

➤ **Vendor: Microsoft**

➤ **Exam Code: AZ-303**

➤ **Exam Name: Microsoft Azure Architect Technologies**

➤ **New Updated Questions from [Braindump2go](#) (Updated in [Nov./2020](#))**

[Visit Braindump2go and Download Full Version AZ-303 Exam Dumps](#)

QUESTION 168

You have an Active Directory forest named contoso.com.

You install and configure Azure AD Connect to use password hash synchronization as the single sign-on (SSO) method. Staging mode is enabled.

You review the synchronization results and discover that the Synchronization Service Manager does not display any sync jobs.

You need to ensure that the synchronization completes successfully.

What should you do?

- A. Run Azure AD Connect and disable staging mode.
- B. From Synchronization Service Manager, run a full import.
- C. Run Azure AD Connect and set the SSO method to Pass-through Authentication.
- D. From Azure PowerShell, run Start-AdSyncSyncCycle -PolicyType Initial.

Answer: A

Explanation:

In staging mode, the server is active for import and synchronization, but it does not run any exports. A server in staging mode is not running password sync or password writeback, even if you selected these features during installation.

When you disable staging mode, the server starts exporting, enables password sync, and enables password writeback.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-sync-staging-server>

<https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-sync-operations>

QUESTION 169

Your on-premises network contains 100 virtual machines that run Windows Server 2019.

You have an Azure subscription that contains an Azure Log Analytics workspace named Workspace1.

You need to collect errors from the Windows event logs on the virtual machines.

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

NOTE: Each correct selection is worth one point.

- A. Create an Azure Event Grid domain.
- B. Deploy the Microsoft Monitoring Agent.
- C. Configure Windows Event Forwarding on the virtual machines.
- D. Create an Azure Sentinel workspace.
- E. Configure the Data Collection settings for Workspace1.

Answer: BE

Explanation:

The Azure Log Analytics agent collects telemetry from Windows and Linux virtual machines in any cloud, on-premises machines, and those monitored by System Center Operations Manager and sends it collected data to your Log Analytics workspace in Azure Monitor.

[AZ-303 Exam Dumps](#) [AZ-303 Exam Questions](#) [AZ-303 PDF Dumps](#) [AZ-303 VCE Dumps](#)

<https://www.braindump2go.com/az-303.html>

Note: You may also see the Log Analytics agent referred to as the Microsoft Monitoring Agent (MMA) or OMS Linux agent.

Data is collected using the Log Analytics agent, which reads various security-related configurations and event logs from the machine and copies the data to your workspace for analysis.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/log-analytics-agent>

<https://docs.microsoft.com/en-us/azure/security-center/security-center-enable-data-collection>

QUESTION 170

You have an Azure subscription named Subscription1.

You deploy a Linux virtual machine named VM1 to Subscription1.

You need to monitor the metrics and the logs of VM1.

What should you use?

- A. Azure HDInsight
- B. Azure Analysis Services
- C. Linux Diagnostic Extension (LAD) 3.0
- D. the AzurePerformanceDiagnostics extension

Answer: D

Explanation:

You can use extensions to configure diagnostics on your VMs to collect additional metric data.

The basic host metrics are available, but to see more granular and VM-specific metrics, you need to install the Azure diagnostics extension on the VM. The Azure diagnostics extension allows additional monitoring and diagnostics data to be retrieved from the VM.

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-machines/linux/tutorial-monitoring>

QUESTION 171

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.

What should you do?

- A. Use the Synchronization Service Manager to modify the Metaverse Designer tab.
- B. Use Azure AD Connect to customize the synchronization options.
- C. Use the Synchronization Rules Editor to create a synchronization rule.
- D. Use Synchronization Service Manager to modify the Active Directory Domain Services (AD DS) Connector.

Answer: C

Explanation:

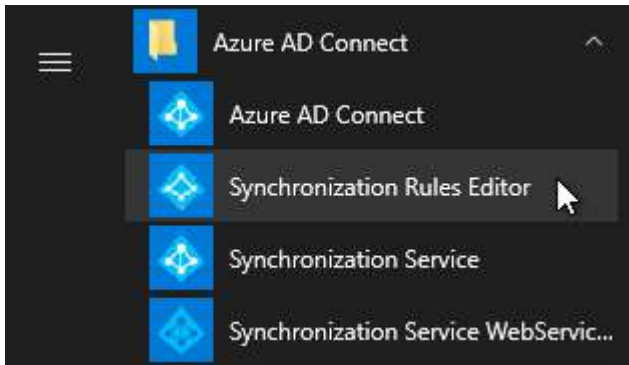
Filtering what objects are synced to Azure AD is a common request and there are many instances where filtering by OU just doesn't cut it. One option is to filter users by their UPN suffix so that only users with the public FQDN as their UPN suffix are synced to Azure AD (e.g., john.doe@acme.com would 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.

