

Microsoft.AZ-400.vSep-2024.by.Limda.269q

Number: AZ-400  
Passing Score: 800  
Time Limit: 120  
File Version: 22.0

Exam Code: AZ-400  
Exam Name: Microsoft Azure DevOps Solutions



## 01 - Develop a Site Reliability Engineering (SRE) strategy

### QUESTION 1

You are designing a configuration management solution to support five apps hosted on Azure App Service. Each app is available in the following three environments: development, test, and production.

You need to recommend a configuration management solution that meets the following requirements:

Supports feature flags

Tracks configuration changes from the past 30 days

Stores hierarchically structured configuration values

Controls access to the configurations by using role-based access control (RBAC) permissions

Stores shared values as key/value pairs that can be used by all the apps

Which Azure service should you recommend as the configuration management solution?

- A. Azure Cosmos DB
- B. Azure App Service
- C. Azure App Configuration
- D. Azure Key Vault

**Correct Answer: A**

**Section:**

**Explanation:**



### QUESTION 2

You have a containerized solution that runs in Azure Container Instances. The solution contains a frontend container named App1 and a backend container named DB1. DB1 loads a large amount of data during startup. You need to verify that DB1 can handle incoming requests before users can submit requests to App1.

What should you configure?

- A. a liveness probe
- B. a performance log
- C. a readiness probe
- D. an Azure Load Balancer health probe

**Correct Answer: C**

**Section:**

**Explanation:**

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:

A: 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.

Reference:

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

### QUESTION 3

You are designing a strategy to monitor the baseline metrics of Azure virtual machines that run Windows Server. You need to collect detailed data about the processes running in the guest operating system.

Which two agents should you deploy? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. the Telegraf agent
- B. the Azure Log Analytics agent
- C. the Azure Network Watcher Agent for Windows
- D. the Dependency agent

**Correct Answer: B, D**

**Section:**

**Explanation:**

The following table provide a quick comparison of the Azure Monitor agents for Windows.

	Azure Monitor agent (preview)	Diagnostics extension (WAD)	Log Analytics agent	Dependency agent
<b>Environments supported</b>	Azure	Azure	Azure Other cloud On-premises	Azure Other cloud On-premises
<b>Agent requirements</b>	None	None	None	Requires Log Analytics agent
<b>Data collected</b>	Event Logs Performance	Event Logs ETW events Performance File based logs IIS logs .NET app logs Crash dumps Agent diagnostics logs	Event Logs Performance File based logs IIS logs Insights and solutions Other services	Process dependencies Network connection metrics
<b>Data sent to</b>	Azure Monitor Logs Azure Monitor Metrics	Azure Storage Azure Monitor Metrics Event Hub	Azure Monitor Logs	Azure Monitor Logs (through Log Analytics agent)



Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/agents-overview>

#### QUESTION 4

You have a build pipeline in Azure Pipelines that occasionally fails.

You discover that a test measuring the response time of an API endpoint causes the failures.

You need to prevent the build pipeline from failing due to the test.

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

NOTE: Each correct selection is worth one point.

- A. Set Flaky test detection to Off.
- B. Clear Flaky tests included in test pass percentage.
- C. Enable Test Impact Analysis (TIA).
- D. Manually mark the test as flaky.
- E. Enable test slicing.

**Correct Answer: B, D**

**Section:**

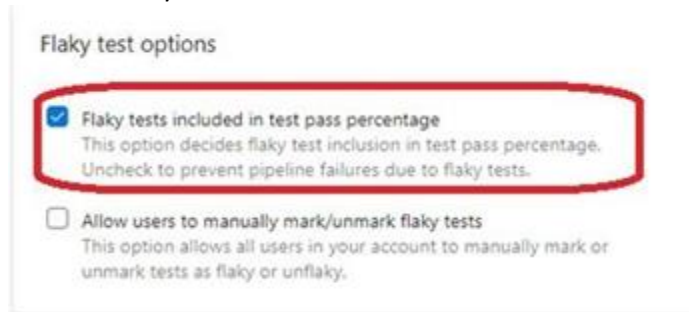
**Explanation:**

D: You can mark or unmark a test as flaky based on analysis or context, by choosing Flaky.

