection

```
{
  startValidationTesting();
}
```

Box 1: RepositoryUpdated -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Box 2: service -

Box 3: imageCollection -
Reference:
https://docs.microsoft.com/en-us/azure/devops/notifications/oob-supported-event-types

halmosi Highly Voted 2 years, 7 months ago

Incorrect...
1. ImagePushed
2. repository
3. topic

https://docs.microsoft.com/hu-hu/azure/event-grid/event-schema-container-registry?tabs=event-grid-event-schema
upvoted 104 times

Shion2009 2 years, 7 months ago

I'm not sure if this is correct. Take a look at Background Overview:
"The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages."

"contentanalysisservice" is not the name of a repository and "contonsoimages" is the name of the ACR.
upvoted 2 times

Shion2009 2 years, 7 months ago

Okay, halmosi's answer seems possible:
https://docs.microsoft.com/de-de/azure/event-grid/event-schema-container-registry?tabs=event-grid-event-schema
upvoted 3 times

Pozz4ever 2 years, 6 months ago

Sorry I Don't understand... as you said contentanalysisservice in the new of the service, I suppose, and contosoimages is the name of the ACR. so how can it be repository and topic?
upvoted 1 times

kapetan 2 years, 6 months ago

An event does not have service, repository or imageCollection as top-level data, it has a topic. Regarding the repository, from the choises available, the target has only repository as property. It is all documented in the already given link:
https://docs.microsoft.com/en-us/azure/event-grid/event-schema-container-registry?tabs=event-grid-event-schema
upvoted 7 times

Pozz4ever 2 years, 6 months ago

thanks
upvoted 1 times

jvyas 2 years, 5 months ago

"contentanalysisservice" is the name of the repository/service. check the line AM08 in application manifest file.

upvoted 3 times

KingChuang 10 months, 1 week ago

On my exam 2022-12-26.
Chose：
1. ImagePushed
2. repository
3. topic

upvoted 2 times

clarionprogrammer 2 years, 6 months ago

In English...
https://docs.microsoft.com/en-us/azure/event-grid/event-schema-container-registry?tabs=event-grid-event-schema

upvoted 2 times

rdemontis 2 years, 7 months ago

thanks for the tip! The answer indicated by the test seems to be completely "off topic" ;-)

upvoted 6 times

coffecold Highly Voted 1 year ago

Ok, I need to start memorizing all schemes of possible event system topics:
Azure API Management
Azure App Configuration
Azure App Service
Azure Blob Storage
Azure Cache for Redis
Azure Communication Services
Azure Container Registry
Azure Event Hubs
Azure FarmBeats
Azure Health Data Services
Azure IoT Hub
Azure Key Vault
Azure Kubernetes Service
Azure Machine Learning
Azure Maps
Azure Media Services
Azure Policy
Azure resource groups
Azure Service Bus
Azure SignalR
Azure subscriptions....

upvoted 11 times

ENGs Most Recent 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

katrang 1 month, 1 week ago

I would go for
1. ImagePushed
2. Repository
3. Topic

Based on the grid event schema example everyone has linked to

[{
"topic": "/subscriptions/<subscription-id>/resourceGroups/<resource-group-name>/providers/Microsoft.ContainerRegistry/registries/<name>",
"eventType": "Microsoft.ContainerRegistry.ImagePushed",
"data": {
"target": {
"repository": "<name>",
},
},
}]

upvoted 1 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 2 times

macobuzi 1 month, 3 weeks ago

What did you choose then?

upvoted 1 times

KGi 8 months, 2 weeks ago

The number of questions with incorrect responses is overwhelming. ExamTopics should mark the correct answers properly to make life easier for us. :(

upvoted 4 times

AlexeyG 8 months, 2 weeks ago

Got this in 16/02/2023

upvoted 4 times

OPT_001122 10 months, 3 weeks ago

ImagePushed
Repository
Topic

upvoted 2 times

gmishra88 1 year ago

How does a new image being pushed mean a new deployment. And they expect me to remember the whole schema for all the services that can send events to EventGrid? And I could even imagine a blob storage being there. This requirement itself is at the edge of darkness and using a repository event to fulfill that requirement can only be thought of really junior hello world programmers of Microsoft.

upvoted 1 times

gmishra88 1 year ago

EUREKA. Did you guys notice "All Internal services must only be accessible from internal Virtual Networks (VNets)." It NEVER said it should be connected to a vnet. It just said it should be ACCESSED from INTERNAL vnet (not Azure vnet). A nice red-herring planted by Microsoft guys (Do not be evil)

THE ANSWER SHOULD BE CONSUMPTION PLAN

upvoted 1 times

TMakki 1 year, 2 months ago

This question is also in Udemy.
1. ImagePushed
2. repository
3. topic

upvoted 2 times

Ciupaz 1 week, 6 days ago

Which course?

upvoted 1 times

Sandeep12093 1 year, 4 months ago

1. Imagepushed - acr has only 4 event types imagepushed, imagedeleted, chartpushed n chart deleted
2. Service ?
3. Repository - this is correct as per MS documentation

upvoted 1 times

Sandeep12093 1 year, 4 months ago

I just checked ACR event schema
3. Topic - since it has id,eventtype,topic,subject,data,dataversion.... there is no repository

upvoted 1 times

Sandeep12093 1 year, 4 months ago

2. REPOSITORY

https://docs.microsoft.com/en-us/azure/container-registry/container-registry-event-grid-quickstart
Check image it has schema written.

upvoted 1 times

SivajiTheBoss 1 year, 7 months ago

100 % correct answer: Cross referenced in Udemy
ImagePushed
Repository
Topic

upvoted 3 times

petitbilly 1 year, 7 months ago

Got this one 03/2022

upvoted 1 times

leonidn 1 year, 9 months ago

ImagePushed
Repository
Topic

https://docs.microsoft.com/en-us/azure/event-grid/event-schema-container-registry?tabs=event-grid-event-schema

upvoted 2 times

mc0re 2 years, 2 months ago

"Image pushed into a repository" is not the same as "image deployed, container is running". One cannot start testing on a pushed image, only after it has been deployed.

upvoted 3 times

MiraA 2 years ago

Hmm, maybe some multi-step task will be used so "ImagePushed" event means the image will be run as a container in a moment? Or is there some functionality in that private ACR which will run/deploy the updated image automatically?

https://docs.microsoft.com/en-us/azure/container-registry/container-registry-tasks-multi-step

upvoted 1 times

Harichakradhar 2 years, 3 months ago

100 % correct
ImagePushed
Repository
Topic

upvoted 4 times

sien 2 years, 6 months ago

Indeed,
ImagePushed
Repository and
Topic

upvoted 7 times

Question #2 *Topic 27*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

Messages are sent to ContentUploadService.

Content is processed by ContentAnalysisService.

After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -

You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02       "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03       "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05       "createdDateTime" : "2019-12-24T06:01:44Z",
AM06       "logoUrl" : null,
AM07       "logoutUrl" : null,
AM08       "name" : "ContentAnalysisService",
AM09
AM10
AM11       "orgRestrictions" : [],
AM12       "parentalControlSettings" : {
AM13        "countriesBlockedForMinors" : [],
AM14         "legalAgeGroupRule" : "Allow"
AM15       },
AM16       "passwordCredentials" : []
AM17 }
```

Question

You need to deploy the CheckUserContent Azure Function. The solution must meet the security and cost requirements.

Which hosting model should you use?

A. Premium plan

B. App Service plan

C. Consumption plan

Correct Answer: *B*

Scenario:

You must minimize costs for all Azure services.

All Internal services must only be accessible from internal Virtual Networks (VNets).

Best for long-running scenarios where Durable Functions can't be used. Consider an App Service plan in the following situations:

☞ You have existing, underutilized VMs that are already running other App Service instances.

☞ You want to provide a custom image on which to run your functions.

☞ Predictive scaling and costs are required.

Note: When you create a function app in Azure, you must choose a hosting plan for your app. There are three basic hosting plans available for Azure Functions:

Consumption plan, Premium plan, and Dedicated (App Service) plan.

Incorrect Answers:

A: A Premium plan would be more costly.

C: Need the VNET functionality.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale

Community vote distribution

B (100%)

ning Highly Voted 2 years, 2 months ago

A is correct! I guess the guy gives the answer confused with which one is cheaper. Premium is cheaper than app service place. It can auto scale, warm starting and with VNET integration. In addition, it can run custom container as well, just the same as app service plan.

upvoted 24 times

xavi1 1 year, 8 months ago

Premium does NOT support VNET in multiple regions: https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#networking-features, i beleive it's B

upvoted 8 times

wangga 1 year, 7 months ago

I agree, VNET and multiple regions only available for AppService plan

upvoted 2 times

troy89 1 year, 8 months ago

Are you sure? When I compare the costs (same ACU, memory, etc.) then App Service Plan is always cheaper.

upvoted 3 times

ning 2 years, 2 months ago

Too add some more stuff here, ASE(application service environment) is very different app service plan and way more expensive, in addition, those two docker images are in azure registry so, definitely can be supported by premium plan.

upvoted 1 times

altafpatel1984 1 year, 10 months ago

but here will use existing app service plan, will not create new one. so using existing app service plan will be cost effective.

upvoted 3 times

j888 Highly Voted 2 years, 2 months ago

I believed the answer is correct.
https://docs.microsoft.com/en-us/azure/azure-functions/dedicated-plan
"You pay for function apps in an App Service Plan as you would for other App Service resources. This differs from Azure Functions Consumption plan or Premium plan hosting..."
So.. ASE can be made available on the App service plan that hosting the function

upvoted 8 times

ENGs Most Recent 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

MikeAWS 1 week, 2 days ago

sorry, what do you mean by "before the update of the Exam"?

upvoted 1 times

NightshadeRC 3 months, 1 week ago

Had this question today: 2023-07-26

upvoted 2 times

appieazure 9 months, 2 weeks ago

Selected Answer: B

On My Exam:2022-12-02

upvoted 1 times

KingChuang 10 months, 1 week ago

Selected Answer: B

On My Exam:2022-12-26
My Choice:
B:App Service plan

upvoted 3 times

OPT_001122 10 months, 3 weeks ago

Selected Answer: B

https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#networking-features

Premium plan Vnet integration is regional hence for this question it is not the correct answer.

Correct Anser - App Service Plan

upvoted 4 times

gmishra88 1 year ago

Another highly debated question. Very good points on whether Premium is costly or (dedicated) app service plan is costly. No idea. Only thing sure is it cannot be consumption because of the hidden wording for vnet integration. But I guess Microsoft guys forget that they added that vnet integration requirement and expected us to answer Consumption plan. But my intelligence will make me answer app service plan and hope for the best

upvoted 2 times

naicud 1 year, 7 months ago

I tink B is correct
The virtual network integration feature is used in Azure App Service dedicated compute pricing tiers. If your app is in an App Service Environment, it's already in a virtual network and doesn't require use of the VNet integration feature to reach resources in the same virtual network:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-networking-options#:~:text=The%20virtual%20network%20integration%20feature,in%20the%20same%20virtual%20network.

upvoted 2 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer: A
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale

upvoted 1 times

Mev4953 1 year, 9 months ago

The given answer and explanation seem correct.

upvoted 1 times

altafpatel1984 1 year, 11 months ago

"If you just want to run your function app in a virtual network, you can do this using the Premium plan."

Reference: https://docs.microsoft.com/en-us/azure/azure-functions/dedicated-plan

upvoted 4 times

phvogel 2 years ago

I'm liking the App Service plan (B). It's 10 cents an hour cheaper than Premium at the S1 level (even cheaper at higher levels) and has the VNet integration the case study requires. The only thing that Premium adds is private endpoints which aren't required here.
https://azure.microsoft.com/en-us/pricing/details/app-service/windows/

upvoted 4 times

Kalaisuran 2 years, 4 months ago

Option: A.
Reason : As per the security rule : All Internal services must only be accessible from internal Virtual Networks (VNets). Since the AZ function consumption plan does n't support vnet, So choose the next option Premium Plan

upvoted 3 times

ZodiaC 2 years, 3 months ago

ANY LINKS?

upvoted 1 times

SWedig 2 years, 3 months ago

https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale

upvoted 1 times

JamesDC 2 years, 3 months ago

https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale
with this document, ASE provides vNet integration. So, answer should be ASE for sure.

upvoted 2 times

Jurgen1234 2 years, 4 months ago

Should'be A Imho, both the UploadService and AnalysisService run in container instances on in app service plans

upvoted 1 times

ZodiaC 2 years, 3 months ago

ANY LINKS ???

upvoted 1 times

SWedig 2 years, 3 months ago

https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale

upvoted 2 times

ning 2 years, 1 month ago

Completely wrong! Container instance is different from app service plan, two very different ways to host container images, there is no way you can have a container instance running on top of a service plan

upvoted 1 times

PhilLI 1 year, 9 months ago

after choosing Docker Image you can still select an App Service Plan ....

upvoted 1 times

Topic 28 - Testlet 5

Question #1 *Topic 28*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

LabelMaker app -

Coho Winery produces, bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends data to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the LabelMaker application to include artwork and text for custom label designs.

Requirements. Data -

You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using SQL.

Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Requirements. Security -

You have the following security requirements:

Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.

▪

External partners must use their own credentials and authenticate with their organization's identity management solution.

External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.

Storage of e-commerce application settings must be maintained in Azure Key Vault.

E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).

Conditional access policies must be applied at the application level to protect company content.

The LabelMaker application must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

Requirements. LabelMaker app -

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure

Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture -

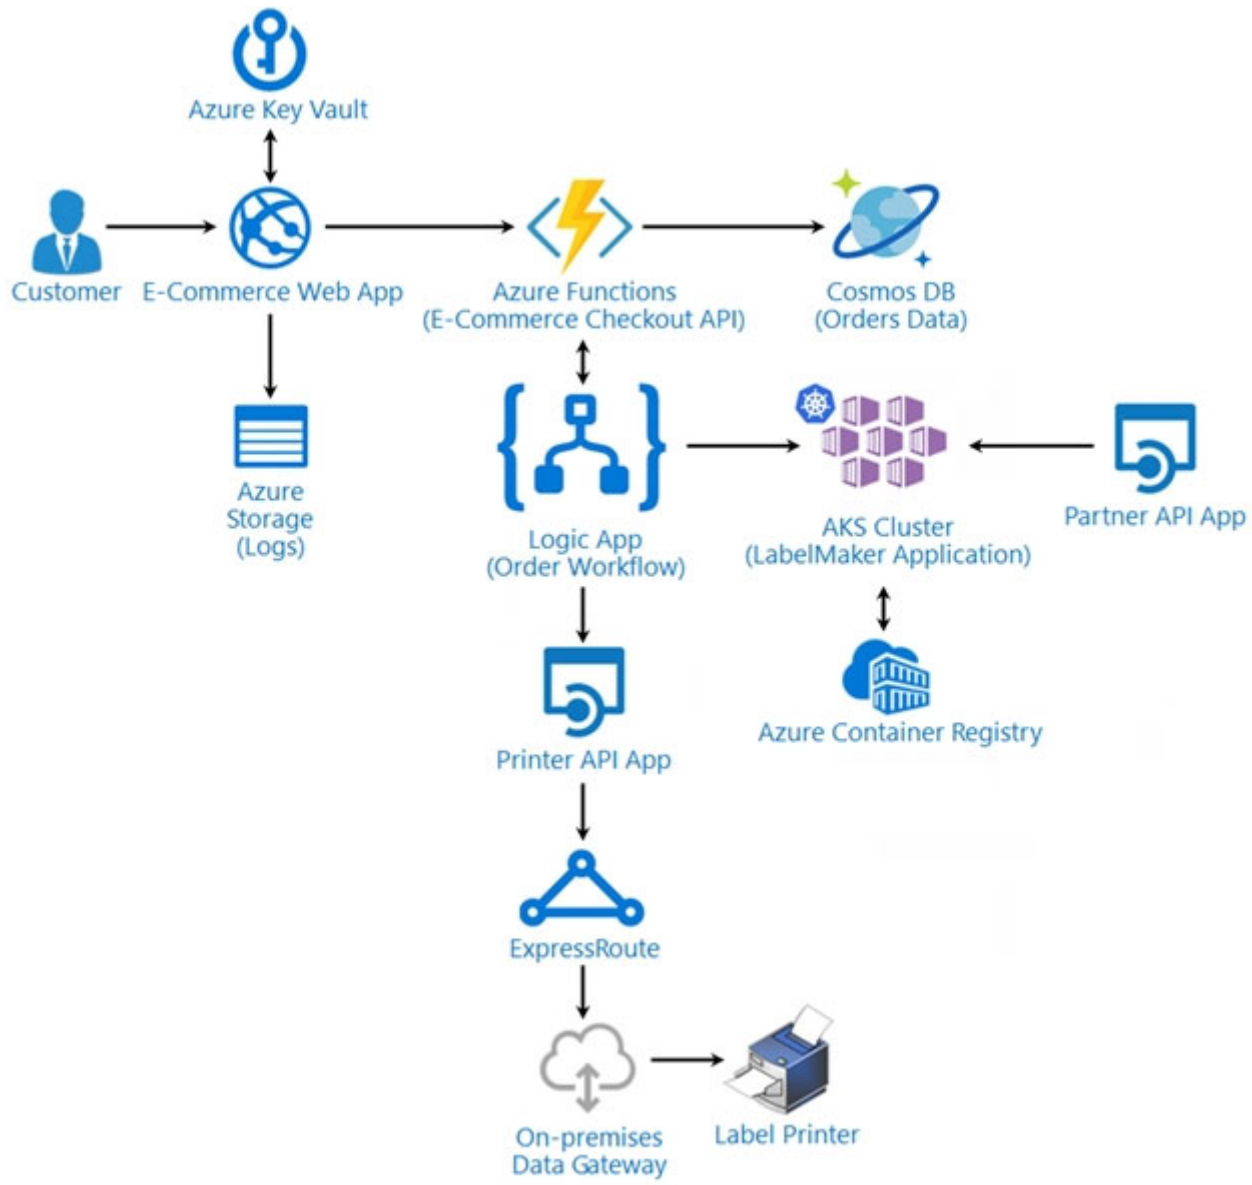

Issues -

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order.json -

Relevant portions of the app files are shown below. Line numbers are included for reference only.

This JSON file contains a representation of the data for an order that includes a single item.

Order.json -

Order.json

