

➤ **Vendor: Microsoft**

➤ **Exam Code: AZ-400**

➤ **Exam Name: Implementing and Operating Cisco Data Center Core Technologies (DCCOR)**

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

Visit Braindump2go and Download Full Version AZ-400 Exam Dumps

QUESTION 199

SIMULATION

You need to create and configure an Azure Storage account named az400lod11566895stor in a resource group named RG1lod11566895 to store the boot diagnostics for a virtual machine named VM1.

To complete this task, sign in to the Microsoft Azure portal.

Answer:

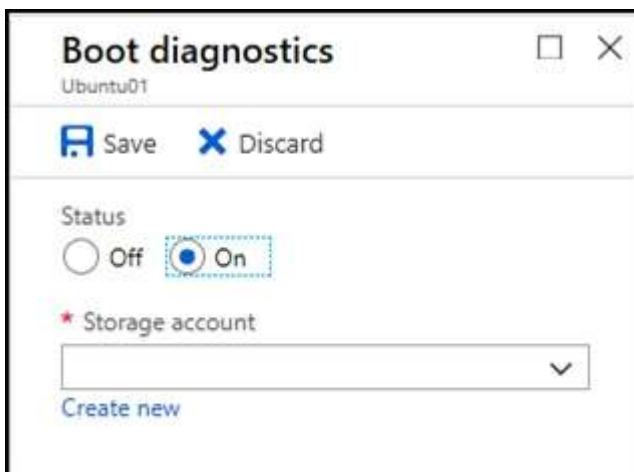
Step 1: To create a general-purpose v2 storage account in the Azure portal, follow these steps:

1. On the Azure portal menu, select All services. In the list of resources, type Storage Accounts. As you begin typing, the list filters based on your input. Select Storage Accounts.
2. On the Storage Accounts window that appears, choose Add.
3. Select the subscription in which to create the storage account.
4. Under the Resource group field, select RG1lod11566895
5. Next, enter a name for your storage account named: az400lod11566895stor
6. Select Create.

Step 2: Enable boot diagnostics on existing virtual machine

To enable Boot diagnostics on an existing virtual machine, follow these steps:

1. Sign in to the Azure portal, and then select the virtual machine VM1.
2. In the Support + troubleshooting section, select Boot diagnostics, then select the Settings tab.
3. In Boot diagnostics settings, change the status to On, and from the Storage account drop-down list, select the storage account az400lod11566895stor.
4. Save the change.



You must restart the virtual machine for the change to take effect.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-create>

<https://docs.microsoft.com/en-us/azure/virtual-machines/troubleshooting/boot-diagnostics>

QUESTION 200**SIMULATION**

You have a web app that connects to an Azure SQL Database named db1.


You need to configure db1 to send Query Store runtime statistics to Azure Log Analytics.

To complete this task, sign in to the Microsoft Azure portal.

Answer:

To enable streaming of diagnostic telemetry for a single or a pooled database, follow these steps:

1. Go to Azure SQL database resource.
2. Select Diagnostics settings.
3. Select Turn on diagnostics if no previous settings exist, or select Edit setting to edit a previous setting. You can create up to three parallel connections to stream diagnostic telemetry.
4. Select Add diagnostic setting to configure parallel streaming of diagnostics data to multiple resources.






5. Enter a setting name for your own reference.

6. Select a destination resource for the streaming diagnostics data: Archive to storage account, Stream to an event hub, or Send to Log Analytics.

