AZ-400 Q&As

Designing and Implementing Microsoft DevOps Solutions

Pass Microsoft AZ-400 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/az-400.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

QUESTION 1

HOTSPOT

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:

1.

Ensure that the secrets are retrieved by Azure DevOps.

2.

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.

Hot Area:

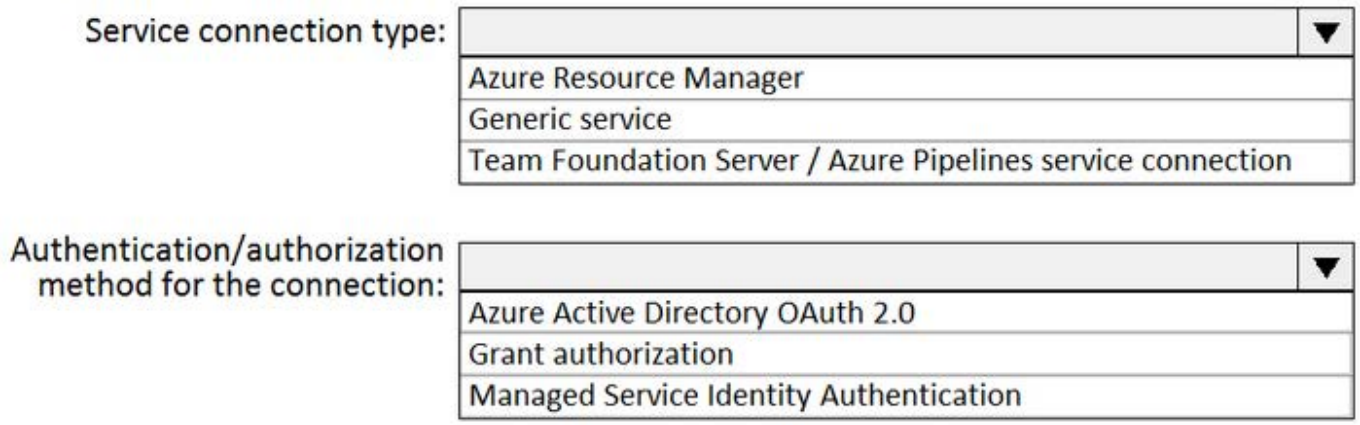

Correct Answer:

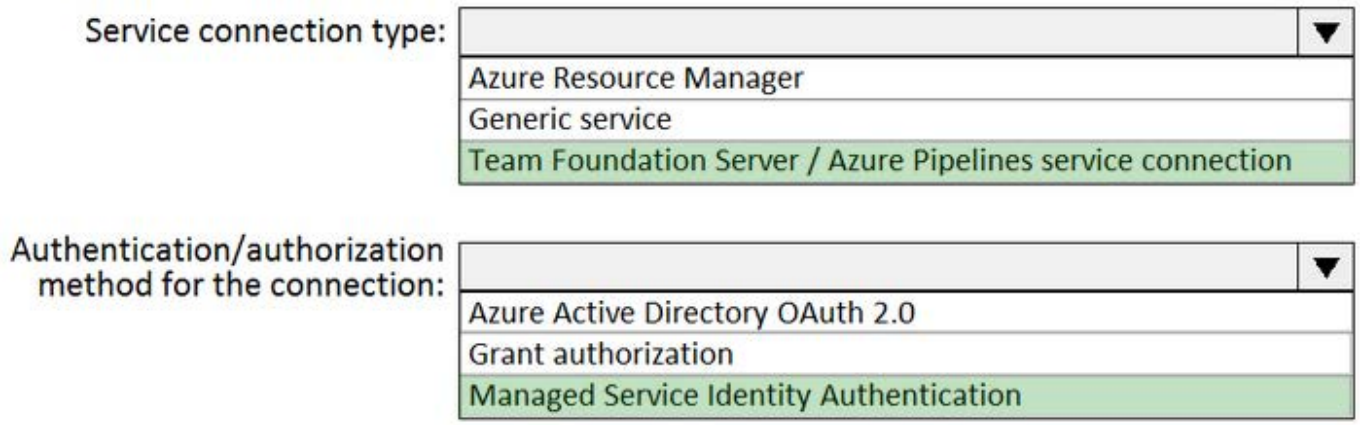

Box 1: Azure Pipelines service connection

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 2

DRAG DROP

You plan to use Azure Kubernetes Service (AKS) to host containers deployed from images hosted in a Docker Trusted Registry.