```
{
  "id" : 1,
  "customers" : [
  {
     "familyName" : "Doe",
     "givenName" : "John",
     "customerid" : 5
   }
  ],
  "line_items" : [
  {
    "fulfillable_quantity" : 1,
    "id": 6,
    "price" : "199.99" ,
    "product_id" : 7513594,
    "quantity": 1,
    "requires_shipping" : true ,
    "sku": "SFC-342-N" ,
    "title" : "Surface Go" ,
    "vendor" : "Microsoft" ,
    "name" : "Surface Go - 8GB" ,
    "taxable" : true ,
    "tax_lines" : [
    {
      "title" : "State Tax" ,
      "price" : "3.98" ,
      "rate" : 0.06
     }
    ],
    "total_discount" : "5.00" ,
    "discount_allocations" : [
     {
       "amount" : "5.00" ,
       "discount_application_index" : 2
       }
    ]
```

Question

DRAG DROP -

You need to deploy a new version of the LabelMaker application to ACR.

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:

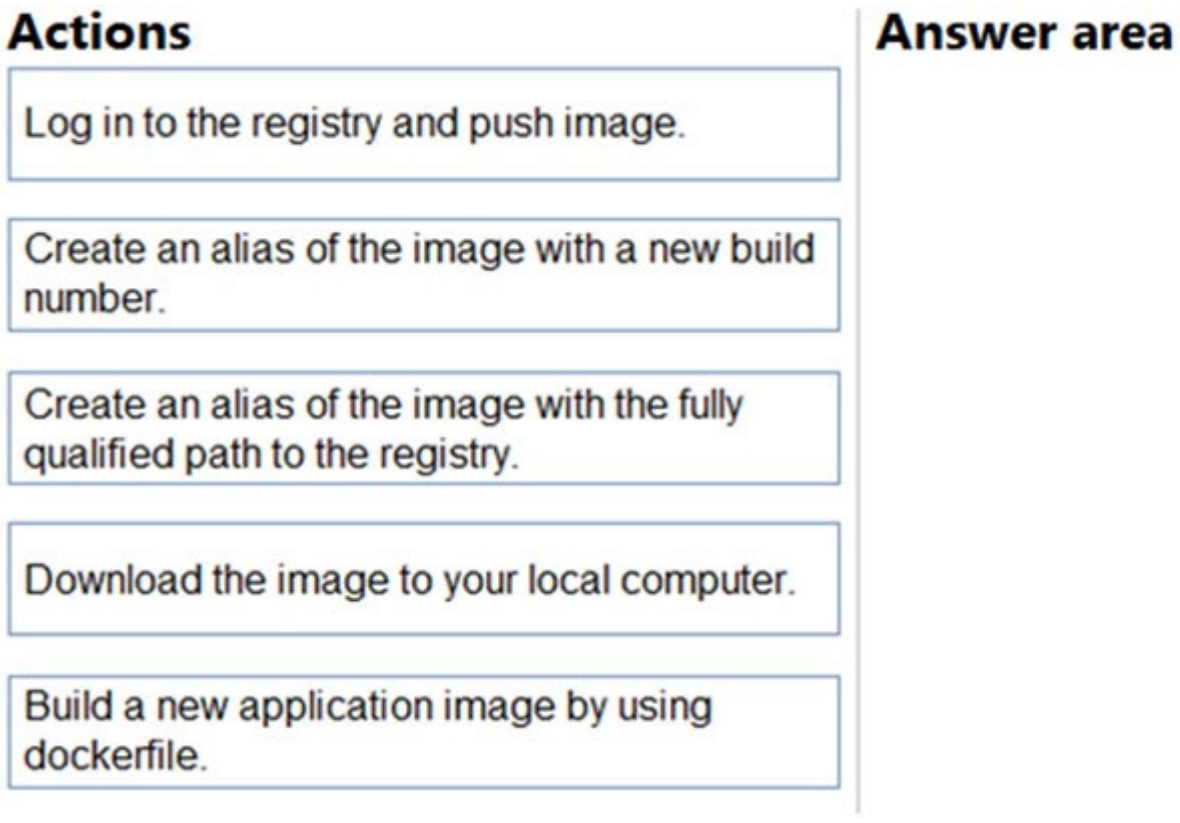

Correct Answer:

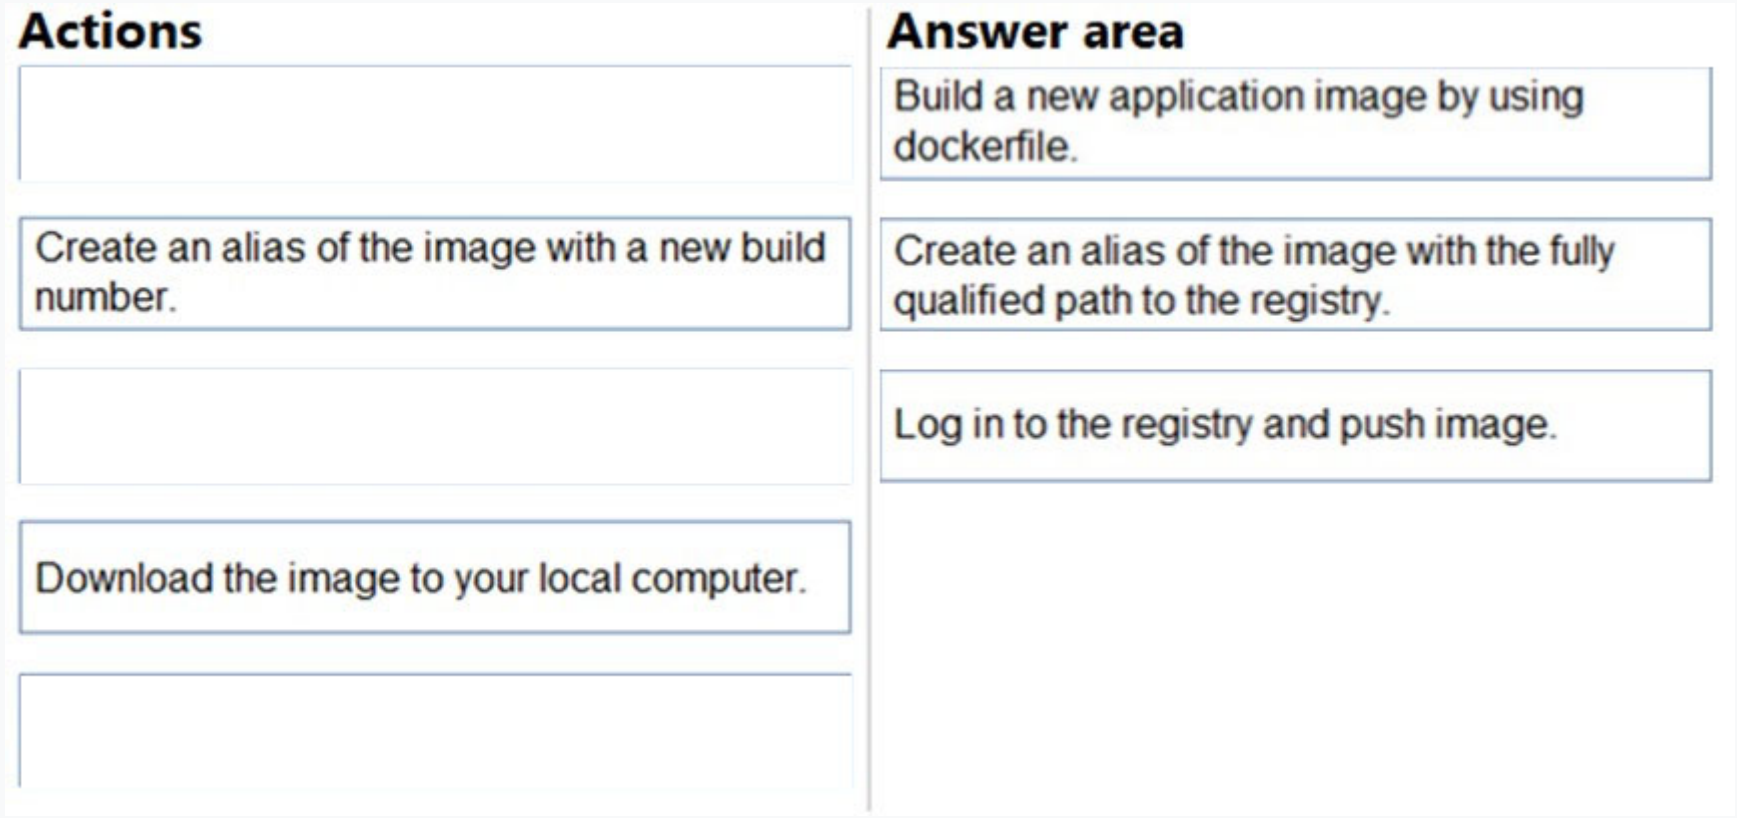

Step 1: Build a new application image by using dockerfile

Step 2: Create an alias if the image with the fully qualified path to the registry

Before you can push the image to a private registry, you've to ensure a proper image name. This can be achieved using the docker tag command. For demonstration purpose, we'll use Docker's hello world image, rename it and push it to ACR.

pulls hello-world from the public docker hub

$ docker pull hello-world

tag the image in order to be able to push it to a private registry

$ docker tag hello-word <REGISTRY_NAME>/hello-world

push the image

$ docker push <REGISTRY_NAME>/hello-world

Step 3: Log in to the registry and push image

In order to push images to the newly created ACR instance, you need to login to ACR form the Docker CLI. Once logged in, you can push any existing docker image to your ACR instance.

Scenario:

Coho Winery plans to move the application to Azure and continue to support label creation.

LabelMaker app -

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure

Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Reference:

https://thorsten-hans.com/how-to-use-a-private-azure-container-registry-with-kubernetes-9b86e67b93b6 https://docs.microsoft.com/en-us/azure/container-registry/container-registry-tutorial-quick-task

SivajiTheBoss Highly Voted 1 year, 7 months ago

Correct answer given.
Step 1: Build a new application image by using dockerfile

FROM node:8.9.3-alpine
RUN mkdir -p /usr/src/app
COPY ./app/ /usr/src/app/
WORKDIR /usr/src/app
RUN npm install
CMD node /usr/src/app/index.js

docker build ./aci-helloworld -t aci-test-app
docker images
docker run -d -p 8080:80 aci-test-app
(Example to build image only and run locally)

Step 2: Create an alias if the image with the fully qualified path to the registry
docker tag mcr.microsoft.com/oss/nginx/nginx:1.15.5-alpine myregistry.azurecr.io/samples/nginx

Step 3: Log in to the registry and push image
docker push myregistry.azurecr.io/samples/nginx

upvoted 16 times

OPT_001122 11 months, 2 weeks ago

nice explanation

upvoted 1 times

jay158 Highly Voted 2 years, 3 months ago

Correct
https://docs.microsoft.com/en-us/azure/container-registry/container-registry-get-started-docker-cli?tabs=azure-cli

upvoted 10 times

AnKiLa 2 years, 3 months ago

Agree
But it's an update - so this link looks better for me:
https://docs.microsoft.com/en-us/azure/container-registry/container-registry-tutorial-deploy-update

upvoted 3 times

whiteblack Most Recent 10 months ago

Very Useful.

upvoted 1 times

gmishra88 1 year ago

You could still alias with a new build number, but yeah, not required. And, yes, the version can be a build number. But if only three can be chosen I will not select that

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022, I chose as below:
Step 1: Build a new application image by using dockerfile
Step 2: Create an alias if the image with the fully qualified path to the registry
Step 3: Log in to the registry and push image

upvoted 3 times

mc0re 2 years, 2 months ago

Strange question. I realize that from the provided options this is the only solution one can assemble. But the project describes LabelMaker app to be a set of images running in AKS. Pushing a new image into ACR won't change a thing, as one needs to update the YAML deployment file for Kubernetes to pick up the image, and then tell AKS to reconcile.

upvoted 4 times

gmishra88 1 year ago

It is a Microsoft question. Keep intelligence out of it

upvoted 4 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

LabelMaker app -

Coho Winery produces, bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends data to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the LabelMaker application to include artwork and text for custom label designs.

Requirements. Data -

You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using SQL.

Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Requirements. Security -

You have the following security requirements:

Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.

▪

External partners must use their own credentials and authenticate with their organization's identity management solution.

External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.

Storage of e-commerce application settings must be maintained in Azure Key Vault.

E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).

Conditional access policies must be applied at the application level to protect company content.

The LabelMaker application must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

Requirements. LabelMaker app -

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure

Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture -

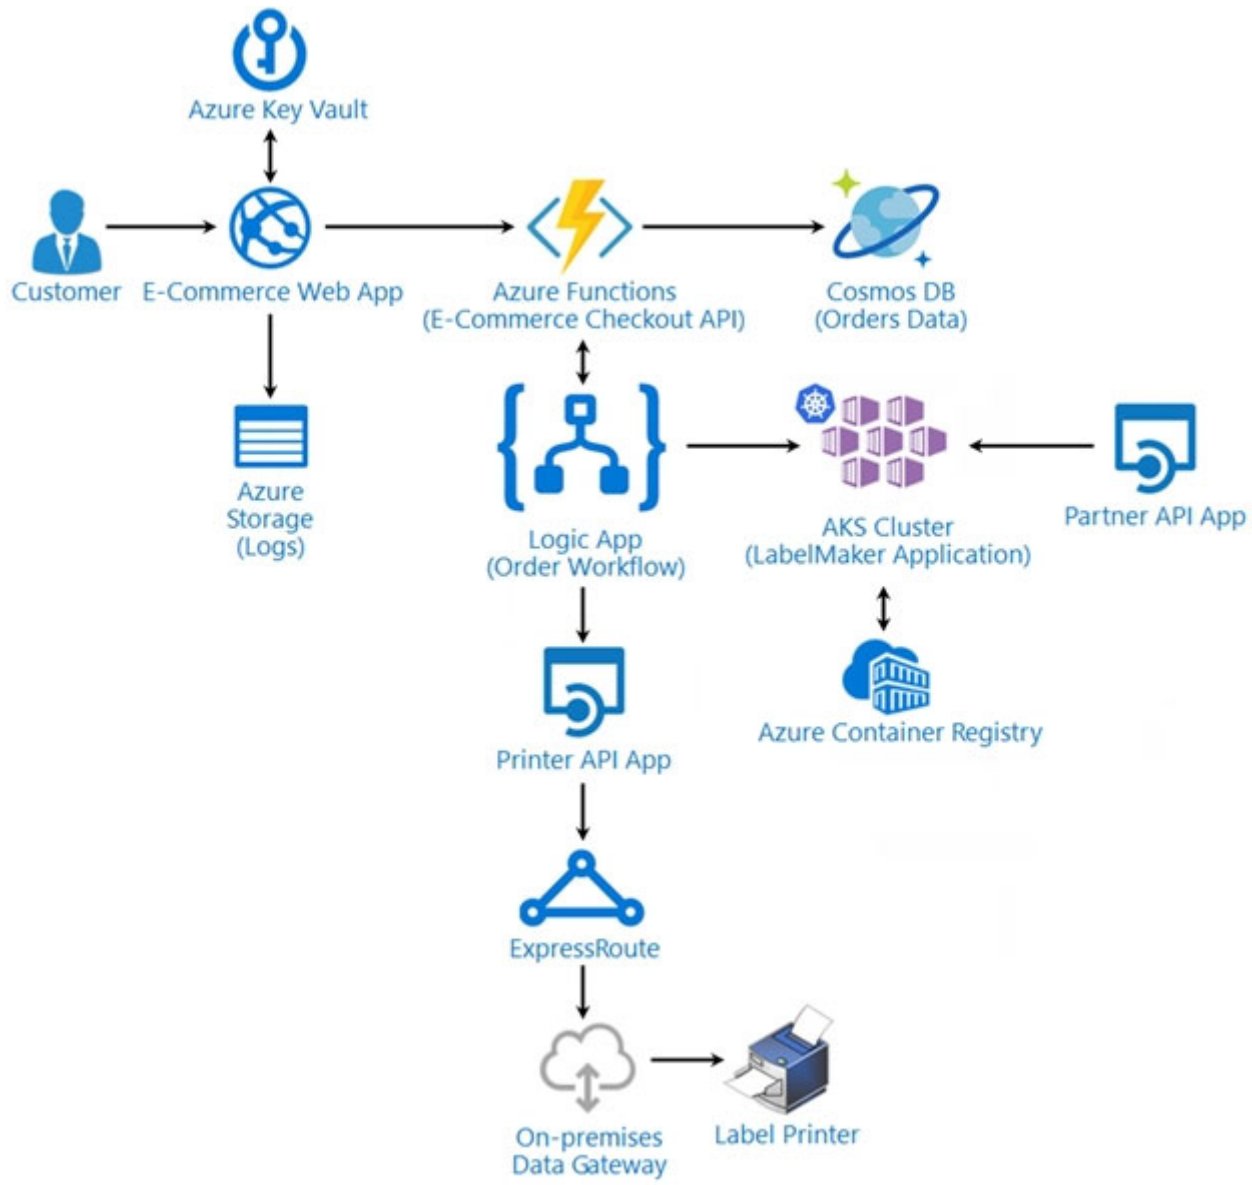

Issues -

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order.json -

Relevant portions of the app files are shown below. Line numbers are included for reference only.

This JSON file contains a representation of the data for an order that includes a single item.

Order.json -

Order.json