To configure flaky test management, choose Project settings, and select Test management in the Pipelines section.

B:

Slide the On/Off button to On.



Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/test/flaky-test-management>



#### QUESTION 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 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: Add a code coverage step to the build pipelines.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead implement Continuous Assurance for the project.

Reference: <https://azsk.azurewebsites.net/04-Continuous-Assurance/Readme.html>

#### 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 Integration for the project.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead implement Continuous Assurance for the project.

Reference:

<https://azsk.azurewebsites.net/04-Continuous-Assurance/Readme.html>

#### QUESTION 7

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**

**Section:**

**Explanation:**

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:

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.

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-Continuous-Assurance/Readme.html>

#### QUESTION 8

HOTSPOT

You have an application named App1 that has a custom domain of app.contoso.com.

You create a test in Azure Application Insights as shown in the following exhibit.



### Create test

^ Basic Information

\* Test name  
 ✓

[Learn more about configuring tests against applications hosted behind a firewall](#)

Test type:  
 ▼

\* URL ⓘ  
 ✓

Parse dependent requests ⓘ

Enable retries for availability test failures: ⓘ

Test frequency ⓘ  
 ▼

▼ Test locations  
 4 location(s) configured

^ Success criteria

Test Timeout ⓘ  
 ▼

HTTP response ⓘ

Status code must equal:

Content match ⓘ

Content must contain:

▼ Alerts  
 Enabled



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.

Hot Area:

## Answer Area

The test will execute [answer choice].

	▼
every 30 seconds at a random location	
every 30 seconds per location	
every five minutes at a random location	
every five minutes per location	

The test will pass if [answer choice] within 30 seconds.

	▼
App1 responds to an ICMP ping	
the HTML of App1 and the HTML from URLs in <a> tags load	
all the HTML, JavaScripts, and images of App1 load	

Answer Area:

## Answer Area

The test will execute [answer choice].

	▼
every 30 seconds at a random location	
every 30 seconds per location	
every five minutes at a random location	
every five minutes per location	

The test will pass if [answer choice] within 30 seconds.

	▼
App1 responds to an ICMP ping	
the HTML of App1 and the HTML from URLs in <a> tags load	
all the HTML, JavaScripts, and images of App1 load	

Section:

Explanation:

Box 1: every five minutes at a random location

Test frequency: Sets how often the test is run from each test location. With a default frequency of five minutes and five test locations, your site is tested on average every minute.

Box 2:

Parse dependent requests: Test requests images, scripts, style files, and other files that are part of the web page under test. The recorded response time includes the time taken to get these files. The test fails if any of these resources cannot be successfully downloaded within the timeout for the whole test.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability>

**QUESTION 9**

DRAG DROP

You use Azure Pipelines to automate Continuous Integration/Continuous Deployment (CI/CD) for an Azure web app named WebApp1.

You configure an Azure Monitor alert that is triggered when WebApp1 generates an error.

You need to configure the alert to forward details of the error to a third-party system. The solution must minimize administrative effort.

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:

**Actions**

**Answer Area**

Select the Recurrence trigger.

Create an Azure event hub.

Create an Azure logic app.

Select the HTTP request trigger.

Update the action group in Azure Monitor.

Select the Sliding Window trigger.



Correct Answer:



## Actions

## Answer Area

Select the Recurrence trigger.		Create an Azure logic app.
Create an Azure event hub.		Select the HTTP request trigger.
	⬅	Update the action group in Azure Monitor.
	➡	
Select the Sliding Window trigger.		



### Section:

### Explanation:

Box 1: Create an Azure logic app.

Box 2: Select the HTTP request trigger.

Box 3: Updated the action group in Azure Monitor.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/action-groups-logic-app>

### QUESTION 10

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 DevOps organization named Contoso and an Azure subscription. The subscription contains an Azure virtual machine scale set named VMSS1 that is configured for autoscaling.

You have a project in Azure DevOps named Project1. Project1 is used to build a web app named App1 and deploy App1 to VMSS1.