You need to recommend a solution for provisioning and connecting to AKS. The solution must ensure that AKS is RBAC-enabled and uses a custom service principal.

Which three commands should you recommend be run in sequence? 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:

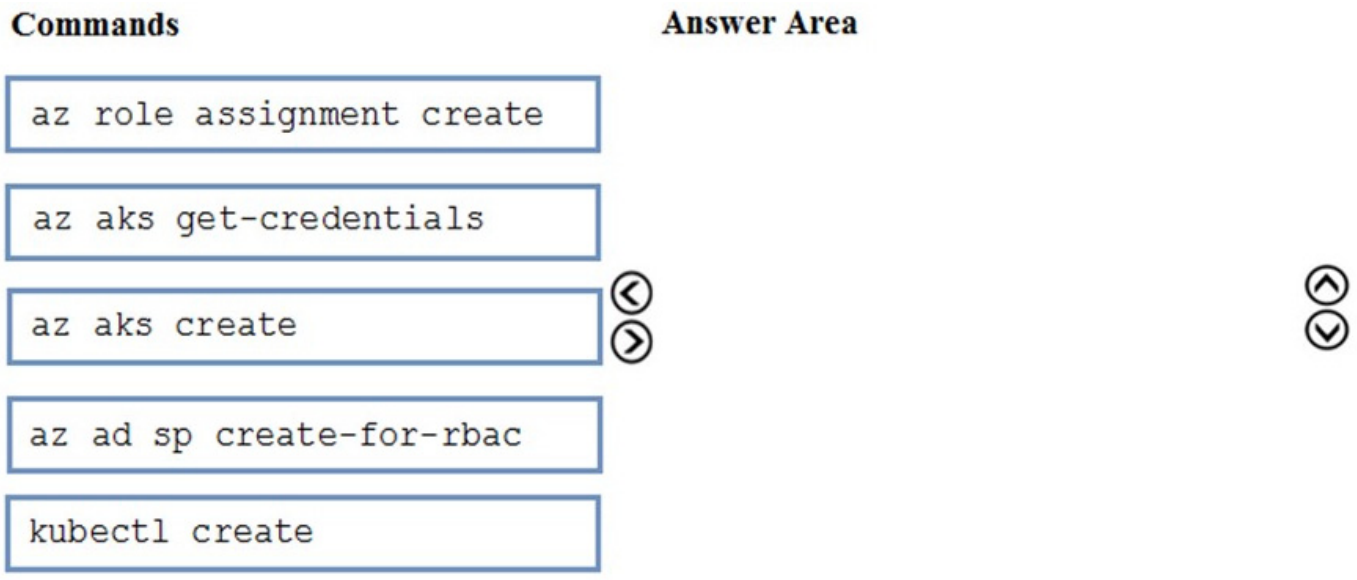

Correct Answer:

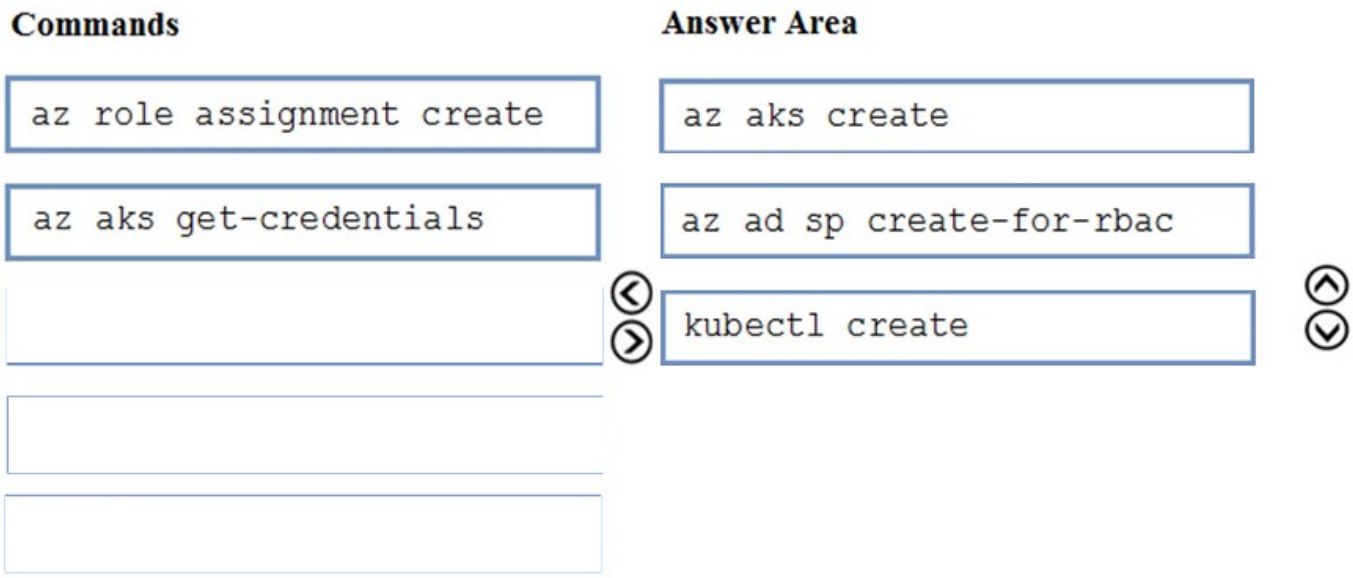

Step 1 : az acr create

An Azure Container Registry (ACR) can also be created using the new Azure CLI.

az acr create

--name

--resource-group

--sku Basic

Step 2: az ad sp create-for-rbac

Once the ACR has been provisioned, you can either enable administrative access (which is okay for testing) or you create a Service Principal (sp) which will provide a client_id and a client_secret.

az ad sp create-for-rbac

--scopes /subscriptions//resourcegroups//providers/Microsoft.ContainerRegistry/registries/

--role Contributor

--name

Step 3: kubectl create

Create a new Kubernetes Secret.

kubectl create secret docker-registry

--docker-server .azurecr.io

--docker-email

--docker-username=

--docker-password

References:

https://thorsten-hans.com/how-to-use-private-azure-container-registry-with-kubernetes

QUESTION 3

You are creating a build pipeline in Azure Pipelines.

You define several tests that might fail due to third-party applications.

You need to ensure that the build pipeline completes successfully if the third-party applications are unavailable.

What should you do?

A. Configure the build pipeline to use parallel jobs

B. Configure flaky tests

C. Increase the test pass percentage

D. Add the Requirements quality widget to your dashboard

Correct Answer: B

Requirements traceability is the ability to relate and document two or more phases of a development process, which can then be traced both forward or backward from its origin. Requirements traceability help teams to get insights into indicators such as quality of requirements or readiness to ship the requirement. A fundamental aspect of requirements traceability is association of the requirements to test cases, bugs and code changes.

Reference: https://docs.microsoft.com/en-us/azure/devops/pipelines/test/requirements-traceability

QUESTION 4

HOTSPOT

You have an Azure Kubernetes Service (AKS) pod.

You need to configure a probe to perform the following actions:

1.

Confirm that the pod is responding to service requests.

2.

Check the status of the pod four times a minute.

3.