```
{
  "id" : 1,
  "customers" : [
  {
     "familyName" : "Doe",
     "givenName" : "John",
     "customerid" : 5
   }
  ],
  "line_items" : [
  {
    "fulfillable_quantity" : 1,
    "id": 6,
    "price" : "199.99" ,
    "product_id" : 7513594,
    "quantity": 1,
    "requires_shipping" : true ,
    "sku": "SFC-342-N" ,
    "title" : "Surface Go" ,
    "vendor" : "Microsoft" ,
    "name" : "Surface Go - 8GB" ,
    "taxable" : true ,
    "tax_lines" : [
    {
      "title" : "State Tax" ,
      "price" : "3.98" ,
      "rate" : 0.06
     }
    ],
    "total_discount" : "5.00" ,
    "discount_allocations" : [
     {
       "amount" : "5.00" ,
       "discount_application_index" : 2
       }
    ]
```

Question

You need to access data from the user claim object in the e-commerce web app.

What should you do first?

A. Write custom code to make a Microsoft Graph API call from the e-commerce web app.

B. Assign the Contributor RBAC role to the e-commerce web app by using the Resource Manager create role assignment API.

C. Update the e-commerce web app to read the HTTP request header values.

D. Using the Azure CLI, enable Cross-origin resource sharing (CORS) from the e-commerce checkout API to the e-commerce web app.

Correct Answer: *C*

Methods to Get User Identity and Claims in a .NET Azure Functions App include:

☞ ClaimsPrincipal from the Request Context

The ClaimsPrincipal object is also available as part of the request context and can be extracted from the HttpRequest.HttpContext.

☞ User Claims from the Request Headers.

App Service passes user claims to the app by using special request headers.

Reference:

https://levelup.gitconnected.com/four-alternative-methods-to-get-user-identity-and-claims-in-a-net-azure-functions-app-df98c40424bb

Community vote distribution

C (100%)

j888 Highly Voted 2 years, 2 months ago

https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-user-identities
The answer C

upvoted 24 times

Jonas_86 2 years, 1 month ago

your link is intersting but I think the answer you provide is FALSE.
As I understood from the link, if you want to access user claims from Azure Function then the C answer is OK but if you want to get it from webapp then the A answer is the good one.
Please, recheck the link above

upvoted 5 times

phvogel 2 years ago

To quote the link provided:
"For all language frameworks, App Service makes the claims in the incoming token (whether from an authenticated end user or a client application) available to your code by injecting them into the request headers. External requests aren't allowed to set these headers, so they are present only if set by App Service...."

So the first thing that happens in the the Web App is to read the information from the request headers that was inserted by the App Service as part of authenticating the user in order to provide the user's claims.

upvoted 6 times

troy89 1 year, 8 months ago

I would say this is correct because the JWT is sent in the header and it contains the claims. No need to access the GraphAPI

upvoted 1 times

hubekpeter 11 months, 1 week ago

Exactly.

upvoted 1 times

wsellmair 1 year, 8 months ago

in Function you can access the user clams from headers, ClaimsPrincipal

For Azure Functions, ClaimsPrincipal.Current is not populated for .NET code, but you can still find the user claims in the request headers, or get the ClaimsPrincipal object from the request context or even through a binding parameter. See working with client identities in Azure Functions for more information.

https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-user-identities

upvoted 1 times

justyoung17 Highly Voted 2 years, 1 month ago

I think it's (A) though because it's a e-commerce 'web app' which is not an 'Azure function'.

upvoted 6 times

ensa 2 years ago

But Azure function will get data from web app and App Service passes user claims to the app by using special request headers. External requests aren't allowed to set these headers, so they are present only if set by the App Service. then
[FunctionName("ClaimsDemo")]
public static IActionResult Run
([HttpTrigger(AuthorizationLevel.Anonymous, "get", "post", Route = null)]
HttpRequest req, ILogger log)
{

//Extract User ID and Claims from the request headers
var principal_name = req.Headers["X-MS-CLIENT-PRINCIPAL-NAME"].FirstOrDefault();
var principal_Id = req.Headers["X-MS-CLIENT-PRINCIPAL-ID"].FirstOrDefault();
So CCCCCCCCCCCCCCCCCCCCC

upvoted 3 times

katrang Most Recent 1 month, 2 weeks ago

Selected Answer: C

App Service makes the claims in the incoming token (whether from an authenticated end user or a client application) available to your code by injecting them into the request headers.

So from this point, you just need to read them in your code

upvoted 1 times

JH81 4 months, 1 week ago

Selected Answer: C

"E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD)." so using Microsoft Graph API call is probably what you want to do but the tricky part is the question asks what should you do first. Read the header to get the info you need to make the Graph API call would be my vote.

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: C

The answer C

upvoted 1 times

wsellmair 1 year, 8 months ago

Selected Answer: C

https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-user-identities

upvoted 1 times

wsellmair 1 year, 8 months ago

in Function you can access the user clams from headers, ClaimsPrincipal

For Azure Functions, ClaimsPrincipal.Current is not populated for .NET code, but you can still find the user claims in the request headers, or get the ClaimsPrincipal object from the request context or even through a binding parameter. See working with client identities in Azure Functions for more information.

https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-user-identities

upvoted 2 times

surprise0011 6 months, 2 weeks ago

Great finding! New link from learn.microsoft:
https://learn.microsoft.com/en-us/azure/app-service/configure-authentication-user-identities

upvoted 1 times

MiraA 2 years ago

Answer is A.
Check this tutorial related to a web app accessing Microsoft Graph to get signed user's display name and his/her photo from Azure AD. It seems it corresponds to "access data from the user claim object" requirement.

https://docs.microsoft.com/en-us/azure/app-service/scenario-secure-app-access-microsoft-graph-as-user

Note: Reading HTTP headers (the answer C) could give basic information only - but the user claim contains more detailed data.

https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-user-identities

upvoted 6 times

gmishra88 1 year ago

The question does not ask anything but claims and does not say read profile (which has detailed information), in which case you might want to access the graphs API if given access.

upvoted 1 times

troy89 1 year, 8 months ago

Yeah, but to access the GraphAPI you need the information from the token which is sent in the header, so in both ways, the first step is to read the header values.

upvoted 1 times

ning 2 years, 2 months ago

I do not know if C is correct, the thing you need is https://docs.microsoft.com/en-us/azure/active-directory/develop/msal-net-migration, you cannot read the http header directly, you need MS library, best thing I can see if graph from A ... But it is not really, the name of the library ...

upvoted 3 times

Topic 29 - Testlet 6

Question #1 *Topic 29*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -
The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
HOTSPOT -
You need to implement the retail store location Azure Function.
How should you configure the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

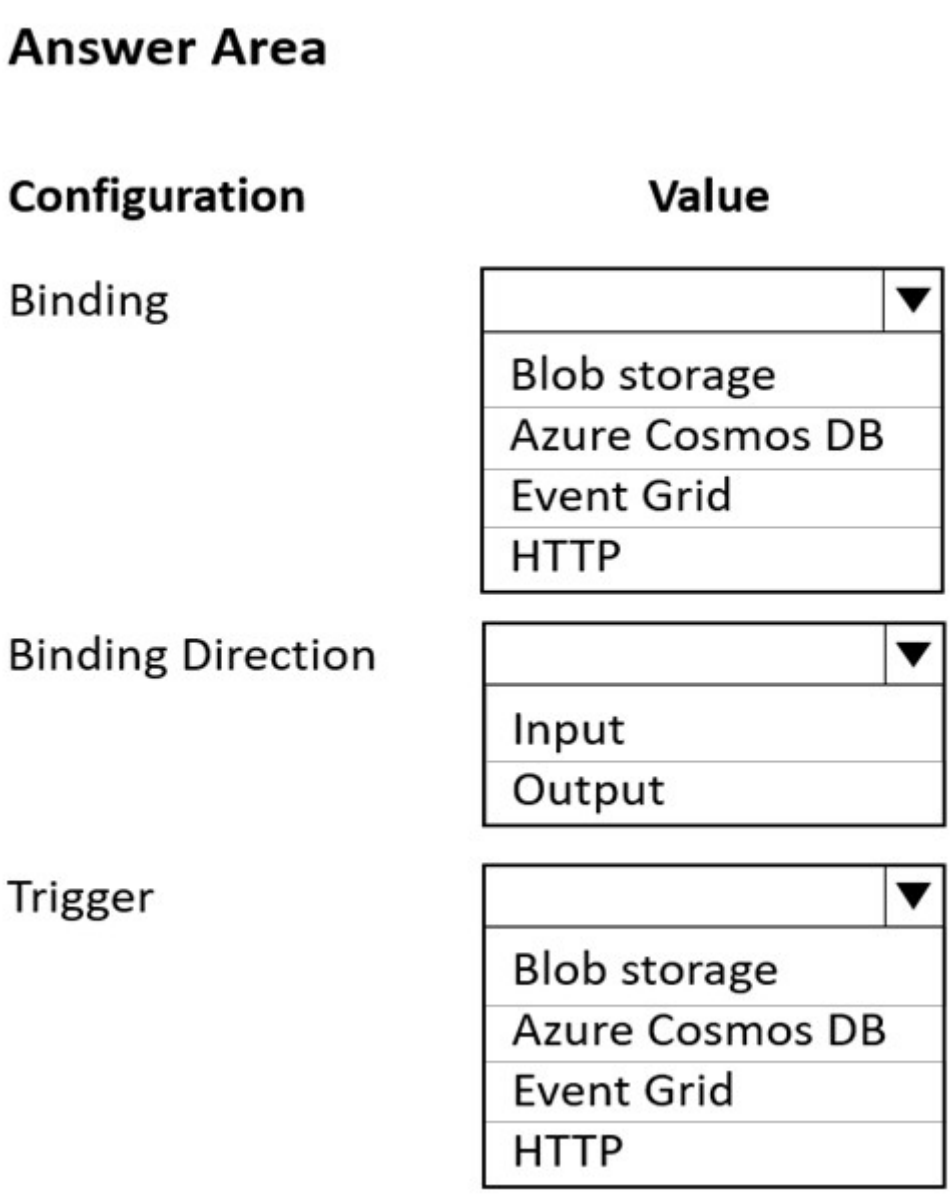

Correct Answer:

Answer Area

Configuration	Value
Binding	Blob storage / Azure Cosmos DB / Event Grid / **HTTP**
Binding Direction	**Input** / Output
Trigger	**Blob storage** / Azure Cosmos DB / Event Grid / HTTP

Scenario: Retail store locations: Azure Functions must process data immediately when data is uploaded to Blob storage.

Box 1: HTTP -
Binding configuration example: https://<storage_account_name>.blob.core.windows.net

Box 2: Input -
Read blob storage data in a function: Input binding

Box 3: Blob storage -
The Blob storage trigger starts a function when a new or updated blob is detected.
Azure Functions integrates with Azure Storage via triggers and bindings. Integrating with Blob storage allows you to build functions that react to changes in blob data as well as read and write values.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger

andrvelich Highly Voted 1 year, 6 months ago

1. azure cosmos
2. output
3. blob storage

upvoted 39 times

Dani_ac7 1 year, 1 month ago

Seems ok, te person who writes answers in exam topics doesn't read questions...

upvoted 10 times

UzrShkl 5 months, 1 week ago

Correct

upvoted 1 times

adilkhan 6 months, 3 weeks ago

100% correct!

upvoted 1 times

SaintBahamut 1 year, 6 months ago

Agreed: We need write to cosmos output, so that deals box1 and 2, and func is triggered by changes on bloob

upvoted 4 times

SenseiJC Highly Voted 11 months ago

1. (binding) = Azure CosmosDB
reason: Azure Functions must update Azure Cosmos DB
2. (direction) = output
reason: all triggers are input, so we are not talking about the trigger binding. we must update Azure Cosmos DB so we need output binding. If input and output were selectable, i'd probably go for that
3. (trigger) = EventGrid
reason: the azure function is triggered by an EventGrid event so that processing happens immediately

upvoted 18 times

macobuzi 1 month, 3 weeks ago

There is no EventGrid mentioned anywhere. You can simply use Azure Function Blob Trigger to get triggered from the Blob storage.

upvoted 1 times

ks321 3 months, 1 week ago

This is the right answer and has to be highest voted

upvoted 1 times

katrang Most Recent 1 month, 2 weeks ago

I would say Cosmos - Output - Blob.

You have a Blob storage trigger (Azure Functions must process data immediately when data is uploaded to Blob storage).
And a Cosmos Output Binding (Azure Functions must update Azure Cosmos DB by using native SQL language queries)

upvoted 1 times

applepie 3 months, 3 weeks ago

I wonder why the answer is not
binding - Blob
direction - in
trigger - event grid.
example scenario 2 - https://learn.microsoft.com/en-us/azure/azure-functions/functions-triggers-bindings?tabs=csharp

upvoted 2 times

macobuzi 1 month, 3 weeks ago

Question stated that Function must update data to Azure Cosmos DB

upvoted 1 times

DenysEx 6 months ago

based on the Topic 21 - Testlet 22, question #1
event routing: BLOB -> EventGrid -> Function

1. CosmosDB
2. Output
3. trigger = EventGrid

upvoted 3 times

AQMA 7 months, 3 weeks ago

1. Azure Cosmos DB
2. Output
3. Event Grid

using Event Grid Trigger for blob events instead of using Blob storage as input to function, provides low latency.
Reference: https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp-script&pivots=programming-language-csharp

according to this reference, it says: "If you require faster or more reliable blob processing, you should instead implement one of the following strategies:
Add the source parameter with a value of EventGrid to your binding definition and create an event subscription on the same container.
OR
Replace the Blob Storage trigger with an Event Grid trigger using an event subscription on the same container."

upvoted 2 times

Nhiendo 9 months, 3 weeks ago

1. azure cosmos
2. output
3. blob storage
Becase
1.Triggers cause a function to run. A trigger defines how a function is invoked and a function must have exactly one trigger. Triggers have associated data, which is often provided as the payload of the function.
2.Binding to a function is a way of declaratively connecting another resource to the function; bindings may be connected as input bindings, output bindings, or both. Data from bindings is provided to the function as parameters.
refer : https://learn.microsoft.com/en-us/azure/azure-functions/functions-triggers-bindings?tabs=csharp

upvoted 2 times

OPT_001122 10 months, 3 weeks ago

1. azure cosmos
2. output
3. blob storage

upvoted 2 times

coffecold 1 year ago

Can someone give a clue from which system this configuration format is?
Azure functions : No
Cosmos DB : No
Azure Event Grid : No
Storage blob : No

upvoted 1 times

coffecold 1 year ago

If it was function.json it should have been someting like:

{
"disabled":false,
"bindings":[
// ... bindings here
{
"type": "EventGridTrigger",
"direction": "in",
"name": "inBinding",
// ... more depending on binding
},
{
"type": "CosmosDBOutput",
"direction": "out",
"name": "outBinding",
// ... more depending on binding
}
]
}

upvoted 6 times

OPT_001122 11 months, 2 weeks ago

this is more convincing !! Thanks

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

this is more convincing !! Thanks

upvoted 1 times

elequiel 1 year ago

got for me in 20/10/2022

upvoted 3 times

darsak10 5 months, 3 weeks ago

Can you confirm what is the correct answer, it looks like the answer is not clear yet.

upvoted 1 times

Ciupaz 1 day, 13 hours ago

Should be:
Azure Cosmos DB
Output
Event Grid

upvoted 1 times

gmishra88 1 year ago

I do not understand this question. Binding is Azure cosmos, then the binding direction should be output. If it is trigger then binding direction should be in. So, I do not understand what this means. I just hope these kind of absurd questions are not asked by Microsoft.
But for the trigger I will use EventGrid because it says it should respond immediately. Blob triggers use polling, while event grid triggers are "event-driven" (basically Microsoft wants us to better use Event Grid for storage events)

upvoted 2 times

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
HOTSPOT -
You need to implement the corporate website.
How should you configure the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Azure	Configuration
Plan	Free Standard Premium Isolated
Service	App Service Web App App Service Static Web App Azure Function App Azure Blob Storage

Correct Answer:

Answer Area

Azure	Configuration
Plan	Free **Standard** Premium Isolated
Service	**App Service Web App** App Service Static Web App Azure Function App Azure Blob Storage

Box 1: Standard -

Below is a high-level comparison of the features as per the pricing tier for the App Service Plan.

SKUs	FREE	SHARED	BASIC	STANDARD	PREMIUM	ISOLATED*
Limits						
Apps	10	100	Unlimited	Unlimited	Unlimited	Unlimited
Disk space	1 GB	1 GB	10 GB	50 GB	250 GB	
Max instances			Up to 3	Up to 10	Up to 20	
SLA			99.95%	99.95%	99.95%	
App Deployment						
Continuous Deployment	Available	Available	Available	Available	Available	Available
Deployment Slots				Available	Available	Available
Development Tools						
Clone App					Available	Available
Site Extensions	Available	Available	Available	Available	Available	Available
Testing in Production				Available	Available	Available

Note: Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Corporate website requirements:
☞ Secure the website by using SSL.
