

[November-2020Real Exam Questions-Braindump2go AZ-303 PDF AZ-303 197Q Download][Q173-Q193

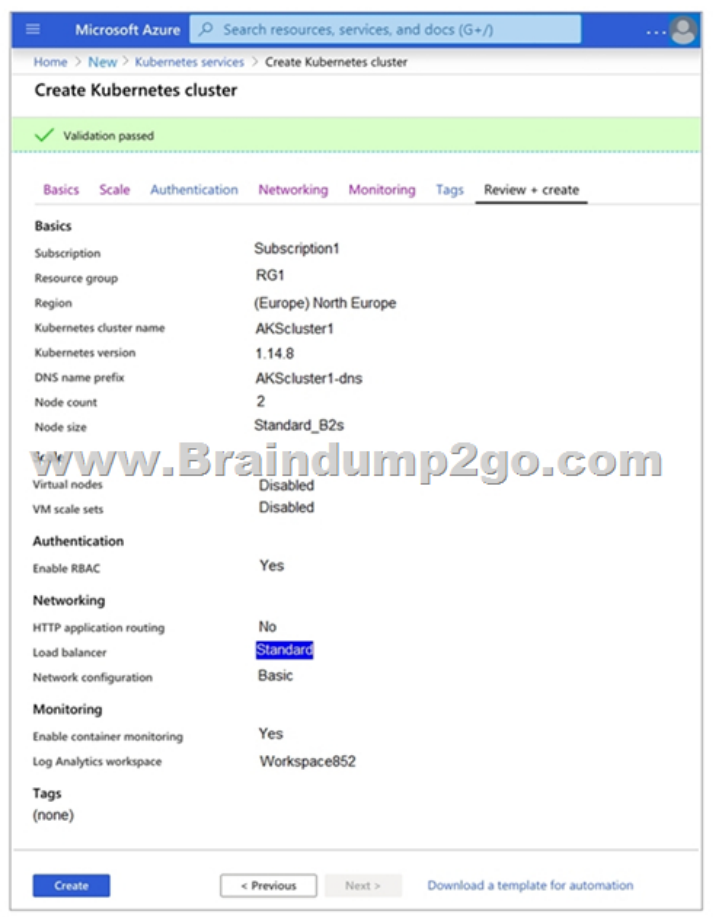
November/2020 Latest Braindump2go AZ-303 Exam Dumps with PDF and VCE Free Updated Today! Following are some new AZ-303 Real Exam Questions!
QUESTION 173You have an Azure Cosmos DB account named Account1. Account1 includes a database named DB1 that contains a container named Container1. The partition key for Container1 is set to /city.You plan to change the partition key for Container1.What should you do first?
A. Delete Container1.
B. Create a new Azure Cosmos DB account.
C. Implement the Azure Cosmos DB.NET.SDK.
D. Regenerate the keys for Account1.
Answer: B
Explanation:The Change Feed Processor and Bulk Executor Library, in Azure Cosmos DB can be leveraged to achieve a live migration of your data from one container to another. This allows you to re-distribute your data to match the desired new partition key scheme, and make the relevant application changes afterwards, thus achieving the effect of "updating your partition key".
Incorrect Answers:A: It is not possible to "update" your partition key in an existing container.
Reference:

<https://devblogs.microsoft.com/cosmosdb/how-to-change-your-partition-key/>
QUESTION 174Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You have an app named App1 that uses data from two on-premises Microsoft SQL Server databases named DB1 and DB2.You plan to move DB1 and DB2 to Azure.You need to implement Azure services to host DB1 and DB2. The solution must support server-side transactions across DB1 and DB2.
Solution: You deploy DB1 and DB2 as Azure SQL databases each on a different Azure SQL Database server.Does this meet the goal?
A. Yes
B. No
Answer: B
Explanation:Instead deploy DB1 and DB2 to SQL Server on an Azure virtual machine.
Note: Understanding distributed transactions.When both the database management system and client are under the same ownership (e.g. when SQL Server is deployed to a virtual machine), transactions are available and the lock duration can be controlled.
Reference:<https://docs.particular.net/nservicebus/azure/understanding-transactionality-in-azure>