Initiate a shutdown if the pod is unresponsive.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
apiVersion: v1
kind: Pod
metadata:
     labels:
          test: readiness-and-liveness
     name: readiness-http
spec:
     containers:
     - name: container1
       image: k8s.ger.io/readiness-and-lieveness
       args:
       - /server
```

[livenessProbe: | readinessProbe: | ShutdownProbe: | startupProbe:]

```
    httpGet:
         path: /checknow
         port: 8123
         httpHeaders:
         - name: Custom-Header
           value: CheckNow
```

[initialDelaySeconds: 15 | periodSeconds: 15 | timeoutSeconds: 15]

Correct Answer:

Answer Area

```
apiVersion: v1
kind: Pod
metadata:
     labels:
          test: readiness-and-liveness
     name: readiness-http
spec:
     containers:
     - name: container1
       image: k8s.ger.io/readiness-and-lieveness
       args:
       - /server
```

livenessProbe:
readinessProbe:
ShutdownProbe:
startupProbe:

```
    httpGet:
         path: /checknow
         port: 8123
         httpHeaders:
         - name: Custom-Header
           value: CheckNow
```

initialDelaySeconds: 15
periodSeconds: 15
timeoutSeconds: 15

Box 1: readinessProbe:

For containerized applications that serve traffic, you might want to verify that your container is ready to handle incoming requests. Azure Container Instances supports readiness probes to include configurations so that your container can\\'t be

accessed under certain conditions.

Incorrect Answers:

livenessProbe: Containerized applications may run for extended periods of time, resulting in broken states that may need to be repaired by restarting the container. Azure Container Instances supports liveness probes so that you can

configure your containers within your container group to restart if critical functionality is not working.

Box 2: periodSeconds: 15

The periodSeconds property designates the readiness command should execute every 15 seconds.

Reference:

https://docs.microsoft.com/en-us/azure/container-instances/container-instances-readiness-probe

QUESTION 5

You manage an Azure web app that supports an e-commerce website.

You need to increase the logging level when the web app exceeds normal usage patterns. The solution must minimize administrative overhead.

Which two resources should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. an Azure Monitor alert that has a dynamic threshold

B. an Azure Automation runbook