7. For the standard, event-based monitoring experience, select the following check boxes for database diagnostics log telemetry: QueryStoreRuntimeStatistics

Diagnostics settings ✕

 Save  Discard  Delete

* Name
service ✓

☐ Archive to a storage account

☐ Stream to an event hub

☒ Send to Log Analytics

Subscription
Workload Insight dev/test subscription ▼

Log Analytics Workspace
sqlanalytics356 (westcentralus) ▼

LOG

☒ SQLInsights

☒ AutomaticTuning

☒ QueryStoreRuntimeStatistics

☒ QueryStoreWaitStatistics

☒ Errors

☒ DatabaseWaitStatistics

☒ Timeouts

☒ Blocks

☒ Deadlocks

METRIC

☒ Basic

8. For an advanced, one-minute-based monitoring experience, select the check box for Basic metrics.

9. Select Save.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-export-configure>

QUESTION 201

[AZ-400 Exam Dumps](#) [AZ-400 Exam Questions](#) [AZ-400 PDF Dumps](#) [AZ-400 VCE Dumps](#)

<https://www.braindump2go.com/az-400.html>

Hotspot Question

Your company uses Git as a source code control system for a complex app named App1.

You plan to add a new functionality to App1.

You need to design a branching model for the new functionality.

Which branch lifetime and branch type should you use in the branching model? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area**Branch lifetime:**

	▼
Long-lived	
Short-lived	

Branch type:

	▼
Master	
Feature	
Integration	

Answer:**Answer Area****Branch lifetime:**

	▼
Long-lived	
Short-lived	

Branch type:

	▼
Master	
Feature	
Integration	

Explanation:

Branch lifetime: Short-lived

Branch type: Feature

Feature branches are used when developing a new feature or enhancement which has the potential of a development lifespan longer than a single deployment. When starting development, the deployment in which this feature will be released may not be known. No matter when the feature branch will be finished, it will always be merged back into the master branch.

References:

<https://gist.github.com/digitaljhelms/4287848>

QUESTION 202

Hotspot Question

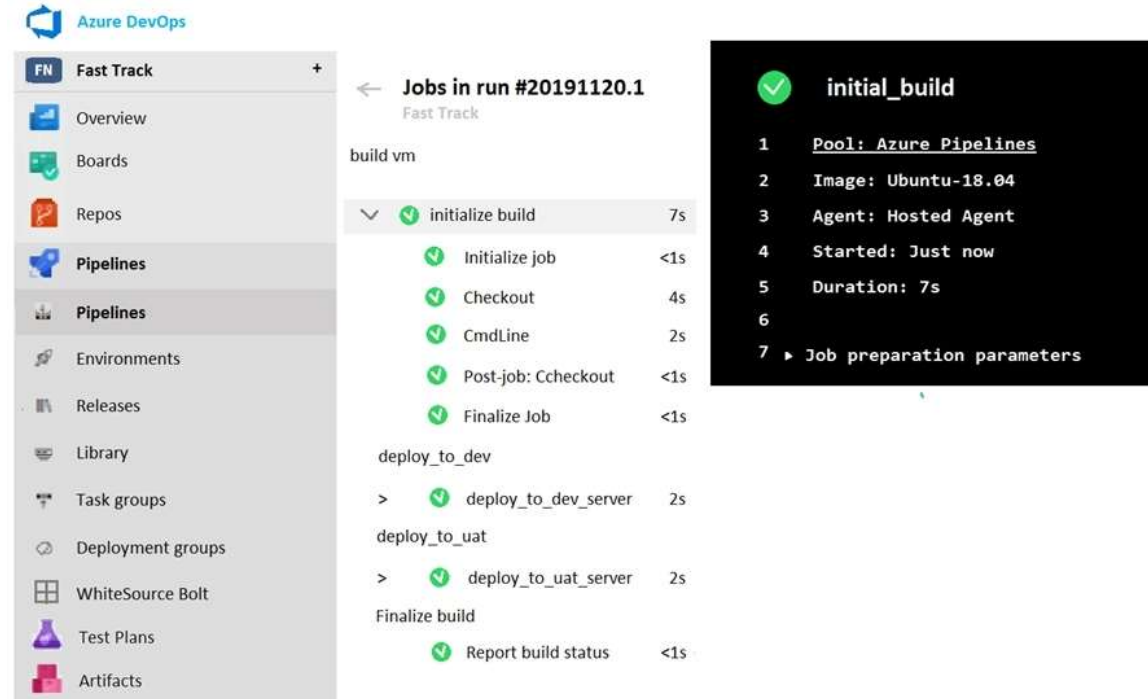
[AZ-400 Exam Dumps](#) [AZ-400 Exam Questions](#) [AZ-400 PDF Dumps](#) [AZ-400 VCE Dumps](#)

<https://www.braindump2go.com/az-400.html>

Your company uses Azure DevOps to deploy infrastructures to Azure.