☞ Minimize costs for data storage and hosting.
☞ Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).
☞ Distribute the website content globally for local use.
☞ Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.
☞ The website must have 99.95 percent uptime.

Box 2: App Service Web App -
A Web App is a web application that is hosted in an App Service. The App Service is the managed service in Azure that enables you to deploy a web application and make it available to your customers on the Internet in a very short amount of time.
Incorrect:
A Static Web Application is any web application that can be delivered directly to an end user's browser without any server-side alteration of the HTML, CSS, or
JavaScript content.
Reference:
https://azure-training.com/2018/12/27/understanding-app-services-app-service-plan-and-ase/ https://docs.microsoft.com/en-us/azure/app-service/overview

AbdulMannan Highly Voted 1 year, 1 month ago
Got this question on 30-Sep-2022 exam.
Answer is correct. Passed with 870 score.
upvoted 8 times

darsak10 Most Recent 5 months, 3 weeks ago
Standard plan with app service web app would be the appropriate option.
upvoted 2 times

SenseiJC 11 months ago
Static web apps only support HTTP triggered functions (https://learn.microsoft.com/en-us/azure/static-web-apps/functions-bring-your-own#link-an-existing-azure-functions-app). The top of the case study reads: "several APIs hosted in Azure Functions."

If we assume that the referenced azure functions here (APIs) are NOT the Azure functions for processing blob storage, then they can be HTTP triggered and we are fine choosing Static Web App.

I think this is a reasonable assumption, but I leave it to the test taker to interpret which is correct.
upvoted 3 times

OPT_001122 11 months, 2 weeks ago
Answer is correct.
upvoted 1 times

OPT_001122 11 months, 2 weeks ago
All below requirements are met with standard plan

Secure the website by using SSL.
Minimize costs for data storage and hosting.
Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).
Distribute the website content globally for local use.
Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.
The website must have 99.95 percent uptime.
upvoted 3 times

coffecold 1 year ago
I think the tier is Standard is because of the Traffic manager support.
requirement : "Distribute the website content globally for local use",
Distributions slots (in the answer given) are not required.
(+ static web app)
upvoted 2 times

Mousavi 1 year, 1 month ago
1-CI/CD
2-including SSL certificate
3-Distribute the website content globally
4-Minimize costs

with these criteria we can use standard Plan and for the web site we have to choose azure app service

upvoted 2 times

gmishra88 1 year ago

I would have said host it in azure static web app (a fine option) that has static part as well as support functions. A perfect option.
If functions are to be created outside in another function app and then it contains only the static part then I will host it in blob-storage to get the 99.95% availability SLA of blob storage hot tier and CDN is also easy.
It is very difficult to choose one when I cannot look into the head of the question creator. Another vague question from the "M$"

upvoted 3 times

gmishra88 1 year ago

It is DEFINITELY static web app and Standard plan. Check the availability SLA. Static web app guarantees 99.95% availability and it supports git workflow. It has Functions capability. And in practice also it is a beautiful solution unlike many solutions from MIcrosoft

upvoted 1 times

gmishra88 1 year ago

See the link for Static web app: https://azure.microsoft.com/en-us/products/app-service/static/#features
I have used this in real life (but do not know if available for enterprise customers)

upvoted 1 times

gmishra88 1 year ago

Yes, another evil question to sow confusion. Microsoft takes it personally when people answer correctly

upvoted 1 times

Topic 30 - Testlet 7

Question #1 *Topic 30*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment -

Corporate website -

The company provides a public website located at http://www.vanarsdelltd.com. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

Retail Store Locations -

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

Requirements -

The application components must meet the following requirements:

Corporate website -

Secure the website by using SSL.

Minimize costs for data storage and hosting.

Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).

Distribute the website content globally for local use.

Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

The website must have 99.95 percent uptime.

Retail store locations -

Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.

Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services -

Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services -
The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security -
All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
Authentication and authorization must use Azure AD and services must use managed identities where possible.

Issues -

Retail Store Locations -
You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question
You need to implement a solution to resolve the retail store location data issue.
Which three Azure Blob features should you enable? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Soft delete

B. Change feed

C. Snapshots

D. Versioning

E. Object replication

F. Immutability

Correct Answer: *ABD*
Scenario: You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Before you enable and configure point-in-time restore, enable its prerequisites for the storage account: soft delete, change feed, and blob versioning.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/point-in-time-restore-manage

Community vote distribution
ABD (100%)

OPT_001122 Highly Voted 11 months, 2 weeks ago
Selected Answer: ABD
a b and d
upvoted 13 times

andrvelich Highly Voted 1 year, 6 months ago
correct
upvoted 7 times

juanckar Most Recent 3 months, 3 weeks ago
This was on the exam (July 2023). Went with highly voted. Scored 917
upvoted 4 times

AQMA 7 months, 3 weeks ago
Selected Answer: ABD
https://learn.microsoft.com/en-us/azure/storage/blobs/point-in-time-restore-overview#prerequisites-for-point-in-time-restore
according to the above reference Point-in-time restore requires that the following Azure Storage features be enabled before you can enable point-in-time restore:

Soft delete
Change feed
Blob versioning
upvoted 5 times

coffecold 1 year ago

In portal
"Select Turn on point-in-time restore. When you select this option, soft delete for blobs, versioning, and change feed are also enabled"
So it needs the three options to do point in time restore. You probably can't clear them, once Turn on point-in-time is chosen.

upvoted 3 times

gmishra88 1 year ago

"Before you enable and configure point-in-time restore, enable its prerequisites for the storage account: soft delete, change feed, and blob versioning"

upvoted 2 times

AbdulMannan 1 year, 1 month ago

Got this question on 30-Sep-2022 exam.
Answer is correct. Passed with 870 score.

upvoted 3 times

Mousavi 1 year ago

I will take exam on 5 October. Any suggestion?

upvoted 3 times

Topic 31 - Testlet 8

Question #1 *Topic 31*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -
You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         -  port: 80
CS12            protocol: TCP
CS13         resources:
CS14           requests:
CS15              cpu: 1.0
CS16              memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14        "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

You need to store the user agreements.

Where should you store the agreement after it is completed?

A. Azure Storage queue

B. Azure Event Hub

C. Azure Service Bus topic

D. Azure Event Grid topic

Correct Answer: *B*

Azure Event Hub is used for telemetry and distributed data streaming.

This service provides a single solution that enables rapid data retrieval for real-time processing as well as repeated replay of stored raw data. It can capture the streaming data into a file for processing and analysis.

It has the following characteristics:

- low latency
- capable of receiving and processing millions of events per second
- at least once delivery

Reference:

https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

Community vote distribution

B (100%)

perry230 Highly Voted 3 years, 4 months ago

Correct: "Information regarding agreements is used by multiple divisions within Contoso, Ltd. User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour."

upvoted 33 times

KingChuang 10 months, 1 week ago

On My Exam:2022-12-26
My Choice:
B:Azure Event Hub

upvoted 2 times

GCMan 3 years ago

I think Azure Event Hub is correct but I see arguments of a couple other options.

upvoted 1 times

Juanlu 2 years, 11 months ago

As @coolest said: B is correct, so the answer is correct. Note: "You can use Event Capture to store the agreements into Azure blob storage for long term storage".

upvoted 1 times

coffecold 1 year ago

"If you need to store data for more than 7 days, a feature of Azure Event Hubs called Capture is the preferred solution for longer-term storage. When configuring Capture, there are two locations where this information can be stored: Azure Blob Storage or Azure Data Lake Store account"
So Event capture just uses a storage solution as well.

upvoted 2 times

coolest Highly Voted 3 years, 2 months ago

B - Azure Event Hub is Correct.

If you are looking at millions of agreements per hour, you need to use a data ingestion service like the Azure Event Hub. You can use Event Capture to store the agreements into Azure blob storage for long term storage.

upvoted 15 times

Ciupaz Most Recent 1 week, 4 days ago

Selected Answer: B

Should be Azure Blob Storage, if was a choice...
Azure Blob Storage is a highly scalable object storage service that can store any type of data. It is a good option for storing user agreements for long periods of time because it is cost-effective and durable.

upvoted 1 times

ENGs 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

BaoNguyen2411 3 months ago

Got this question on 29/06/2023

upvoted 1 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: B

Azure Event Hub

upvoted 2 times

gmishra88 1 year ago

Another brilliant question from Microsoft. The wording is very vague. Where to store after it is completed. After it is completed, store it in Archive till the retention period. But because that is not in the answer, and I guess Microsoft does not potentially understand the meaning of "complete" properly, and guessing that word is added just to add some: entertainment: I choose event-hub
But you should not use any of these options for storing information. But yes, you can use events capture with storage blob, but that is a lot of thinking

upvoted 6 times

coffecold 1 year ago

Storing is storing, it is not the 7-days data that is kept in retention in the event hub. Completed is if the data is finally stored in a storage solution. So I go for option A

upvoted 1 times

meoukg 1 year, 7 months ago

Got it on 03/2022!

upvoted 2 times

aruni_mishra 1 year, 9 months ago

Azure Event Hubs enables you to automatically capture the streaming data in Event Hubs in an Azure Blob storage or Azure Data Lake Storage Gen 1 or Gen 2 account of your choice, with the added flexibility of specifying a time or size interval.
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-capture-overview
Ans: Azure Event Hub

upvoted 1 times

altafpatel1984 1 year, 11 months ago

Azure Storage Queue can process 2000 messages per second. i.e. 72 million messages per hour. Since message is to be processed here, it cannot be event and hence Storage Queue will be used to store data and hence answer is A - Storage Queue

https://docs.microsoft.com/en-us/learn/modules/communicate-between-apps-with-azure-queue-storage/2-create-the-azure-storage-infrastructure

upvoted 3 times

Bartimaeus 1 year, 8 months ago

It's 7.2 million, but it's only target throughput - max is 20_000/s = 72 million

upvoted 1 times

paru123456789 2 years, 8 months ago

Answer: B

upvoted 3 times

cbn 2 years, 8 months ago

"Information regarding agreements is used by multiple divisions within Contoso, Ltd."
- This needs multiple subscribers. Storage queue cannot be used for this.

"The volume of agreements is expected to be in the millions per hour."
- This leads to choice for Event Hub / Service bus topic

"When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version"
- This is not about agreements, however an Event Hub supports this scenario as well.

I will go with C (Event Hub)
upvoted 5 times

cbn 2 years, 8 months ago
Sorry, I mean B (Event Hub)
upvoted 2 times

mvr 2 years, 9 months ago
Ok, yeah, so we need to handle millions. So Storage Queue or Event Hub.
But then, do we need messages or events?
It says we need to store (maybe temporarily, until processed), and events do not contain content to be stored. So I'd say Storage Queue. Right?
upvoted 5 times

matejka 2 years, 9 months ago
Azure Event Hub is correct as it is able to handle milions of events per second. But the questions is very confusing as it states "to store..." and Event Hub is not designed to persistent store data. It stores the events and should route those to eg. storage account or so.
upvoted 2 times

khoant 2 years, 9 months ago
should be B
upvoted 1 times

Cornholioz 2 years, 10 months ago
I still don't know for sure. The scenario gives a few facts and the question asks to store the agreements. Couldn't find a strong argument to say Event Hub is better than Queue in this case. I can eliminate Service Bus Topic and Event Grid Topic, because of the traffic numbers. But between the other two, it's hard.
Even after all the research, I'm only making a guess here: Event Hub
upvoted 1 times

luppittegui 2 years, 10 months ago
Who talk about store? The case says: "When a user submits content, they must agree to a user agreement", so ALWAYS a user submits content, they MUST agree.
"Event Hubs Standard tier currently supports a maximum retention period of seven days." so it's more than enough for an approval process
upvoted 1 times

Question #2 *Topic 31*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -

You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14        "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

HOTSPOT -

You need to implement the bindings for the CheckUserContent function.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
public static class CheckUserContent
{
   [FunctionName("CheckUserContent")]
   public static void Run(
      [QueueTrigger("userContent")] / [BlobTrigger("userContent/{name}")] / [CosmosDBTrigger("content", "userContent")] / [Table("content", "userContent", "{name}")]
      string content,
      [Queue("userContent")] / [CosmosDB("content", "userContent")] / [Table("content", "userContent", "{name}")] / [Blob("userContent/{name}", FileAccess.Write)]
      Stream output)
   {
      ...
   }
}
```