C. an Azure Monitor alert that uses an action group that has an email action

D. the Azure Monitor autoscale settings

E. an Azure Monitor alert that has a static threshold

Correct Answer: AB

You can use Azure Monitor to monitor base-level metrics and logs for most services in Azure. You can call Azure Automation runbooks by using action groups or by using classic alerts to automate tasks

based on alerts.

Metric Alert with Dynamic Thresholds detection leverages advanced machine learning (ML) to learn metrics\\' historical behavior, identify patterns and anomalies that indicate possible service issues. It provides support of both a simple UI and

operations at scale by allowing users to configure alert rules through the Azure Resource Manager API, in a fully automated manner.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-dynamic-thresholds

https://docs.microsoft.com/en-us/azure/automation/automation-create-alert-triggered- runbook

QUESTION 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 manage a project in Azure DevOps.

You need to prevent the configuration of the project from changing over time.

Solution: Implement Continuous Assurance for the project.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

The basic idea behind Continuous Assurance (CA) is to setup the ability to check for "drift" from what is considered a secure snapshot of a system. Support for Continuous Assurance lets us treat security truly as a \\'state\\' as opposed to a \\'point in time\\' achievement. This is particularly important in today\\'s context when \\'continuous change\\' has become a norm.

There can be two types of drift:

1.

Drift involving \\'baseline\\' configuration: This involves settings that have a fixed number of possible states (often pre-defined/statically determined ones). For instance, a SQL DB can have TDE encryption turned ON or OFF...or a Storage Account may have auditing turned ON however the log retention period may be less than 365 days.

2.

Drift involving \\'stateful\\' configuration: There are settings which cannot be constrained within a finite set of well-known states. For instance, the IP addresses configured to have access to a SQL DB can be any (arbitrary) set of IP addresses. In such scenarios, usually human judgment is initially required to determine whether a particular configuration should be considered \\'secure\\' or not. However, once that is done, it is important to ensure that there is no "stateful drift" from the attested configuration. (E.g., if, in a troubleshooting session, someone adds the IP address of a developer machine to the list, the Continuous Assurance feature should be able to identify the drift and generate notifications/ alerts or even trigger \\'auto-remediation\\' depending on the severity of the change).