Pipelines are developed by using YAML.

You execute a pipeline and receive the results in the web portal for Azure Pipelines as shown in the following exhibit.



The screenshot shows the Azure DevOps Pipelines interface. On the left is a navigation pane with options like Overview, Boards, Repos, Pipelines, Environments, Releases, Library, Task groups, Deployment groups, WhiteSource Bolt, Test Plans, and Artifacts. The main area displays 'Jobs in run #20191120.1' for a 'Fast Track' pipeline. It lists stages: 'build_vm' (containing 'initialize build' with sub-jobs: Initialize job, Checkout, CmdLine, Post-job: Ccheckout, Finalize Job), 'deploy_to_dev' (containing 'deploy_to_dev_server'), 'deploy_to_uat' (containing 'deploy_to_uat_server'), and 'Finalize build' (containing 'Report build status'). On the right, a detailed view of the 'initial_build' job is shown, listing parameters: Pool: Azure Pipelines, Image: Ubuntu-18.04, Agent: Hosted Agent, Started: Just now, Duration: 7s, and Job preparation parameters.

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer Area

The pipeline contains

	▼
one stage	
two stages	
three stages	
four stages	
five stages	

Build_vm contains

	▼
one job	
two jobs	
three jobs	
four jobs	
five jobs	

Answer:

Answer Area

The pipeline contains

	▼
one stage	
two stages	
three stages	
four stages	
five stages	

Build_vm contains

	▼
one job	
two jobs	
three jobs	
four jobs	
five jobs	

Explanation:

<https://dev.to/rajikaimal/azure-devops-ci-cd-yaml-pipeline-4glj>

QUESTION 203

Drag and Drop Question

Your company has an Azure subscription named Subscription1. Subscription1 is associated to an Azure Active Directory tenant named contoso.com.

You need to provision an Azure Kubernetes Services (AKS) cluster in Subscription1 and set the permissions for the cluster by using RBAC roles that reference the identities in contoso.com.

Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

Answer Area**Objects**

a system-assigned managed identity

a cluster

an application registration in contoso.com

an RBAC binding

Answer:

Answer Area	
Objects	
	a cluster
	a system-assigned managed identity
	an RBAC binding
an application registration in contoso.com	

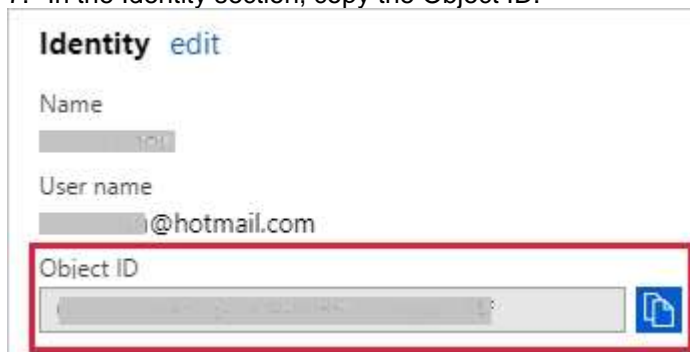
Explanation:

Step 1: Create an AKS cluster

Step 2: a system-assigned managed identity

To create an RBAC binding, you first need to get the Azure AD Object ID.

1. Sign in to the Azure portal.
2. In the search field at the top of the page, enter Azure Active Directory.
3. Click Enter.
4. In the Manage menu, select Users.
5. In the name field, search for your account.
6. In the Name column, select the link to your account.
7. In the Identity section, copy the Object ID.