Correct Answer:

Answer Area

```
public static class CheckUserContent
{
    [FunctionName("CheckUserContent")]
    public static void Run(
        [QueueTrigger("userContent")]
        [BlobTrigger("userContent/{name}")]
        [CosmosDBTrigger("content", "userContent")]
        [Table("content", "userContent", "{name}")]
        string content,
        [Queue("userContent")]
        [CosmosDB("content", "userContent")]
        [Table("content", "userContent", "{name}")]
        [Blob("userContent/{name}", FileAccess.Write)]
        Stream output)
    {
        ...
    }
}
```

Box 1: [BlobTrigger(..)]

Box 2: [Blob(..)]

Azure Blob storage output binding for Azure Functions. The output binding allows you to modify and delete blob storage data in an Azure Function.

The attribute's constructor takes the path to the blob and a FileAccess parameter indicating read or write, as shown in the following example:

```
[FunctionName("ResizeImage")]
public static void Run(
[BlobTrigger("sample-images/{name}")] Stream image,
[Blob("sample-images-md/{name}", FileAccess.Write)] Stream imageSmall)
{
...
}
```

Scenario: You must create an Azure Function named CheckUserContent to perform the content checks.

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-output

MohmmadFayez Highly Voted 2 years, 3 months ago

1- Queue Trigre[]
2- Blob[]

upvoted 40 times

warchoon 6 months, 3 weeks ago

agree
This is the only one reason to transform a string to a stream in the binding function

upvoted 1 times

KingChuang 10 months, 1 week ago

On My Exam:2022-12-26
My Answer:
Quere Trigger
Blob

upvoted 1 times

MiraA 2 years ago

The "content" parameter is of type "string" so it must be QueueTrigger.
For example BlobTrigger uses Stream type, CosmosDBTrigger uses IReadOnlyList<> type, it seems the Table Storage has no trigger binding.

https://docs.microsoft.com/en-us/azure/azure-functions/functions-triggers-bindings?tabs=csharp#supported-bindings
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue-trigger

https://docs.microsoft.com/cs-cz/azure/azure-functions/functions-bindings-storage-blob-trigger
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-cosmosdb-v2-trigger
upvoted 16 times

Mev4953 1 year, 9 months ago

BlobTrigge can also use string parameter

[FunctionName("BlobTriggerCSharp")]
public static void Run([BlobTrigger("samples-workitems/{name}")] Stream myBlob, string name, ILogger log)
{

https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp
upvoted 4 times

PhilLI 1 year, 9 months ago

Agree with Azure Queue Storage trigger indeed (see example on https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue-trigger?tabs=csharp#example)
upvoted 1 times

Lucario95 1 year, 11 months ago

This is correct according to the data type of "content" for the first binding and "output" for the second.
upvoted 1 times

clarionprogrammer Highly Voted 2 years, 6 months ago

Correct.
Ref: https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-output?tabs=csharp
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-input?tabs=csharp
upvoted 14 times

Dianahu Most Recent 3 months ago

https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-input?tabs=python-v2%2Cin-process&pivots=programming-language-csharp
[FunctionName("BlobInput")]
public static void Run(
[QueueTrigger("myqueue-items")] string myQueueItem,
[Blob("samples-workitems/{queueTrigger}", FileAccess.Read)] Stream myBlob,
ILogger log)
{
log.LogInformation($"BlobInput processed blob\n Name:{myQueueItem} \n Size: {myBlob.Length} bytes");
}
upvoted 1 times

JH81 4 months, 1 week ago

Terrible requirements that leave us questioning the actual implementation. But I have to agree that QueueTrigger is the first answer because the requirements state "accepts user generated content as a string". That can only be the QueueTrigger. Blob for the output is not in question.
upvoted 1 times

gmishra88 1 year ago

Another CSI Miami question from Microsoft.
Yes, detectives, the content has to be a Stream if it is Blob trigger.
Microsoft (don't be evil) did not give any information on how the function is triggered, but stereotypically put the red-herring in the description: " website shows the content from azure storage blob". Of course that does not mean the function is triggered from it. Next my only clue is the type of this parameter.
Microsoft says .Net knowledge is required and any SDK is enough. But how will a javascript person know this that it is Stream and not the content given as String?
upvoted 8 times

le129 1 year, 1 month ago

Blob trigger seems correct. trigger binding can be string if the content is small.
Binding to string, or Byte[] is only recommended when the blob size is small. This is recommended because the entire blob contents are loaded into memory. For most blobs, use a Stream or CloudBlockBlob type. For more information, see Concurrency and memory usage.

https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=in-process%2Cextensionv5&pivots=programming-language-csharp
upvoted 1 times

shawnz 1 year, 8 months ago

the question said accept user upload content as a string.
and user upload to blob to trigger the content check.
so blobtrigger and blob are correct.
upvoted 3 times

leonidn 1 year, 9 months ago

"content" is of type string. Then QueueTrigger is the only what that is applicable. Agree with "blob".
upvoted 3 times

gfiorini 1 year, 11 months ago

It is really annoying the way the question is worded. Where in the specs is written the source and the target of the Function ?!? It is true that we can infer it from the possible answers. But why dont ask then "which are valid input / output" for the function ?