Reference: https://azsk.azurewebsites.net/04-Continous-Assurance/Readme.html

QUESTION 7

Your company is currently making use of Team Foundation Server 2013 (TFS 2013), but intend to migrate to Azure DevOps.

You have been tasked with supplying a migration approach that allows for the preservation of Team Foundation Version Control changesets dates, as well as the changes dates of work items revisions. The approach should also allow for the

migration of all TFS artifacts, while keeping migration effort to a minimum.

You have suggested upgrading TFS to the most recent RTW release.

Which of the following should also be suggested?

A. Installing the TFS kava SDK

B. Using the TFS Database Import Service to perform the upgrade.

C. Upgrading PowerShell Core to the latest version.

D. Using the TFS Integration Platform to perform the upgrade.

Correct Answer: B

In Phase 3 of your migration project, you will work on upgrading your Team Foundation Server to one of the supported versions for the Database Import Service in Azure Devops Services.

QUESTION 8

DRAG DROP

You manage the Git repository for a large enterprise application.

During the development of the application, you use a file named Config.json.

You need to prevent Config.json from being committed to the source control whenever changes to the application are committed.

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:

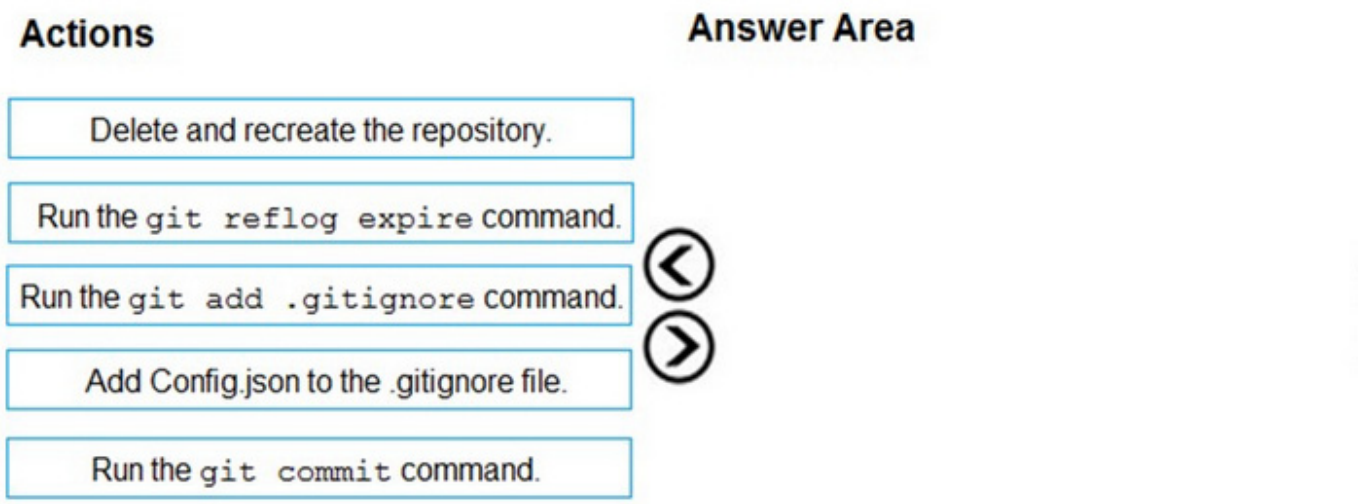

Correct Answer:

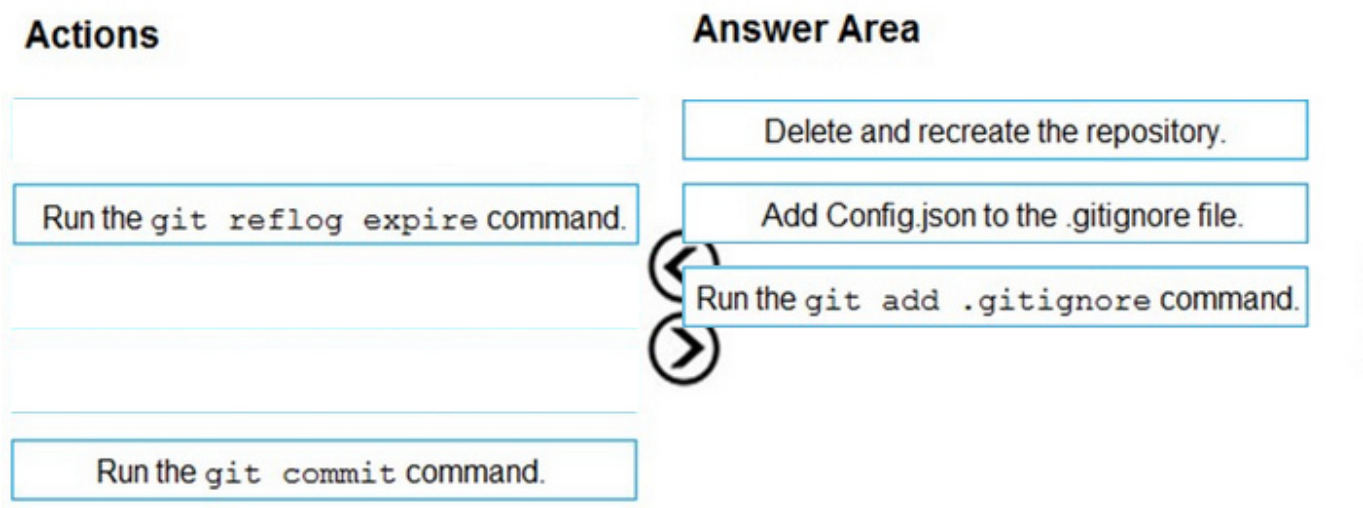

Step 1: Delete and recreate the repository.

Step 2: Add Config.json to the .gitignore file

Each line in the .gitignore excludes a file or set of files that match a pattern.

Example:

ignore a single file

Config.json

Step 3: Run the git add .gitignore command

At the initial commit we want basically move from Untracked to Staged, for staging we have to indicate which file we want to move or specify a pattern, as example:

Reference:

http://hermit.no/how-to-find-the-best-gitignore-for-visual-studio-and-azure-devops/

https://geohernandez.net/how-to-add-an-existing-repository-into-azure-devops-repo-with-git/

QUESTION 9