You need to ensure that an email alert is generated whenever VMSS1 scales in or out.

Solution: From Azure DevOps, configure the Notifications settings for Project1.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Notifications help you and your team stay informed about activity that occurs within your projects in Azure DevOps. You can get notified when changes occur to the following items:

- work items
- code reviews
- pull requests
- source control files
- builds

**Reference:**

<https://docs.microsoft.com/en-us/azure/devops/notifications/about-notifications?view=azure-devops>

**QUESTION 11**

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 DevOps organization named Contoso and an Azure subscription. The subscription contains an Azure virtual machine scale set named VMSS1 that is configured for autoscaling.

You have a project in Azure DevOps named Project1. Project1 is used to build a web app named App1 and deploy App1 to VMSS1.

You need to ensure that an email alert is generated whenever VMSS1 scales in or out.

Solution: From Azure DevOps, configure the Service hooks settings for Project1.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**



**QUESTION 12**

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 DevOps organization named Contoso and an Azure subscription. The subscription contains an Azure virtual machine scale set named VMSS1 that is configured for autoscaling.

You have a project in Azure DevOps named Project1. Project1 is used to build a web app named App1 and deploy App1 to VMSS1.

You need to ensure that an email alert is generated whenever VMSS1 scales in or out.

Solution: From Azure Monitor, create an action group.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section:**

**Explanation:**

An action group is a collection of notification preferences defined by the owner of an Azure subscription. Azure Monitor, Service Health and Azure Advisor alerts use action groups to notify users that an alert has been triggered.

**Reference:**

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/action-groups>

**QUESTION 13**

**DRAG DROP**

You are planning projects for three customers. Each customer's preferred process for work items is shown in the following table.

Customer name	Preferred process
Litware, Inc.	Track product backlog items (PBIs) and bugs on the Kanban board. Break the PBIs down into tasks on the task board.
Contoso, Ltd.	Track user stories and bugs on the Kanban board. Track the bugs and tasks on the task board.
A. Datum Corporation	Track requirements, change requests, risks, and reviews.

The customers all plan to use Azure DevOps for work item management.

Which work item process should you use for each customer? To answer, drag the appropriate work item process to the correct customers. Each work item process 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:**

**Processes**

- Agile
- CMMI
- Scrum
- XP

**Answer Area**

- Litware
- Contoso:
- A. Datum:



**Correct Answer:**

**Processes**

- 
- 
- 
- XP

**Answer Area**

- Litware
- Contoso:
- A. Datum:

**Section:**

**Explanation:**

Box 1: Scrum

Choose Scrum when your team practices Scrum. This process works great if you want to track product backlog items (PBIs) and bugs on the Kanban board, or break PBIs and bugs down into tasks on the taskboard.

Box 2: Agile

Choose Agile when your team uses Agile planning methods, including Scrum, and tracks development and test activities separately. This process works great if you want to track user stories and (optionally) bugs on the Kanban board, or track bugs and tasks on the taskboard.

Box 3: CMMI

Choose CMMI when your team follows more formal project methods that require a framework for process improvement and an auditable record of decisions. With this process, you can track requirements, change requests, risks, and reviews.

Incorrect Answers:

XP:

The work tracking objects contained within the default DevOps processes and DevOps process templates are Basic, Agile, CMMI, and Scrum XP (Extreme Programming) and DevOps are different things. They don't contradict with each other, they can be used together, but they have different base concepts inside them.

References:

<https://docs.microsoft.com/en-us/azure/devops/boards/work-items/guidance/choose-process?view=azure-devops>

#### QUESTION 14

Your company hosts a web application in Azure. The company uses Azure Pipelines for the build and release management of the application. Stakeholders report that the past few releases have negatively affected system performance.

You configure alerts in Azure Monitor.

You need to ensure that new releases are only deployed to production if the releases meet defined performance baseline criteria in the staging environment first. What should you use to prevent the deployment of releases that fall to meet the performance baseline?