QUESTION 175
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You have an app named App1 that uses data from two on-premises Microsoft SQL Server databases named DB1 and DB2.You plan to move DB1 and DB2 to Azure.You need to implement Azure services to host DB1 and DB2. The solution must support server-side transactions across DB1 and DB2.
Solution: You deploy DB1 and DB2 as Azure SQL databases on the same Azure SQL Database server.Does this meet the goal?
A. Yes
B. No
Answer: B
Explanation:Instead deploy DB1 and DB2 to SQL Server on an Azure virtual machine.
Note: Understanding distributed transactions.When both the database management system and client are under the same ownership (e.g. when SQL Server is deployed to a virtual machine), transactions are available and the lock duration can be controlled.
Reference:

<https://docs.particular.net/nservicebus/azure/understanding-transactionality-in-azure>
QUESTION 176Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You have a server named Server1 that runs Windows Server 2019. Server1 is a container host.You are creating a Dockerfile to build a container image.You need to add a file named File1.txt from Server1 to a folder named C:Folder1 in the container image.
Solution: You add the following line to the Dockerfile.Copy-Item File1.txt C:Folder1File1.txtYou then build the container image.
Does this meet the goal?
A. Yes
B. No
Answer: B
Explanation:Copy is the correct command to copy a file to the container image.
Reference:https://docs.docker.com/develop/develop-images/dockerfile_best-practices/#add-or-copy

<https://docs.docker.com/engine/reference/builder/>
QUESTION 177You create an Azure Kubernetes Service (AKS) cluster configured as shown in the exhibit. (Click the Exhibit tab.)



You deploy a containerized application named App1 to the agentPool node pool. You need to create a containerized application named App2 that runs on four nodes of size DS3 v2. What should you do first? A. Upgrade the AKS cluster. B. Create a new node pool. C. Modify the autoscaling settings for the agentPool node. D. Enable virtual nodes for the AKS cluster. Answer: B
Explanation: Changing the agent size is not allowed. In the future Microsoft plans to support multiple node pools wherein you can create different pools with different VM sizes. Reference: <https://github.com/Azure/AKS/issues/132> QUESTION 178 You create an Azure Kubernetes Service (AKS) cluster and an Azure Container Registry. You need to perform continuous deployments of a containerized application to the AKS cluster as soon as the image updates in the registry. What should you use to perform the deployments? A. an Azure Automation runbook B. a kubectl script from a CRON job C. an Azure Resource Manager template D. an Azure Pipelines release pipeline Answer: D Explanation: You can implement a Continuous Deployment pipeline. Example:

Pipeline Az
Recent Char
Stage View
WWW
Average st
(Average full run t
420
Jan 23 No
18:00 Changes

What the pipeline accomplishes : Stage 1: The code gets pushed in the Github. The Jenkins job gets triggered automatically. The Dockerfile is checked out from Github. Stage 2: Docker builds an image from the Dockerfile and then the image is tagged with the build number. Additionally, the latest tag is also attached to the image for the containers to use. Stage 3: We have default deployment and service YAML files stored on the Jenkins server. Jenkins makes a copy of the default YAML files, make the necessary changes according to the build and put them in a separate folder. Stage 4: kubectl was initially configured at the time of setting up AKS on the

Jenkins server. The YAML files are fed to the kubectl util which in turn creates pods and services. Reference:

<https://medium.com/velotio-perspectives/continuous-deployment-with-azure-kubernetes-service-azure-container-registry-jenkins-ca337940151b>