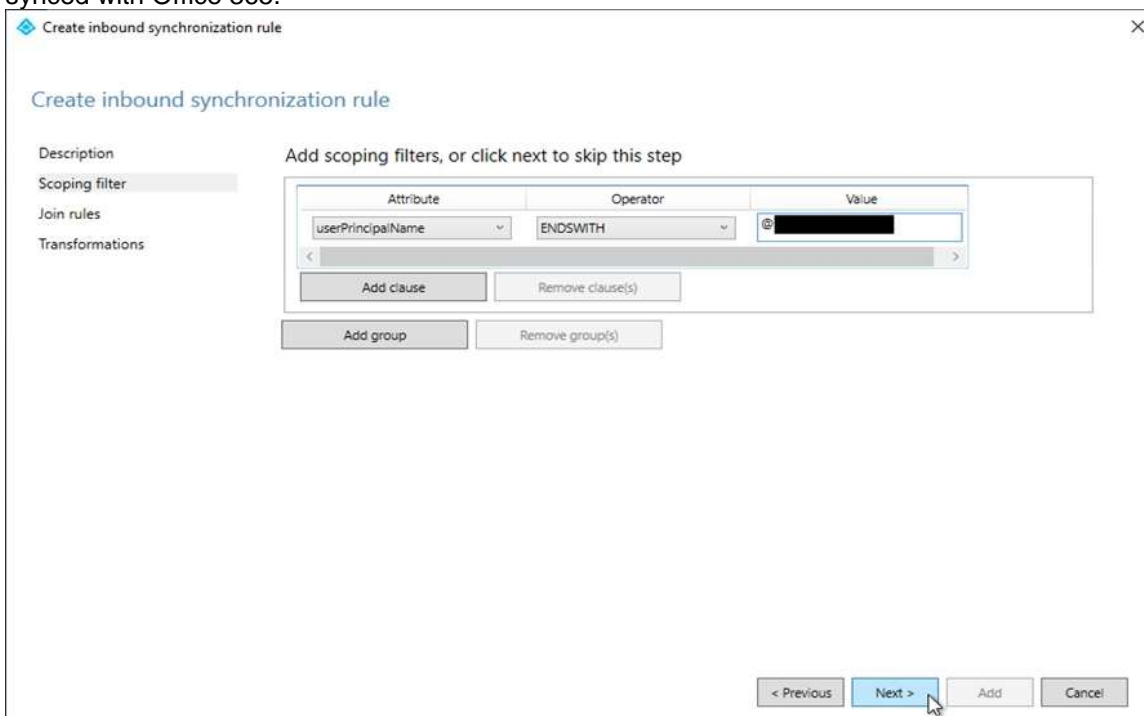
3. Fill out the appropriate fields on the Description tab and click Next >.

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 172

You have an Azure SQL database named DB1.

You plan to create the following four tables in DB1 by using the following code.

Table1.

```
CREATE TABLE Table1
(
    StudentId INT IDENTITY PRIMARY KEY,
    PersonId INT REFERENCES Table4 (PersonId),
    Email NVARCHAR(256)
}
```

Table2.

```
CREATE TABLE Table2
(
    StudentId INT REFERENCES Table1 (StudentId),
    CourseId INT REFERENCES Table3 (CourseId),
    Grade DECIMAL(5,2) CHECK (Grade <= 100.00),
    Attempt TINYINT
)
```

Table3.

```
CREATE TABLE Table3
(
    CourseId INT IDENTITY PRIMARY KEY,
    Name NVARCHAR(56) NOT NULL,
    Teacher NVARCHAR(256) NOT NULL
)
```

Table4.

```
CREATE TABLE Table4
(
    PersonId INT IDENTITY PRIMARY KEY,
    FirstName NVARCHAR(128) NOT NULL,
    MiddleInitial NVARCHAR(10),
    LastName NVARCHAR(128) NOT NULL,
    DateOfBirth DATE NOT NULL
)
```

You need to identify which table must be created last.
What should you identify?

- A. Table1
- B. Table2
- C. Table3
- D. Table4

Answer: B

Explanation:

Table1 references Table4. Therefore Table4 must be created before Table1.

Table2 references Table1 and Table3. Therefore Table1 and Table3 must be created before Table2.