Identity edit

Name
[redacted]

User name
[redacted]@hotmail.com

Object ID
[redacted]

Step 3: a RBAC binding

Reference:

<https://docs.microsoft.com/en-us/azure/developer/ansible/aks-configure-rbac>

QUESTION 204

Hotspot Question

You manage build and release pipelines by using Azure DevOps. Your entire managed environment resides in Azure. You need to configure a service endpoint for accessing Azure Key Vault secrets. The solution must meet the following requirements:

- Ensure that the secrets are retrieved by Azure DevOps.
- Avoid persisting credentials and tokens in Azure DevOps.

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

NOTE: Each correct selection is worth one point.

Answer Area

Service connection type:

Azure Resource Manager
Generic service
Team Foundation Server / Azure Pipelines service connection

Authentication/authorization method for the connection:

Azure Active Directory OAuth 2.0
Grant authorization
Managed Service Identity Authentication

Answer:

Answer Area

Service connection type:

Azure Resource Manager
Generic service
Team Foundation Server / Azure Pipelines service connection

Authentication/authorization method for the connection:

Azure Active Directory OAuth 2.0
Grant authorization
Managed Service Identity Authentication

Explanation:

Box 1: Azure Pipelines service connection

Box 2: Managed Service Identity Authentication

The managed identities for Azure resources feature in Azure Active Directory (Azure AD) provides Azure services with an automatically managed identity in Azure AD. You can use the identity to authenticate to any service that supports Azure AD authentication, including Key Vault, without any credentials in your code.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/tasks/deploy/azure-key-vault>

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

QUESTION 205

Drag and Drop Question

You are creating a container for an ASP.NET Core app.

You need to create a Dockerfile file to build the image. The solution must ensure that the size of the image is minimized.

How should you configure the file? To answer, drag the appropriate values to the correct targets. Each value must 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.

Answer Area

Values	
dotnet publish -c Release -o out	FROM <input type="text"/> As build-env
dotnet restore	COPY . /app/
microsoft/dotnet:2.2-aspnetcore-runtime	WORKDIR /app
Microsoft/dotnet:2.2-sdk	RUN <input type="text"/>
	FROM <input type="text"/>
	COPY --from=build-env /app/out /app
	WORKDIR /app
	ENTRYPOINT ["dotnet", "MvcMovie.dll"]

Answer:

Values

dotnet publish -c Release -o out

Answer Area

```

FROM Microsoft/dotnet:2.2-sdk As build-env
COPY . /app/
WORKDIR /app
RUN dotnet restore
FROM microsoft/dotnet:2.2-aspnetcore-runtime
COPY --from=build-env /app/out /app
WORKDIR /app
ENTRYPOINT ["dotnet", "MvcMovie.dll"]

```

Explanation:

Box 1: microsoft.com/dotnet/sdk:2.3

The first group of lines declares from which base image we will use to build our container on top of. If the local system does not have this image already, then docker will automatically try and fetch it. The mcr.microsoft.com/dotnet/core/sdk:2.1 comes packaged with the .NET core 2.1 SDK installed, so it's up to the task of building ASP .NET core projects targeting version 2.1

Box 2: dotnet restore

The next instruction changes the working directory in our container to be /app, so all commands following this one execute under this context.

COPY *.csproj /

RUN dotnet restore

Box 3: microsoft.com/dotnet/2.2-aspnetcore-runtime

When building container images, it's good practice to include only the production payload and its dependencies in the container image. We don't want the .NET core SDK included in our final image because we only need the .NET core runtime, so the dockerfile is written to use a temporary container that is packaged with the SDK called build-env to build the app.