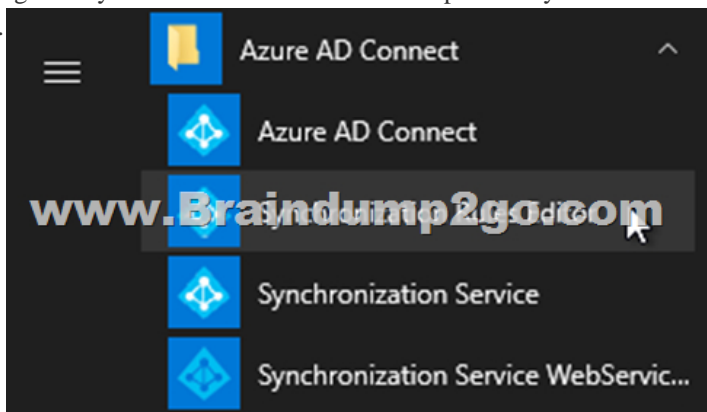
QUESTION 179 You have an Azure web app that runs in a Premium App Service plan. Developers plan to update the app weekly. You need to ensure that the app can be switched from the current version to the new version. The solution must meet the following requirements:- Provide the developers with the ability to test the app in Azure prior to switching versions.- Testing must use the same app instance.- Ensure that the app version can be rolled back.- Minimize downtime. What should you do? A. Create a deployment slot. B. Copy the App Service plan. C. Add an instance of the app to the scale set. D. Create an Azure Active Directory (Azure AD) enterprise application. **Answer: A** Explanation: Azure Functions deployment slots allow your function app to run different instances called "slots". Slots are different environments exposed via a publicly available endpoint. One app instance is always mapped to the production slot, and you can swap instances assigned to a slot on demand. There are a number of advantages to using deployment slots. The following scenarios describe common uses for slots: Different environments for different purposes: Using different slots gives you the opportunity to differentiate app instances before swapping to production or a staging slot. Easy fallbacks: After a swap with production, the slot with a previously staged app now has the previous production app. If the changes swapped into the production slot aren't as you expect, you can immediately reverse the swap to get your "last known good instance" back. Pre-warming Reference: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-deployment-slots>

QUESTION 180

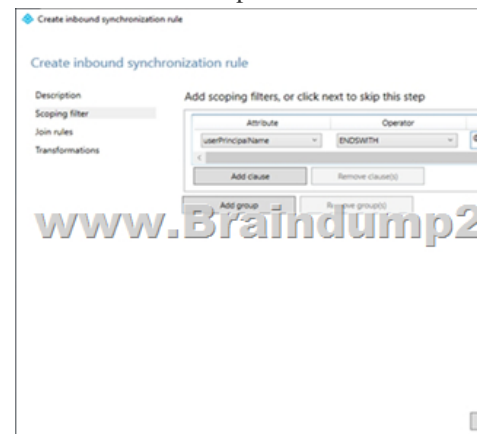
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer 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 manage an Active Directory domain named contoso.local. You install Azure AD Connect and connect to an Azure Active Directory (Azure AD) tenant named contoso.com without syncing any accounts. You need to ensure that only users who have a UPN suffix of contoso.com in the contoso.local domain sync to Azure AD. **Solution:** You use Azure AD Connect to customize the synchronization options. Does this meet the goal? A. Yes B. No **Answer: B** Explanation: Instead use Synchronization Rules Editor to create a synchronization rule. **Note:** Filtering what objects are synced to Azure AD is a common request and there are many instances where filtering by OU just doesn't cut it. One option is to filter users by their UPN suffix so that only users with the public FQDN as their UPN suffix are synced to Azure AD (e.g., john.doe@acme.com would be synced while jane.doe@internal.acme.com would not). Filtering can be configured using either the GUI (Synchronization Rules Editor) or PowerShell. Reference: <https://www.sidekicktech.com/blog/field-notes/2019/upn-suffix-filtering-ad-connect/>

QUESTION 181

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 manage an Active Directory domain named contoso.local. You install Azure AD Connect and connect to an Azure Active Directory (Azure AD) tenant named contoso.com without syncing any accounts. You need to ensure that only users who have a UPN suffix of contoso.com in the contoso.local domain sync to Azure AD. **Solution:** You use Synchronization Rules Editor to create a synchronization rule. Does this meet the goal? A. Yes B. No **Answer: A** Explanation: Filtering what objects are synced to Azure AD is a common request and there are many instances where filtering by OU just doesn't cut it. One option is to filter users by their UPN suffix so that only users with the public FQDN as their UPN suffix are synced to Azure AD (e.g., john.doe@acme.com would be synced while jane.doe@internal.acme.com would not). Filtering can be configured using either the GUI or PowerShell. Through GUI: Using The Synchronization Rules Editor 1. Open the Synchronization Rules Editor on the server where Azure AD Connect is installed.