You have a project in Azure DevOps named Project1 that contains two environments named environment1 and environment2.

When a new version of Project is released, the latest version is deployed to environment2, and the previous version is redeployed to environment1.

You need to distribute users across the environments. The solution must meet the following requirements:

New releases must be available to only a subset of the users.

You must gradually increase the number of users that can access environment2.

What should you use?

A. VIP swaping

B. web app deployment slots

C. Azure Load Balancer

D. Azure Traffic Manager

Correct Answer: B

Explanation:

Use deployment slots

Whenever possible, use deployment slots when deploying a new production build. When using a Standard App Service Plan tier or better, you can deploy your app to a staging environment, validate your changes, and do smoke tests. When

you are ready, you can swap your staging and production slots. The swap operation warms up the necessary worker instances to match your production scale, thus eliminating downtime.

Reference:

https://docs.microsoft.com/en-us/azure/app-service/deploy-best-practices#use-deployment-slots

QUESTION 10

You plan to deploy a solution that will include multiple microservices.

You need to recommend a deployment strategy for the microservices. The solution must meet the following requirements:

1.

Enable users to test new features by using a specific URL.

2.

Minimize the effort required to promote a test version to production.

3.

Minimize the effort required to revert production code to the previous version. Which strategy should you recommend?

A. A/B

B. feature toggle

C. progressive exposure

D. blue/green

Correct Answer: D

QUESTION 11

You manage package feeds by using Azure Artifacts.

You plan to create a new package feed that will include the following views:

1.

@Local

2.

@Latest

3.

@Release

4.

@Prerelease

Which view should you create manually?

A. @Local

B. @Latest

C. @Release

D. @Prerelease

Correct Answer: B

QUESTION 12

You use Azure Repos to manage source code and Azure Pipelines to implement continuous integration and continuous deployment (CI/CD).

You need to ensure that all comments on pull requests are resolved before the pull requests are included in a build. The solution must minimize administrative effort.

What should you include in the solution?

A. a custom action

B. a post-deployment gate