- A. an Azure Scheduler job
- B. a trigger
- C. a gate
- D. an Azure function

**Correct Answer: C**

**Section:**

**Explanation:**

Scenarios and use cases for gates include:

- Quality validation. Query metrics from tests on the build artifacts such as pass rate or code coverage and deploy only if they are within required thresholds.

Quality validation. Query metrics from tests on the build artifacts such as pass rate or code coverage and deploy only if they are within required thresholds.

Use Quality Gates to integrate monitoring into your pre-deployment or post-deployment. This ensures that you are meeting the key health/performance metrics (KPIs) as your applications move from dev to production and any differences in the infrastructure environment or scale is not negatively impacting your KPIs.

Note: Gates allow automatic collection of health signals from external services, and then promote the release when all the signals are successful at the same time or stop the deployment on timeout. Typically, gates are used in connection with incident management, problem management, change management, monitoring, and external approval systems.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/continuous-monitoring>

<https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/gates?view=azure-devops>

#### QUESTION 15

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: Perform a Subscription Health scan when packages are created.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**



**Section:****Explanation:**

Instead implement Continuous Assurance for the project.

Note: The Subscription Security health check features in AzSK contains a set of scripts that examines a subscription and flags off security issues, misconfigurations or obsolete artifacts/settings which can put your subscription at higher risk.

Reference: <https://azsk.azurewebsites.net/04-Continuous-Assurance/Readme.html>

**QUESTION 16**

Your company uses the following resources:

Windows Server 2019 container images hosted in an Azure Container Registry.

Azure virtual machines that run the latest version of Ubuntu

An Azure Log Analytics workspace

Azure Active Directory (Azure AD)

An Azure key vault

For which two resources can you receive vulnerability assessments in Azure Security Center? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. the Azure Log Analytics workspace
- B. the Azure key vault
- C. the Azure virtual machines that run the latest version of Ubuntu
- D. Azure Active Directory (Azure AD)
- E. The Windows Server 2019 container images hosted in the Azure Container Registry.

**Correct Answer: C, E**

**Section:****Explanation:****QUESTION 17**

You use Azure Pipelines to manage build pipelines, GitHub to store source code, and Dependabot to manage dependencies. You have an app named App1.

Dependabot detects a dependency in App1 that requires an update.

What should you do first to apply the update?

- A. Create a pull request.
- B. Approve the pull request.
- C. Create a branch.
- D. Perform a commit.

**Correct Answer: B**

**Section:****Explanation:**

Dependabot is a useful tool to regularly check for dependency updates. By helping to keep your project up to date, Dependabot can reduce technical debt and immediately apply security vulnerabilities when patches are released. How does Dependabot work?

1. Dependabot regularly checks dependencies for updates
2. If an update is found, Dependabot creates a new branch with this upgrade and Pull Request for approval
3. You review the new Pull Request, ensure the tests passed, review the code, and decide if you can merge the change

Reference:

<https://samlearnsazure.blog/2019/12/20/github-using-dependabot/>

## 02 - Develop a Site Reliability Engineering (SRE) strategy

### 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.

### Overview

Litware, Inc. is an independent software vendor (ISV). Litware has a main office and five branch offices.

### Existing Environment

#### Application Architecture

The company's primary application is a single monolithic retirement fund management system based on ASP.NET web forms that use logic written in VB.NET. Some new sections of the application are written in C#.

Variations of the application are created for individual customers. Currently, there are more than 80 live code branches in the application's code base.

The application was developed by using Microsoft Visual Studio. Source code is stored in Team Foundation Server (TFS) in the main office. The branch offices access the source code by using TFS proxy servers.

#### Architectural Issues

Litware focuses on writing new code for customers. No resources are provided to refactor or remove existing code. Changes to the code base take a long time, as dependencies are not obvious to individual developers.

Merge operations of the code often take months and involve many developers. Code merging frequently introduces bugs that are difficult to locate and resolve.

Customers report that ownership costs of the retirement fund management system increase continually. The need to merge unrelated code makes even minor code changes expensive.

Customers report that bug reporting is overly complex.

### Requirements

#### Planned Changes

Litware plans to develop a new suite of applications for investment planning. The investment planning applications will require only minor integration with the existing retirement fund management system.