2. Click the Add new rule button on the View and manage your synchronization rules window. 3. Fill out the appropriate fields on the Description tab and click Next >. 4. On the Scoping filter tab, click Add group, then Add clause, add a userPrincipalName attribute filter, and click Next >. Attribute: userPrincipalName Operator: ENDSWITH Value: Your internal UPN suffix prefixed with @ (e.g., @internal.acme.com). Users with this UPN suffix will NOT be synced with Office 365.



Reference: <https://www.sidekicktech.com/blog/field-notes/2019/upn-suffix-filtering-ad-connect/> QUESTION 182 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 manage an Active Directory domain named contoso.local. You install Azure AD Connect and connect to an Azure Active Directory (Azure AD) tenant named contoso.com without syncing any accounts. You need to ensure that only users who have a UPN suffix of contoso.com in the contoso.local domain sync to Azure AD. Solution: You use the Synchronization Service Manager to modify the Active Directory Domain Services (AD DS) Connector. Does this meet the goal? A. Yes B. No Answer: B Explanation: Instead use Synchronization Rules Editor to create a synchronization rule. Note: Filtering what objects are synced to Azure AD is a common request and there are many instances where filtering by OU just doesn't cut it. One option is to filter users by their UPN suffix so that only users with the public FQDN as their UPN suffix are synced to Azure AD (e.g., john.doe@acme.com would be synced while jane.doe@internal.acme.com would not). Filtering can be configured using either the GUI (Synchronization Rules Editor) or PowerShell. Reference: <https://www.sidekicktech.com/blog/field-notes/2019/upn-suffix-filtering-ad-connect/>

QUESTION 183 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an app named App1 that uses data from two on-premises Microsoft SQL Server databases named DB1 and DB2. You plan to move DB1 and DB2 to Azure. You need to implement Azure services to host DB1 and DB2. The solution must support server-side transactions across DB1 and DB2. Solution: You deploy DB1 and DB2 to SQL Server on an Azure virtual machine. Does this meet the goal? A. Yes B. No Answer: A Explanation: Understanding distributed transactions. When both the database management system and client are under the same ownership (e.g. when SQL Server is deployed to a virtual machine), transactions are available and the lock duration can be controlled. Reference: <https://docs.particular.net/nservicebus/azure/understanding-transactionality-in-azure>

QUESTION 184 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an Azure Cosmos DB database that contains a container named Container1. The partition key for Container1 is set to /day. Container1 contains the items shown in the following table.

Name	Content
Item1	<pre>{ "id": "1", "day": "Mon", "value": "10" }</pre>
Item2	<pre>{ "id": "2", "day": "Mon", "value": "15" }</pre>
Item3	<pre>{ "id": "3", "day": "Tue", "value": "10" }</pre>
Item4	<pre>{ "id": "4", "day": "Wed", "value": "15" }</pre>

You need to programmatically query Azure Cosmos DB and retrieve Item1 and Item2 only. Solution: You run the following query. `SELECT id FROM c WHERE c.day = "Mon" OR c.day = "Tue"` You set the `EnableCrossPartitionQuery` property to `False`. Does this meet the goal? A. Yes B. No Answer: B Explanation: Returns Item1 only as `EnableCrossPartitionQuery` property to `False`. If `EnableCrossPartitionQuery` property is set to `true`, it will return Item1, Item2, and Item3. Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql-query-where>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.documents.client.feedoptions.enablecrosspartitionquery?view=azure-dotnet>