C. a branch policy

D. a pre-deployment gate

Correct Answer: C

QUESTION 13

You are making use of Azure DevOps manage build pipelines, and also deploy pipelines.

The development team is quite large, and is regularly added to.

You have been informed that the management of users and licenses must be automated when it can be.

Which of the following is a task that can\\'t be automated?

A. Group membership changes

B. License assignment

C. Assigning entitlements

D. License procurement

Correct Answer: D

QUESTION 14

DRAG DROP

You have an Azure Kubernetes Service (AKS) implementation that is RBAC-enabled.

You plan to use Azure Container Instances as a hosted development environment to run containers in the AKS implementation.

You need to configure Azure Container Instances as a hosted environment for running the containers in AKS.

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:

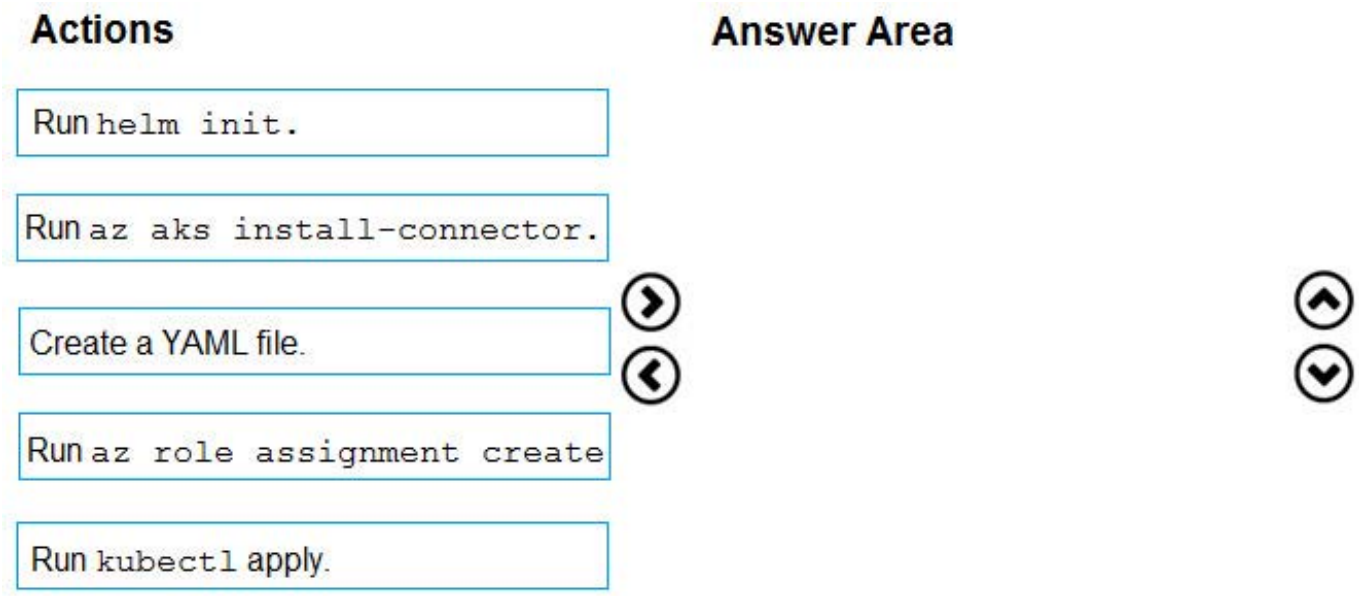

Correct Answer:

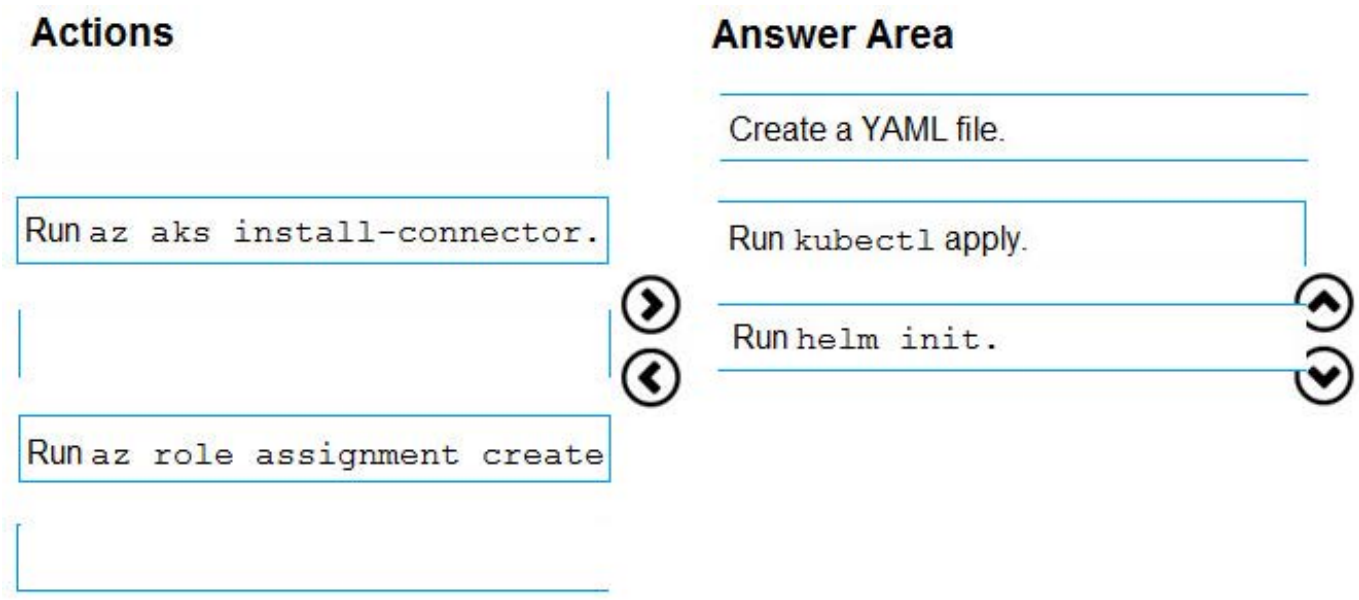

Step 1: Create a YAML file.

If your AKS cluster is RBAC-enabled, you must create a service account and role binding for use with Tiller. To create a service account and role binding, create a file named rbac-virtual-kubelet.yaml

Step 2: Run kubectl apply.

Apply the service account and binding with kubectl apply and specify your rbac-virtual-kubelet.yaml file.

Step 3: Run helm init.

Configure Helm to use the tiller service account:

helm init --service-account tiller

You can now continue to installing the Virtual Kubelet into your AKS cluster.

References: https://docs.microsoft.com/en-us/azure/aks/virtual-kubelet

QUESTION 15

You have an Azure subscription that contains four Azure virtual machines.

You need to configure the virtual machines to use a single identity. The solution must meet the following requirements:

1.

Ensure that the credentials for the identity are managed automatically.

2.

Support granting privileges to the identity. Which type of identity should you use?

A. a system-assigned managed identity

B. a user-assigned managed identity

C. a service principal

D. a user account

Correct Answer: B

Managed Identity is suitable for scenarios where a single resource needs to access another Azure resource, while Service Principal is suitable for more complex scenarios where multiple resources need to access multiple Azure resources.

Reference: https://www.c-sharpcorner.com/blogs/azure-managed-identity-vs-azure-service-principal