The investment planning applications suite will include one multi-tier web application and two iOS mobile applications. One mobile application will be used by employees; the other will be used by customers.

Litware plans to move to a more agile development methodology. Shared code will be extracted into a series of packages.

Litware has started an internal cloud transformation process and plans to use cloud-based services whenever suitable.

Litware wants to become proactive in detecting failures, rather than always waiting for customer bug reports.

#### Technical Requirements

The company's investment planning applications suite must meet the following requirements:

New incoming connections through the firewall must be minimized.

Members of a group named Developers must be able to install packages.

The principle of least privilege must be used for all permission assignments.

A branching strategy that supports developing new functionality in isolation must be used.

Members of a group named Team Leaders must be able to create new packages and edit the permissions of package feeds. Visual Studio App Center must be used to centralize the reporting of mobile application crashes and device types in use. By default, all releases must remain available for 30 days, except for production releases, which must be kept for 60 days. Code quality and release quality are critical. During release, deployments must not proceed between stages if any active bugs are logged against the release. The mobile applications must be able to call the share pricing service of the existing retirement fund management system. Until the system is upgraded, the service will only support basic authentication over HTTPS. The required operating system configuration for the test servers changes weekly. Azure Automation State Configuration must be used to ensure that the operating system on each test server is configured the same way when the servers are created and checked periodically.

#### Current Technical Issue

The test servers are configured correctly when first deployed, but they experience configuration drift over time. Azure Automation State Configuration fails to correct the configurations.

Azure Automation State Configuration nodes are registered by using the following command.

```
Register-AzureRmAutomationDscNode
-ResourceGroupName 'TestResourceGroup'
-AutomationAccountName 'LitwareAutomationAccount'
-AzureVMName $vmname
-ConfigurationMode 'ApplyOnly'
```

## QUESTION 1

### 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 question on 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 sections 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 on 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.

#### Application Architecture

The company's primary application is a single monolithic retirement fund management system based on ASP.NET web forms that use logic written in VB.NET. Some new sections of the application are written in C#.

Variations of the application are created for individual customers. Currently, there are more than 80 live code branches in the application's code base.

The application was developed by using Microsoft Visual Studio. Source code is stored in Team Foundation Server (TFS) in the main office. The branch offices access the source code by using TFS proxy servers.

#### Architectural Issues

Litware focuses on writing new code for customers. No resources are provided to refactor or remove existing code. Changes to the code base take a long time, as dependencies are not obvious to individual developers.

Merge operations of the code often take months and involve many developers. Code merging frequently introduces bugs that are difficult to locate and resolve. Customers report that ownership costs of the retirement fund management system increase continually. The need to merge unrelated code makes even minor code changes expensive.

Customers report that bug reporting is overly complex.

#### Planned changes

Litware plans to develop a new suite of applications for investment planning. The investment planning applications will require only minor integration with the existing retirement fund management system.

The investment planning applications suite will include one multi-tier web application and two iOS mobile application. One mobile application will be used by employees; the other will be used by customers.

Litware plans to move to a more agile development methodology. Shared code will be extracted into a series of packages.

Litware has started an internal cloud transformation process and plans to use cloud-based services whenever suitable.

Litware wants to become proactive in detecting failures, rather than always waiting for customer bug reports.

#### Technical requirements

The company's investment planning applications suite must meet the following requirements:

New incoming connections through the firewall must be minimized.

Members of a group named Developers must be able to install packages.

The principle of least privilege must be used for all permission assignments.

A branching strategy that supports developing new functionality in isolation must be used.

Members of a group named Team Leaders must be able to create new packages and edit the permissions of package feeds.

Visual Studio App Center must be used to centralize the reporting of mobile application crashes and device types in use.

By default, all releases must remain available for 30 days, except for production releases, which must be kept for 60 days.

Code quality and release quality are critical. During release, deployments must not proceed between stages if any active bugs are logged against the release.

The mobile applications must be able to call the share pricing service of the existing retirement fund management system. Until the system is upgraded, the service will only support basic authentication over HTTPS.