Note: FOREIGN KEY REFERENCES is a constraint that provides referential integrity for the data in the column or columns. FOREIGN KEY constraints require that each value in the column exists in the corresponding referenced column or columns in the referenced table. FOREIGN KEY constraints can reference only columns that are PRIMARY KEY or UNIQUE constraints in the referenced table or columns referenced in a UNIQUE INDEX on the referenced table.

Incorrect Answers:

A: Table1 is referenced by Table2 and should be crated before Table2.

C: Table3 is referenced by Table2 and should be crated before Table2.

D: Table4 is referenced by Table1 and should be crated before Table1.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-transact-sql?view=sql-server-ver15>

QUESTION 173

You have an Azure Cosmos DB account named Account1. Account1 includes a database named DB1 that contains a

[AZ-303 Exam Dumps](#) [AZ-303 Exam Questions](#) [AZ-303 PDF Dumps](#) [AZ-303 VCE Dumps](#)

<https://www.braindump2go.com/az-303.html>

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 174

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 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.

[AZ-303 Exam Dumps](#) [AZ-303 Exam Questions](#) [AZ-303 PDF Dumps](#) [AZ-303 VCE Dumps](#)

<https://www.braindump2go.com/az-303.html>

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 176

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 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:\Folder1\File1.txt
```

You 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.