QUESTION 185 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 manage an Active Directory domain named `contoso.local`. You install Azure AD Connect and connect to an Azure Active Directory (Azure AD) tenant named `contoso.com` without syncing any accounts. You need to ensure that only users who have a UPN suffix of `contoso.com` in the `contoso.local` domain sync to Azure AD. Solution: You use the Synchronization Service Manager to modify the Metaverse Designer tab. Does this meet the goal? A. Yes B. No Answer: B Explanation: Instead use Synchronization Rules Editor to create a synchronization rule. Note: Filtering what objects are synced to Azure AD is a common request and there are many instances where filtering by OU just doesn't cut it. One option is to filter users by their UPN suffix so that only users with the public FQDN as their UPN suffix are synced to Azure AD (e.g., `john.doe@acme.com` would be synced while `jane.doe@internal.acme.com` would not). Filtering can be configured using either the GUI (Synchronization Rules Editor) or PowerShell. Reference: <https://www.sidekicktech.com/blog/field-notes/2019/upn-suffix-filtering-ad-connect/>

QUESTION 186 Your network contains an on-premises Active Directory and an Azure Active Directory (Azure AD) tenant. You deploy Azure AD Connect and configure pass-through authentication. Your Azure subscription contains several web apps that are accessed from the Internet. You plan to use Azure Multi-Factor Authentication (MFA) with the Azure Active Directory tenant. You need to recommend a solution to prevent users from being prompted for Azure MFA when they access the web apps from the on-premises network. What should you include in the recommendation? A. an Azure policy B. trusted IPs C. a site-to-site VPN between the on-premises network and Azure D. an Azure ExpressRoute circuit Answer: B Explanation: The Trusted IPs feature of Azure Multi-Factor Authentication is used by administrators of a managed or federated tenant. The feature bypasses two-step verification for users who sign in from the company intranet. The feature is available with the full version of Azure Multi-Factor Authentication, and not the free version for administrators. Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-mfasettings#trusted-ips> QUESTION 187 You have the following Azure Active Directory (Azure AD) tenants:- Contoso.onmicrosoft.com: Linked to a Microsoft Office 365 tenant and syncs to an Active Directory forest named contoso.com by using password hash synchronization- Contosoazure.onmicrosoft.com: Linked to an Azure subscription named Subscription1 You need to ensure that you can assign the users in contoso.com access to the resources in Subscription1. What should you do? A. Configure contoso.onmicrosoft.com to use pass-through authentication. B. Create guest accounts for all the contoso.com users in contosoazure.onmicrosoft.com. C. Deploy a second Azure AD Connect server and sync contoso.com to contosoazure.onmicrosoft.com. D. Configure Active Directory Federation Services (AD FS) federation between contosoazure.onmicrosoft.com and contoso.com. Answer: C Explanation: Azure AD Connect allows you to quickly onboard to Azure AD and Office 365. Note: The most common topology is a single on-premises forest, with one or multiple domains, and a single Azure AD tenant. For Azure AD authentication, password hash synchronization is used. The express installation of Azure AD Connect supports only this topology. Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/hybrid/plan-connect-topologies> QUESTION 188 You have several Azure web apps that use access keys to access databases. You plan to migrate the access keys to Azure Key Vault. Each app must authenticate by using Azure Active Directory (Azure AD) to gain access to the access keys. What should you create in Azure to ensure that the apps can access the access keys? A. managed identities B. managed applications C. Azure policies D. an App Service plan Answer: A Explanation: Azure Key Vault provides a way to securely store credentials and other secrets, but your code needs to authenticate to Key Vault to retrieve them. Managed identities for Azure resources overview helps to solve this problem by giving Azure services an automatically managed identity in 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. Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/general/tutorial-net-create-vault-azure-web-app> QUESTION 189 You have an Azure key vault named KV1. You need to implement a process that will digitally sign the blobs stored in Azure Storage. What is required in KV1 to sign the blobs? A. a key B. a secret C. a certificate Answer: B Explanation: Use an Azure key vault secret to key of your blob storage account container. Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/general/integrate-databricks-blob-storage> QUESTION 190 You set the multi-factor authentication status for a user named admin1@contoso.com to Enabled. Admin1 accesses the Azure portal by using a web browser. Which additional security verifications can Admin1 use when accessing the Azure portal? A. a phone call, an email message that contains a verification code, and a text message that contains an app password. B. an app password, a text message that contains a verification code, and a verification code sent from the Microsoft Authenticator app. C. an app password, a text message that contains a verification code, and a notification sent from the Microsoft Authenticator app. D. a phone call, a text message that contains a verification code, and a notification or a verification code sent from the Microsoft Authenticator app. Answer: D Explanation: The Microsoft Authenticator app can help prevent unauthorized access to accounts and stop fraudulent transactions by pushing a notification to your smartphone or tablet. Users view the notification, and if it's legitimate, select Verify. Otherwise, they can select Deny. Reference: <https://docs.microsoft.com/en-us/azure/active-directory/authentication/concept-authentication-methods>