Reference:

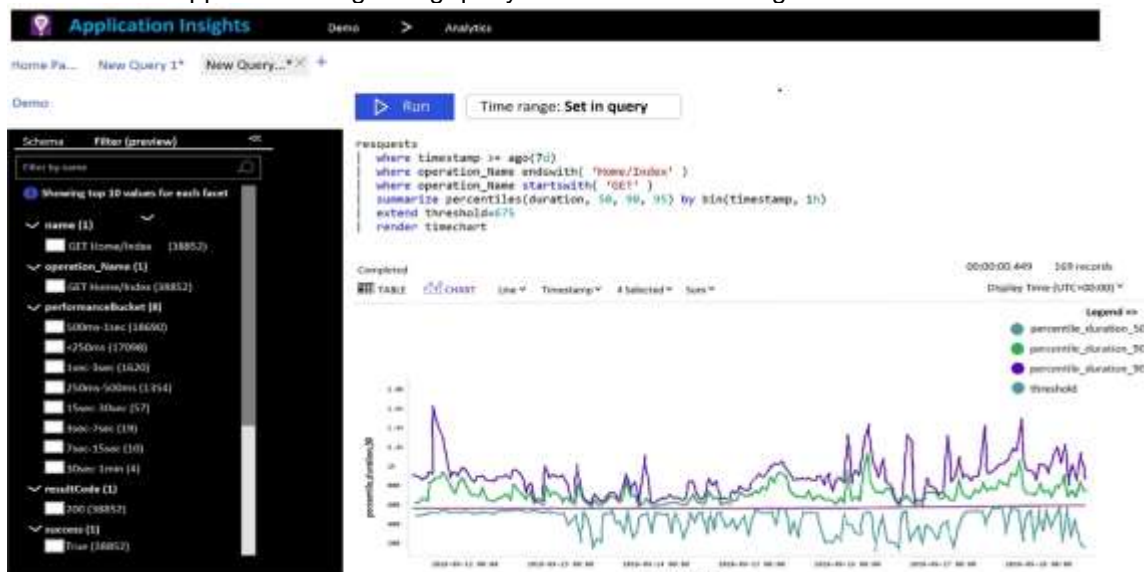
<https://docs.microsoft.com/de-DE/virtualization/windowscontainers/quick-start/building-sample-app>

QUESTION 206

Hotspot Question

You plan to create alerts that will be triggered based on the page load performance of a home page.

You have the Application Insights log query shown in the following exhibit.



Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer Area

To create an alert based on the page load experience of most users, the alerting level must be based on [answer choice].

	▼
percentile_duration_50	
percentile_duration_90	
percentile_duration_95	
threshold	

To only create an alert when authentication error occurs on the server, the query must be filtered on [answer choice].

	▼
item Type	
resultCode	
source	
success	

Answer:

Answer Area

To create an alert based on the page load experience of most users, the alerting level must be based on [answer choice].

	▼
percentile_duration_50	
percentile_duration_90	
percentile_duration_95	
threshold	

To only create an alert when authentication error occurs on the server, the query must be filtered on [answer choice].

	▼
item Type	
resultCode	
source	
success	

Explanation:

Box 1: percentile_duration_95

Box 2: success

For example –
requests

| project name, url, success

| where success == "False"

This will return all the failed requests in my App Insights within the specified time range.

Reference:

<https://devblogs.microsoft.com/premier-developer/alerts-based-on-analytics-query-using-custom-log-search/>

QUESTION 207

Drag and Drop Question

You are configuring the settings of a new Git repository in Azure Repos.

You need to ensure that pull requests in a branch meet the following criteria before they are merged:

- Committed code must compile successfully.
- Pull requests must have a Quality Gate status of Passed in SonarCloud.

Which policy type should you configure for each requirement? To answer, drag the appropriate 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.

Answer Area**Policy Types**

Committed code must compile successfully:

Pull requests must have a Quality Gate status
of Passed in SonarCloud:**Answer:****Answer Area****Policy Types**

Committed code must compile successfully:

Pull requests must have a Quality Gate status
of Passed in SonarCloud:**Explanation:**

Box 1: A check-in policy

Administrators of Team Foundation version control can add check-in policy requirements. These check-in policies require the user to take actions when they conduct a check-in to source control. By default, the following check-in policy types are available:

Builds Requires that the last build was successful before a check-in.

Code Analysis Requires that code analysis is run before check-in.

Work Items Requires that one or more work items be associated with the check-in.

Box 2: Build policy

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/tfvc/add-check-policies><https://azuredevopslabs.com/labs/vstsextend/sonarcloud/>