upvoted 7 times

gmishra88 1 year ago

It is Microsoft

upvoted 2 times

phvogel 2 years ago

The question has nothing to do with the architecture or design. It's actually asking if you can recognize valid trigger and output bindings. The only valid bindings are Queue Trigger (string) and Blob output (File.Write supports writing to a stream).

upvoted 3 times

ning 2 years, 2 months ago

I believe this is blobtrigger --> user upload contents into blob storage ...
Need send to a service ... So I am think send to a queue for processing ...
Just from the given info ...

upvoted 2 times

mcanic 1 year, 10 months ago

Blob trigger recieves stream, Service Bus Queue trigger recieves string, therefore the first option is not blob trigger but a queue trigger

upvoted 2 times

ning 2 years, 2 months ago

Depends on your interpretation for what is final step for azure function, whether to send the contents to the service, or remove the contents from the blob storage ... The requirements are not clear ...

upvoted 2 times

j888 2 years, 2 months ago

This is tricky, storage will be my general choice, however this statement "Messages are sent to contentuploadservice" leading me to believe this is a service bus.

upvoted 1 times

j888 2 years, 2 months ago

I was wrong.
The service bus trigger would be
public static void Run([ServiceBusTrigger("CustomerQueue"....

upvoted 1 times

SuperPetey 2 years, 5 months ago

I believe the answer is QueueTrigger and Queue. This is an architectural design issue -- the entire reason to do user content auditing is to prohibit bad content from being stored and used. Therefore, use a queue to store these messages until they are verified. No need to store content in a blob until then -- Azure Storage Queue perfect solution.

upvoted 6 times

warchoon 6 months, 3 weeks ago

So why did you transform stream to string and then transform it to stream again?

upvoted 1 times

Spooky7 2 years, 4 months ago

Well, based what information was given it seems that entire design looks little bit different. User post content and it is saved in BlobStorage and available immediately. After that check is done and if content is invalid then it is replaced. So given answer is correct.

upvoted 5 times

kwaazaar 2 years, 6 months ago

Correct

upvoted 6 times

Question #3 *Topic 31*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

Overview -

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

* Messages are sent to ContentUploadService.
* Content is processed by ContentAnalysisService.
* After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory -

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements -

ContentAnalysisService -

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs -

You must minimize costs for all Azure services.

Manual review -

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using

React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability -

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring -

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security -

You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.

User agreements -
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing -
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues -
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code -

ContentUploadService -

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06     containers:
CS07     - name: service
CS08         properties:
CS09         image: contoso/contentUploadService:latest
CS10         ports:
CS11         - port: 80
CS12           protocol: TCP
CS13         resources:
CS14           requests:
CS15             cpu: 1.0
CS16             memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19     ip: 10.23.121.112
CS20     ports:
CS21       - port: 80
CS22         protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest -

```
AM01 {
AM02      "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03      "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05      "createdDateTime" : "2019-12-24T06:01:44Z",
AM06      "logoUrl" : null,
AM07      "logoutUrl" : null,
AM08      "name" : "ContentAnalysisService",
AM09
AM10
AM11      "orgRestrictions" : [],
AM12      "parentalControlSettings" : {
AM13       "countriesBlockedForMinors" : [],
AM14        "legalAgeGroupRule" : "Allow"
AM15      },
AM16      "passwordCredentials" : []
AM17 }
```

Question

You need to configure the ContentUploadService deployment.

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

NOTE: Each correct selection is worth one point.

A. Add the following markup to line CS23: type: Private

B. Add the following markup to line CS24: osType: Windows

C. Add the following markup to line CS24: osType: Linux

D. Add the following markup to line CS23: type: Public

Correct Answer: *A*

Scenario: All Internal services must only be accessible from Internal Virtual Networks (VNets)

There are three Network Location types ג€" Private, Public and Domain

Reference:

https://devblogs.microsoft.com/powershell/setting-network-location-to-private/

Community vote distribution

A (50%) C (50%)

Mo_Mo_01 Highly Voted 2 years, 7 months ago

A

https://docs.microsoft.com/en-us/azure/container-instances/container-instances-virtual-network-concepts#unsupported-networking-scenarios

C

https://docs.microsoft.com/en-us/azure/container-instances/container-instances-region-availability#windows-container-groups

upvoted 40 times

KingChuang 10 months, 1 week ago

On My Exam:2022-12-26

My Answer:

A、C

upvoted 2 times

jokergester 2 years, 7 months ago

yep currently, container groups are not supported on Windows containers

upvoted 6 times

SnakePlissken Highly Voted 2 years, 5 months ago

Besides, IP address is in the private address range 10.0.0.0/24, so it can only be private.

upvoted 8 times

ENGs Most Recent 3 weeks, 4 days ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

BaoNguyen2411 3 months ago

go this question on 29/06/2023
My answer: A/ C

upvoted 1 times

JH81 4 months, 1 week ago

Selected Answer: C

I get that we can infer its private because of the IP address range 10. or 192. are private but having to know only that container groups are only supported on Linux. Why?

upvoted 1 times

CODE_STS 8 months ago

Got this in the exam today! Feb 28, 2023

upvoted 2 times

AlexeyG 8 months, 2 weeks ago

Got this in 16/02/2023

upvoted 3 times

rotimislaw 9 months, 4 weeks ago

Selected Answer: A

A. Private - it's internal app
C. Linux - container groups not supported on Win

upvoted 4 times

OPT_001122 11 months, 2 weeks ago

Selected Answer: A

private
linux

upvoted 1 times

coffecold 11 months, 4 weeks ago

Type belongs to the ipAddress settings
OSType belongs to the container group
https://learn.microsoft.com/en-us/azure/container-instances/container-instances-reference-yaml#schema

upvoted 1 times

Eltooth 1 year, 4 months ago

Selected Answer: C

A & C are correct answers.

upvoted 4 times

petitbilly 1 year, 7 months ago

Got this one 03/2022

upvoted 2 times

SivajiTheBoss 1 year, 7 months ago

Both A and C

Currently, only Linux containers are supported in a container group deployed to a virtual network.

upvoted 1 times

leonidn 1 year, 9 months ago

A & C
Container group deployment to a virtual network is generally available for Linux containers. https://docs.microsoft.com/en-us/azure/container-instances/container-instances-virtual-network-concepts#unsupported-networking-scenarios

All Internal services must only be accessible from internal Virtual Networks (VNets).

upvoted 2 times

Mev4953 1 year, 9 months ago

10.23.121.112 is private range IP

upvoted 1 times

MiraA 2 years ago

What about this tutorial with sample YAML file?
https://docs.microsoft.com/en-us/azure/container-instances/container-instances-multi-container-yaml#configure-a-yaml-file

There is "type: Public" but no "ip" value together with "type: Microsoft.ContainerInstance/containerGroups".

Does it mean the "Public" option is supported now?
Or specifying "ip" with value "10.x.x.x" in the assignment still means the private network?

upvoted 1 times

mc0re 2 years, 2 months ago

Where did the "container group" requirement come from?
upvoted 2 times

gmishra88 1 year ago
Good question, I had to cross check to see it. It is an eye test not a developer certification.
upvoted 1 times

MiraA 2 years ago
See line CS02:
type: Microsoft.ContainerInstance/containerGroups
upvoted 4 times

Topic 32 - Testlet 9

Question #1 *Topic 32*

Introductory Info

Case study -

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

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

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

To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background -

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment -

Architecture overview -

The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and

Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram -

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

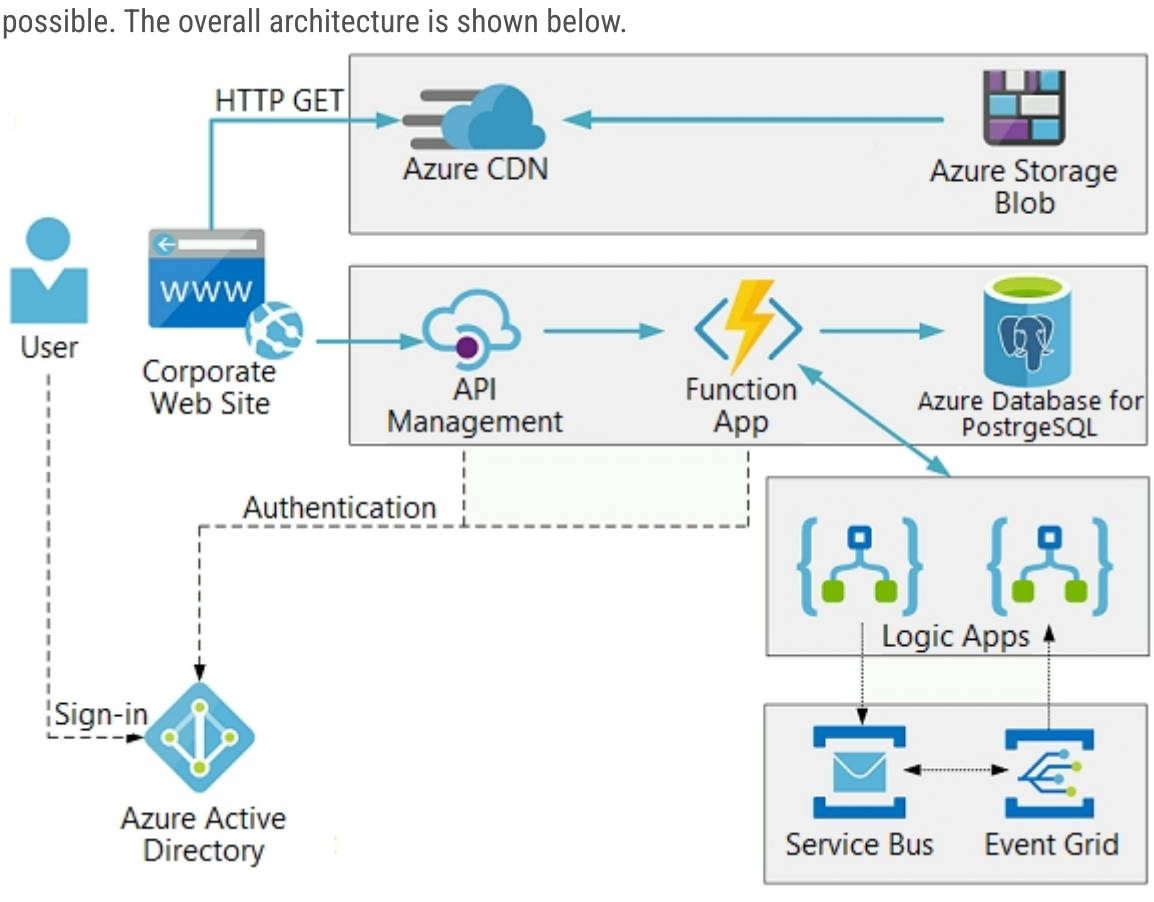

User authentication -

The following steps detail the user authentication process:
1. The user selects Sign in in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements -

Corporate website -
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.

Azure Database for PostgreSQL -
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn
Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security -
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance -
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Issues -

Corporate website -
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.

Function app -
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs -
| where FunctionName = = "RequestUserApproval"

Logic app -
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code -

Corporate website -

Security.cs:

```
SC01 public class Security
SC02 {
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);
SC05 var certName = cert.FriendlyName;
SC06 }
```