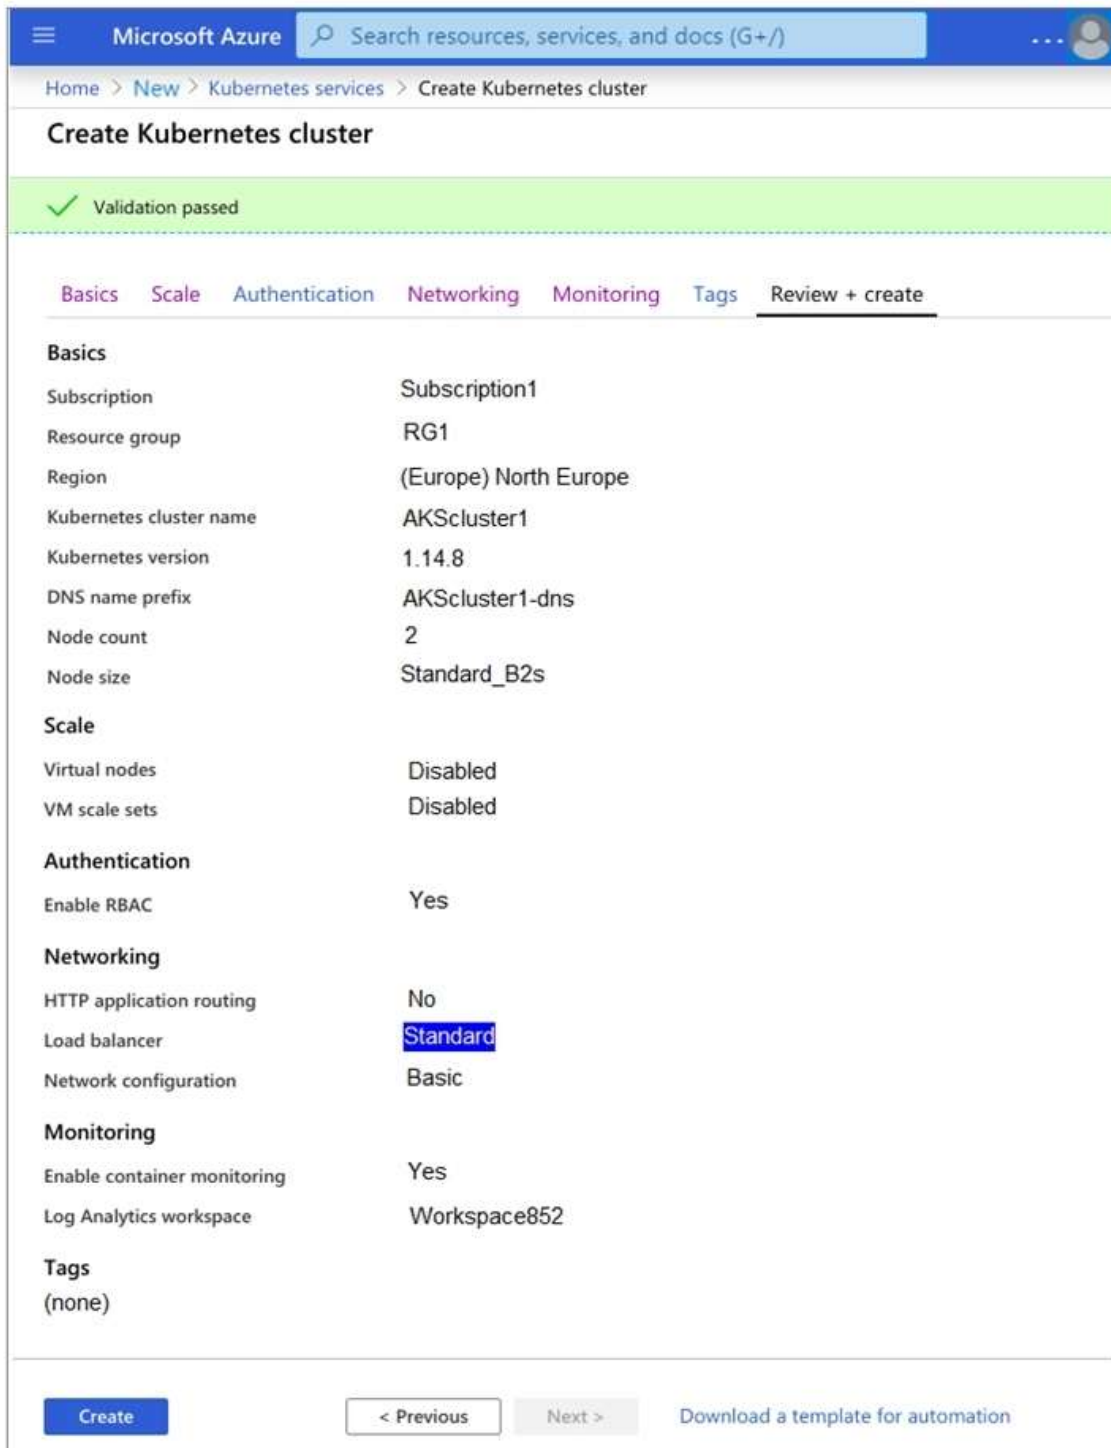
Reference:

https://docs.docker.com/develop/develop-images/dockerfile_best-practices/#add-or-copy

<https://docs.docker.com/engine/reference/builder/>

QUESTION 177

You create an Azure Kubernetes Service (AKS) cluster configured as shown in the exhibit. (Click the Exhibit tab.)



Microsoft Azure Search resources, services, and docs (G+)

Home > New > Kubernetes services > Create Kubernetes cluster

Create Kubernetes cluster

✓ Validation passed

Basics Scale Authentication Networking Monitoring Tags **Review + create**

Basics	
Subscription	Subscription1
Resource group	RG1
Region	(Europe) North Europe
Kubernetes cluster name	AKScluster1
Kubernetes version	1.14.8
DNS name prefix	AKScluster1-dns
Node count	2
Node size	Standard_B2s
Scale	
Virtual nodes	Disabled
VM scale sets	Disabled
Authentication	
Enable RBAC	Yes
Networking	
HTTP application routing	No
Load balancer	Standard
Network configuration	Basic
Monitoring	
Enable container monitoring	Yes
Log Analytics workspace	Workspace852
Tags	
(none)	

Create < Previous Next > Download a template for automation

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

[AZ-303 Exam Dumps](#) [AZ-303 Exam Questions](#) [AZ-303 PDF Dumps](#) [AZ-303 VCE Dumps](#)

<https://www.braindump2go.com/az-303.html>

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:

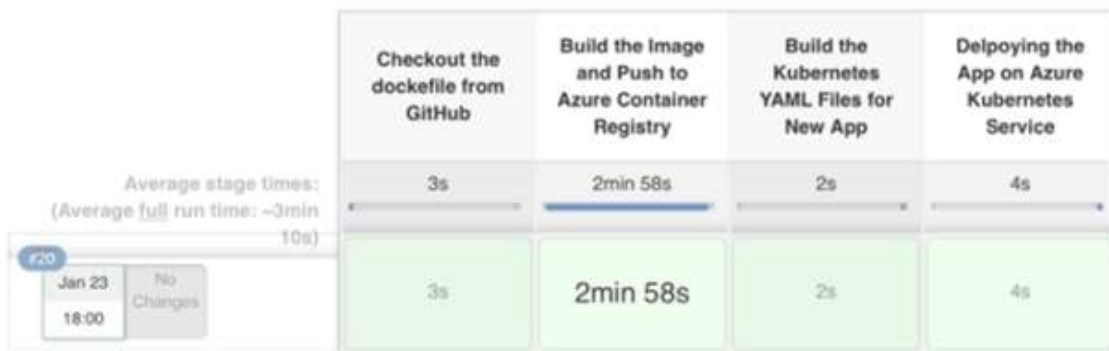
You can implement a Continuous Deployment pipeline.

Example:

Pipeline AzurePipeline



Stage View



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>