QUESTION 191 You have an Azure Cosmos DB account named Account1. Account1 includes a database named DB1 that contains a container named Container1. The partition key for Container1 is set to /city. You plan to change the partition key for Container1. What should you do first? A. Delete Container1. B. Create a new container in DB1. C. Implement the Azure Cosmos DB.NET.SDK. D. Regenerate the keys for Account1. Answer: B Explanation: The Change Feed Processor and Bulk Executor Library, in Azure Cosmos DB can be leveraged to achieve a live migration of your data from one container to another. This allows you to re-distribute your data to match the desired new partition key scheme, and make the relevant application changes afterwards, thus achieving the effect of "updating your partition key". Incorrect Answers: A: It is not possible to "update" your partition key in an existing container. Reference: <https://devblogs.microsoft.com/cosmosdb/how-to-change-your-partition-key/>

QUESTION 192 You plan to automate the deployment of a virtual machine scale set that uses the Windows Server 2016 Datacenter image. You need to ensure that when the scale set virtual machines are provisioned, they have web server components installed. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point. A. Upload a configuration script. B. Create an Azure policy. C. Modify the extensionProfile section of the Azure Resource Manager template. D. Create a new virtual machine scale set in the Azure portal. E. Create an automation account. Answer: C D Explanation:

<https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/tutorial-install-apps-template> QUESTION 193 Hotspot Question You plan to deploy five virtual machines to a virtual network subnet. Each virtual machine will have a public IP address and a private IP address. Each virtual machine requires the same inbound and outbound security rules. What is the minimum number of network interfaces and network security groups that you require? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Minimum number of network interfaces:

▼
5
10
15
20

Minimum number of network security groups:

▼
1
2
5
10

Answer:

Answer Area

Minimum number of network interfaces:

▼
5
10
15
20

Minimum number of network security groups:

▼
1
2
5
10

Explanation: Box 1: 5 We have five virtual machines. Each virtual machine will have a public IP address and a private IP address. Each will require a network interface. Box 2: 1 Each virtual machine requires the same inbound and outbound security rules. We can add them to one group. Reference:

<https://blogs.msdn.microsoft.com/igorpag/2016/05/14/azure-network-security-groups-nsg-best-practices-and-lessons-learned/>
<https://docs.microsoft.com/en-us/azure/virtual-network/security-overview> Resources From: 1. 2020 Latest Braindump2go AZ-303 Exam Dumps (PDF & VCE) Free Share: <https://www.braindump2go.com/az-303.html> 2. 2020 Latest Braindump2go AZ-303 PDF and VCE Dumps Free Share: <https://drive.google.com/drive/folders/114-Ncx3vdm9Ra2pN5d9Lnjv3pxbJpxZB?usp=sharing> 3. 2020 Free Braindump2go AZ-303 PDF Download: [https://www.braindump2go.com/free-online-pdf/AZ-303-PDF\(190-197\).pdf](https://www.braindump2go.com/free-online-pdf/AZ-303-PDF(190-197).pdf)
[https://www.braindump2go.com/free-online-pdf/AZ-303-PDF-Dumps\(168-178\).pdf](https://www.braindump2go.com/free-online-pdf/AZ-303-PDF-Dumps(168-178).pdf)
[https://www.braindump2go.com/free-online-pdf/AZ-303-VCE-Dumps\(179-189\).pdf](https://www.braindump2go.com/free-online-pdf/AZ-303-VCE-Dumps(179-189).pdf) Free Resources from Braindump2go, We Devoted to Helping You 100% Pass All Exams!