Function app -

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval
RA02 {
RA03 [FunctionName("RequestUserApproval")]
RA04 public static async Task<IActionResult> Run(
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req.
ILogger log)
RA06 {
RA07    log.LogInformation("RequestUserApproval function processed a request.");
RA08 ...
RA09 return ProcessRequest(req)
RA10    ? (ActionResult)new OkObjectResult($"User approval processed")
RA11    : new BadRequestObjectResult("Failed to process user approval");
RA12 }
RA13 private static bool ProcessRequest(HttpRequest req)
RA14 {
RA15    ...
RA16 }
RA17 }
```

Question

HOTSPOT -

You need to configure the Account Kind, Replication, and Access tier options for the corporate website's Azure Storage account.

How should you complete the configuration? To answer, select the appropriate options in the dialog box in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

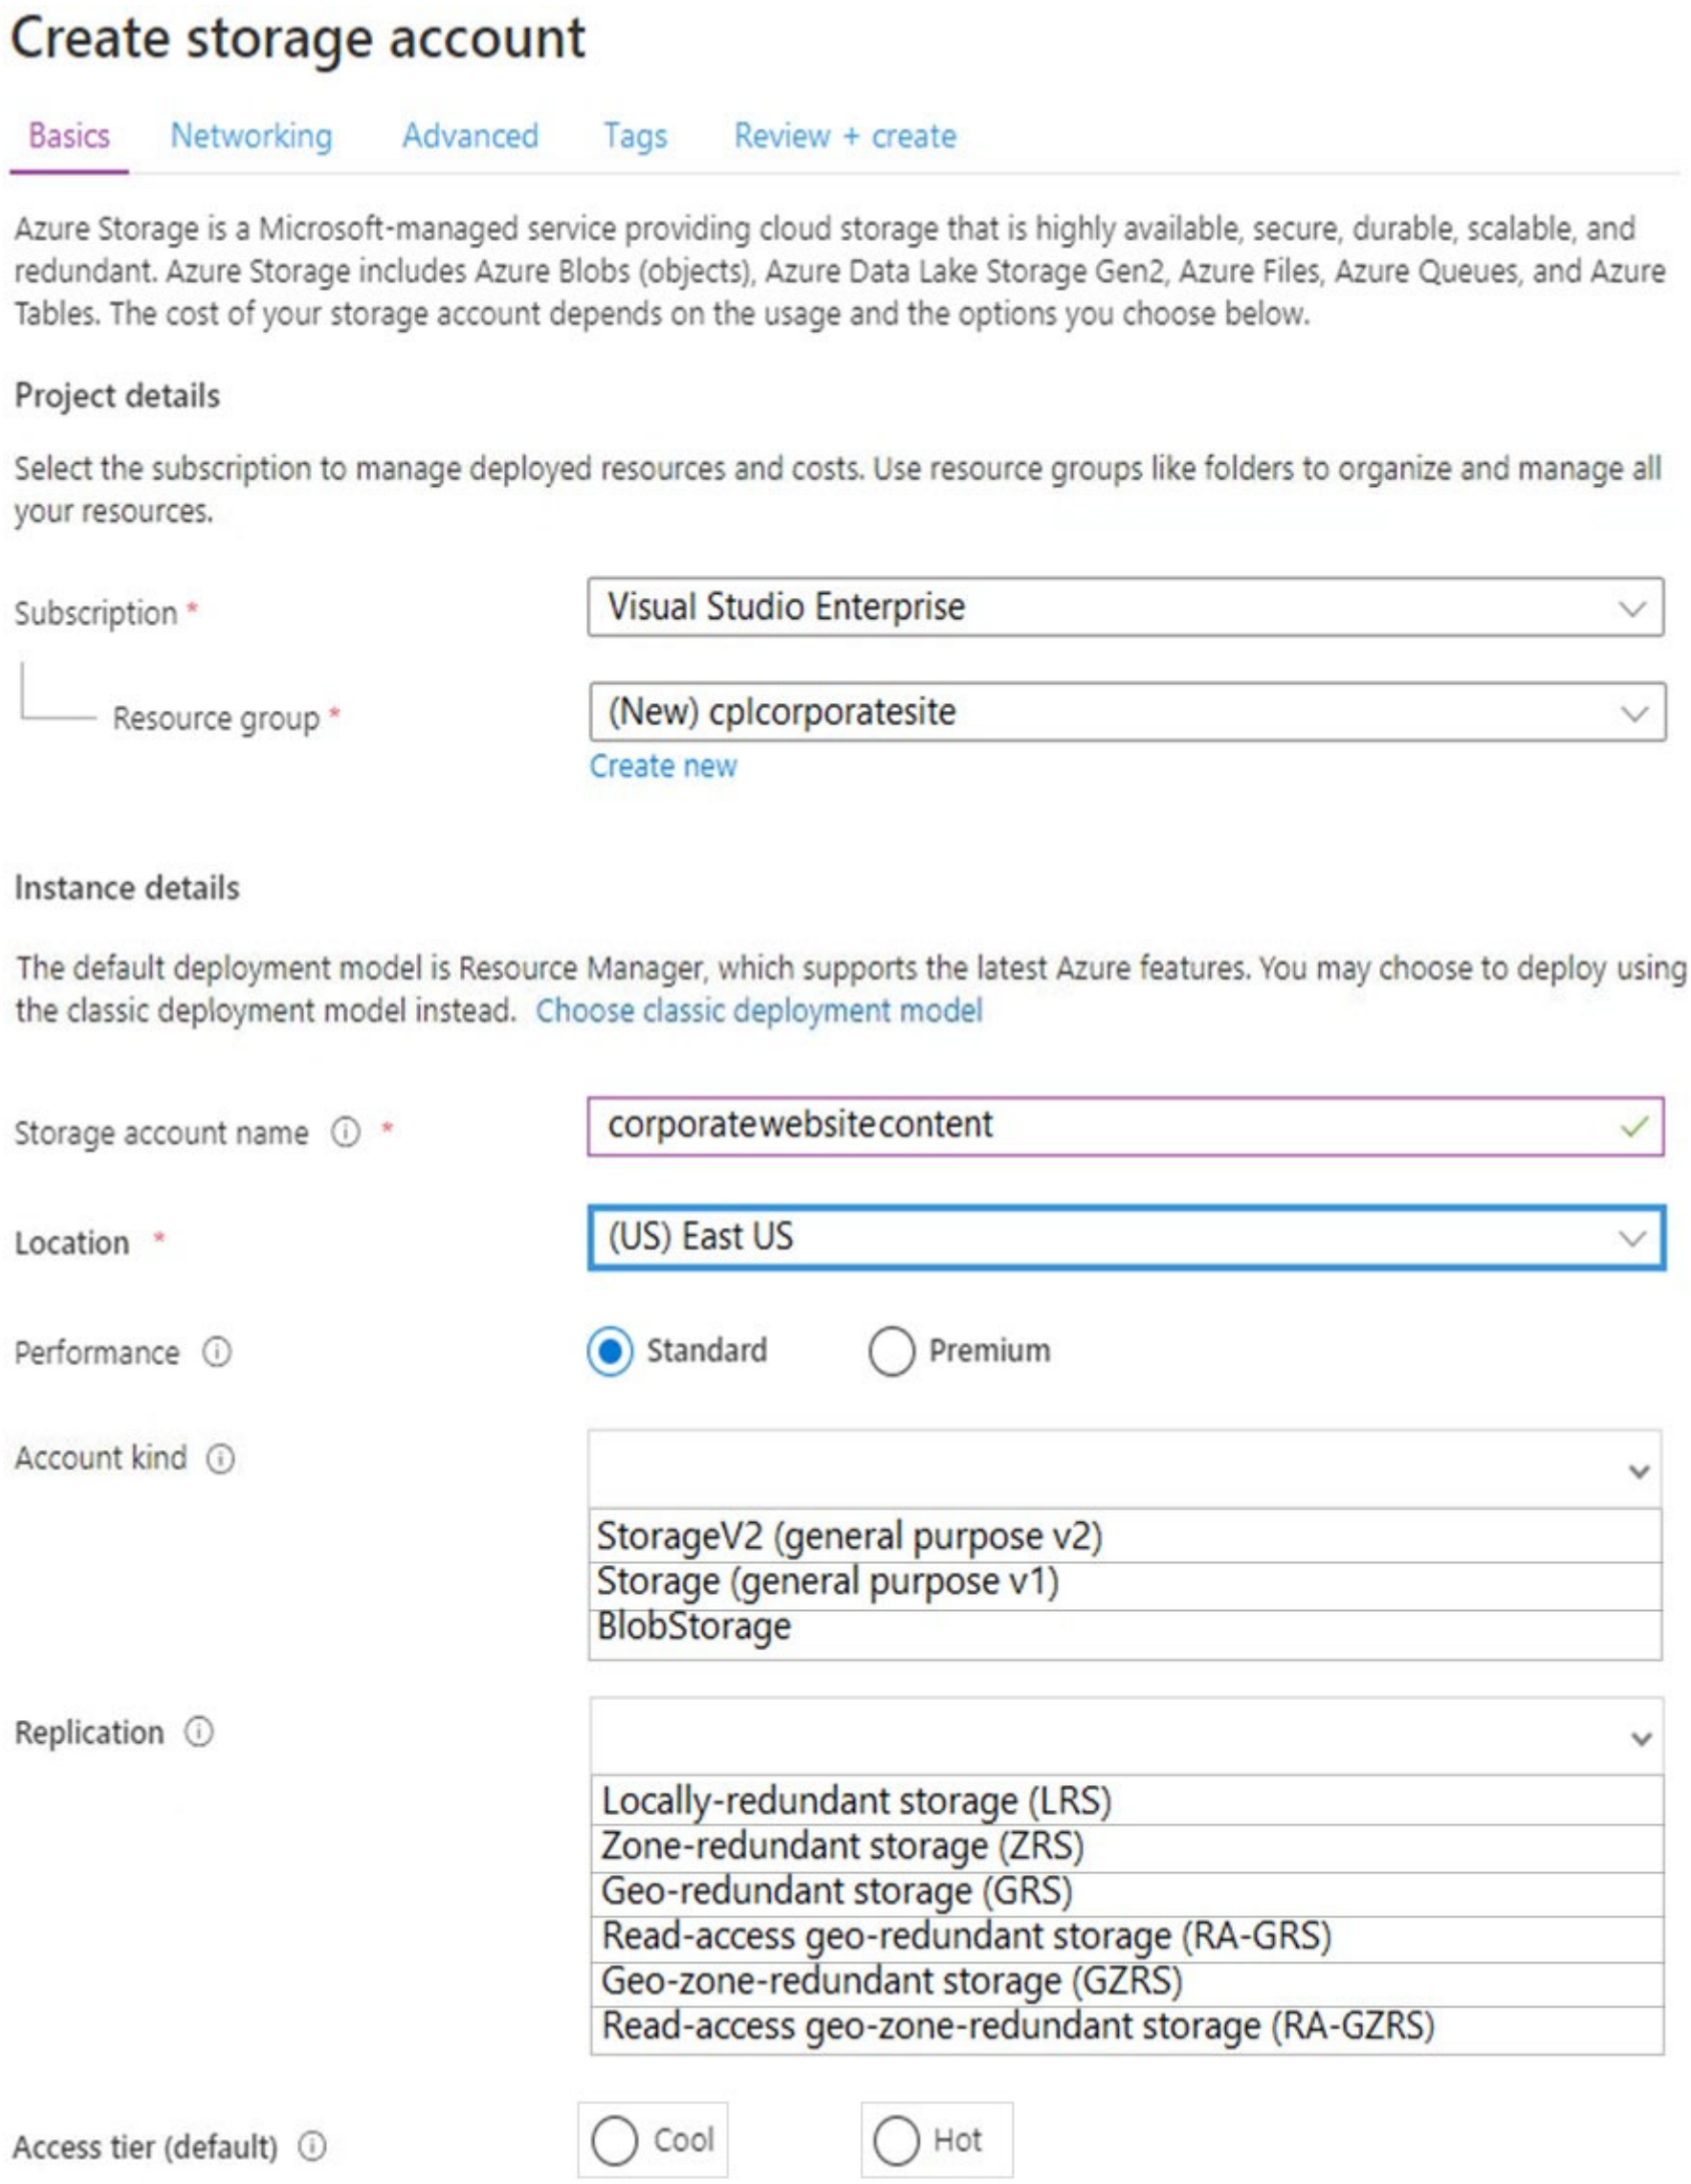

Correct Answer:

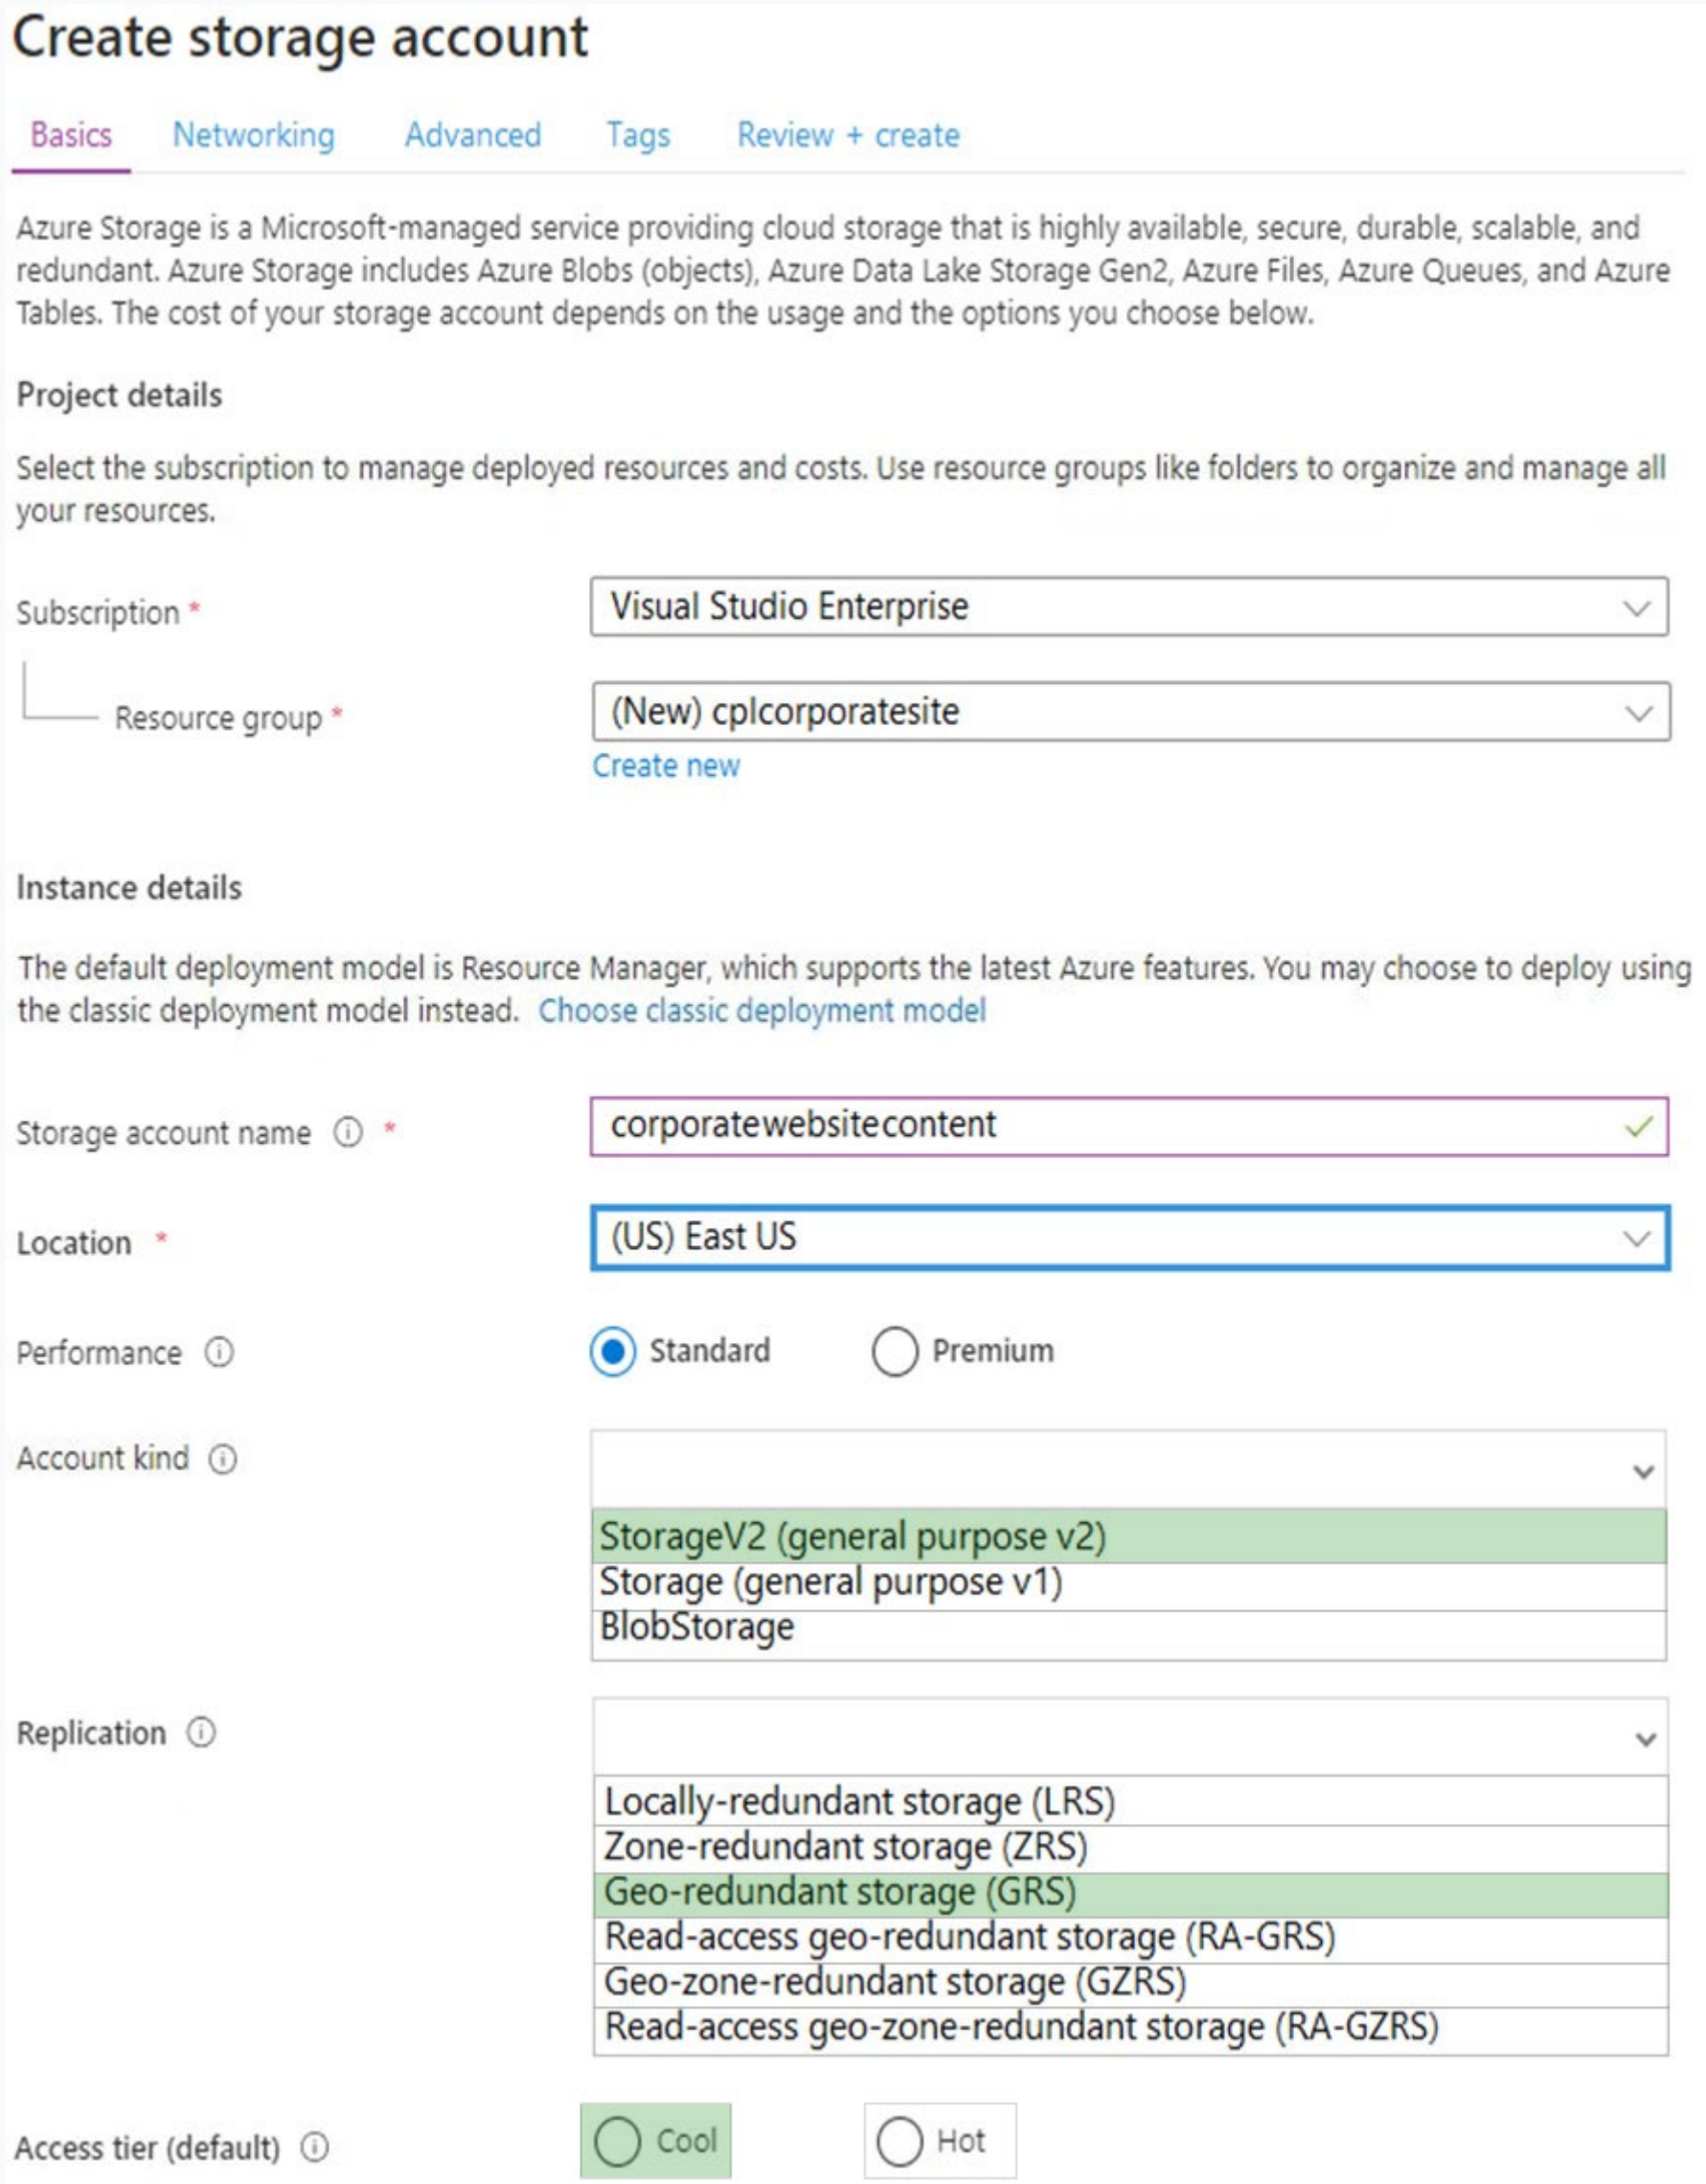

Account Kind: StorageV2 (general-purpose v2)

Scenario: Azure Storage blob will be used (refer to the exhibit). Data storage costs must be minimized.

General-purpose v2 accounts: Basic storage account type for blobs, files, queues, and tables. Recommended for most scenarios using Azure Storage.

Incorrect Answers:

- BlockBlobStorage accounts: Storage accounts with premium performance characteristics for block blobs and append blobs. Recommended for scenarios with high transactions rates, or scenarios that use smaller objects or require consistently low storage latency.
- General-purpose v1 accounts: Legacy account type for blobs, files, queues, and tables. Use general-purpose v2 accounts instead when possible.

Replication: Geo-redundant Storage

Scenario: Data must be replicated to a secondary region and three availability zones.

Geo-redundant storage (GRS) copies your data synchronously three times within a single physical location in the primary region using LRS. It then copies your data asynchronously to a single physical location in the secondary region.

Incorrect Answers:
Geo-zone-redundant storage (GZRS), but it would be more costly.

Access tier: Cool -
Data storage costs must be minimized.
Note: Azure storage offers different access tiers, which allow you to store blob object data in the most cost-effective manner. The available access tiers include:
Hot - Optimized for storing data that is accessed frequently.
Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-account-overview https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal

Paolo93 Highly Voted 2 years, 3 months ago
Replication should be GZRS.
Requirements ask "Data must be replicated to a secondary region and three availability zones". GRS option doesn't copy data in different availability zones
upvoted 35 times

gunz123 2 years, 3 months ago
Geo-redundant storage (GRS) copies your data synchronously three times within a single physical location in the primary region using LRS. It then copies your data asynchronously to a single physical location in the secondary region. Within the secondary region, your data is copied synchronously three times using LRS.

Geo-zone-redundant storage (GZRS) copies your data synchronously across three Azure availability zones in the primary region using ZRS. It then copies your data asynchronously to a single physical location in the secondary region. Within the secondary region, your data is copied synchronously three times using LRS.
upvoted 3 times

gunz123 2 years, 3 months ago
GZRS is correct answer

https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy
upvoted 5 times

MrXBasit 2 years, 3 months ago
I also agree the answer should be gpv2, gzrs and cool
upvoted 2 times

Fengtong 2 years, 2 months ago
Why is not RA-GZRS? Isn't it much cheaper than GZRS?
upvoted 2 times

vokep77043 2 years, 2 months ago
No, RA is additional option which costs MORE.
upvoted 3 times

jkes80 Highly Voted 2 years, 3 months ago
What about cool vs hot in the last box? Shouldn't this be hot, because the site gets accessed frequently?
upvoted 12 times

ning 2 years, 2 months ago
That is also my question? Any where mention performance requirements? Only thing to make it cool, is that behind CDN, cached inside CDN maybe?
upvoted 1 times

alperc 1 year, 11 months ago
No, even though its behind CDN, it should still be hot, because the CDN will most likely ask Storage Account several times a day. Cool will only be considered if the read frequency of storage account will be less than 1 per month.
upvoted 5 times

ReniRechner 1 year, 8 months ago
But is only important to minimize storage costs, not storage access costs.
If they had said "minimize total storage costs" I also would have chosen hot.
upvoted 1 times

ENGs Most Recent 3 weeks, 4 days ago
On my exam 2023-10 before the Update of the Exam
upvoted 1 times

kakajos 1 month ago

The right answers are off-course: Storage V2 (general purpose v2), GZRS, Cool

upvoted 2 times

Vipul123 4 months ago

This question was there on 29th June exam. Went with storage v2, GZRS and Cool.

upvoted 2 times

Chris2349 4 months, 2 weeks ago

Got this case study on 15th of June.
Passed with 887
On this one I answered: StorageV2, GZRS and Cool.

upvoted 1 times

Vmwarevirtual 5 months ago

Had this question on 27 May 2023.
I chose StorageV2 (general-purpose v2) , GZRS , Cool tier .

upvoted 1 times

surprise0011 6 months, 2 weeks ago

received 2023-17-04 went with gpv2,GZRS,cool , score 926

upvoted 5 times

mounir2 7 months ago

1. V2
2. GZRS
3. Hot (website should be visit a least a few times a month...)
microsoft questions are poorly worded as always

upvoted 2 times

st0rmtrooperx 10 months, 2 weeks ago

Got this on Dec 16th, 2022. Scored 921 and answered StorageV2, GZRS and Cool.

upvoted 7 times

SivajiTheBoss 1 year, 7 months ago

Correct Answer: StorageV2, GZRS & Cool.

upvoted 5 times

leonidn 1 year, 9 months ago

1. StorageV2 is a recommended kind.
2. GZRS because "Data must be replicated to a secondary region and three availability zones." and are not required for reading access in the secondary region. Due to GZRS is cheaper than RA-GZRS, we select it.
3. There are no specific requirements that lead us to which access tier is better. Cool has an advantage when we care about storage cost and we do not change files. We store SPA. It is not supposed to be huge. But due to we expect that the modern development process may be turned into multiple releases per day, I believe that hot is preferable. And again, because we do not have specific requirements and default is used to be "hot", I select "hot":

upvoted 9 times

lugospod 1 year, 9 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 4 times

Lucario95 1 year, 11 months ago

I agree with StorageV2, GZRS and COOL

upvoted 1 times

Lucario95 1 year, 9 months ago

Thinking back, HOT should be better because of @alperc reply.

upvoted 1 times

phvogel 2 years ago

GZRS to get the data copied both to three zones within the region AND to another region.
General Storage V2 because BlobStorage is more expensive and doesn't support GZRS
Cool storage gives you the cheapest storage but higher data transfer costs. However, the storage account is only accessed to load the CDN. After that, all content is retrieved from the CDN rather than the Storage account so transfer costs aren't as important as storage accounts.

upvoted 1 times

avenger34 1 year, 11 months ago

General Storage V2 supports GZRS
"Only general-purpose v2 storage accounts support GZRS and RA-GZRS"

https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy#geo-zone-redundant-storage

upvoted 3 times

Mev4953 1 year, 9 months ago

No...Look the ref.
https://docs.microsoft.com/en-us/azure/storage/common/storage-account-overview

upvoted 1 times

j888 2 years, 2 months ago

Somehow I reckon the answer is RA-GRS, coz it replicates 3 zones and saves cost.

upvoted 2 times

8boltsi 2 years, 1 month ago

RA-GRS is more expensive compared to GZRS

upvoted 1 times