The required operating system configuration for the test servers changes weekly. Azure Automation State Configuration must be used to ensure that the operating system on each test server is configured the same way when the servers are created and checked periodically.

#### Current Technical Issue

The test servers are configured correctly when first deployed, but they experience configuration drift over time. Azure Automation State Configuration fails to correct the configurations.

Azure Automation State Configuration nodes are registered by using the following command.

```
Register-AzureRmAutomationDscNode  
-ResourceGroupName 'TestResourceGroup'  
-AutomationAccountName 'LitwareAutomationAccount'  
-AzureVMName $vmname  
-ConfigurationMode 'ApplyOnly'
```

How should you complete the code to initialize App Center in the mobile application? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

```
MSAppCenter.start  
( "{Your App Secret}",  
  withServices:  
)
```

[MSAnalytics.self, [MSDistribute.self, [MSPush.self,	MSAnalytics.self] MSCrashes.self] MSDistribute.self]

Answer Area:

## Answer Area

```
MSAppCenter.start  
( "{Your App Secret}",  
  withServices:  
)
```

[MSAnalytics.self, [MSDistribute.self, [MSPush.self,	MSAnalytics.self] MSCrashes.self] MSDistribute.self]

Section:



**Explanation:**

Scenario: Visual Studio App Center must be used to centralize the reporting of mobile application crashes and device types in use. In order to use App Center, you need to opt in to the service(s) that you want to use, meaning by default no services are started and you will have to explicitly call each of them when starting the SDK. Insert the following line to start the SDK in your app's AppDelegate class in the didFinishLaunchingWithOptions method. `MSAppCenter.start("{Your App Secret}", withServices: [MSAnalytics.self, MSCrashes.self])` References: <https://docs.microsoft.com/en-us/appcenter/sdk/getting-started/ios>

**03 - Develop a Site Reliability Engineering (SRE) strategy**

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.

Overview

Contoso, Ltd. is a manufacturing company that has a main office in Chicago.

Existing Environment

Contoso plans to improve its IT development and operations processes by implementing Azure DevOps principles. Contoso has an Azure subscription and creates an Azure DevOps organization.

The Azure DevOps organization includes:

The Docker extension

A deployment pool named Pool7 that contains 10 Azure virtual machines that run Windows Server 2019

The Azure subscription contains an Azure Automation account.

Requirements

Planned changes

Contoso plans to create projects in Azure DevOps as shown in the following table.



Project name	Project details
Project 1	Project1 will provide support for incremental builds and third-party SDK components
Project 2	Project2 will use an automatic build policy. A small team of developers named Team2 will work independently on changes to the project. The Team2 members will not have permissions to Project2.
Project 3	Project3 will be integrated with SonarQube
Project 4	Project4 will provide support for a build pipeline that creates a Docker image and pushes the image to the Azure Container Registry. Project4 will use an existing Dockerfile.
Project 5	Project5 will contain a Git repository in Azure Repos and a continuous integration trigger that will initiate a build in response to any change except for changes within /folder1 of the repository.
Project 6	Project6 will provide support for build and deployment pipelines. Deployment will be allowed only if the number of current work items representing active software bugs is 0.
Project 7	Project7 will contain a target deployment group named Group7 that maps to Pool7. Project7 will use Azure Automation State Configuration to maintain the desired state of the computers in Group7.

Technical requirements

Contoso identifies the following technical requirements:

Implement build agents for Project1.

Whenever possible, use Azure resources.

Avoid using deprecated technologies.

Implement a code flow strategy for Project2 that will:

- Enable Team2 to submit pull requests for Project2.
- Enable Team2 to work independently on changes to a copy of Project2.

- Ensure that any intermediary changes performed by Team2 on a copy of Project2 will be subject to the same restrictions as the ones defined in the build policy of Project2. Whenever possible, implement automation and minimize administrative effort.

Implement Project3, Project5, Project6, and Project7 based on the planned changes.

Implement Project4 and configure the project to push Docker images to Azure Container Registry.

### QUESTION 1

You add the virtual machines as managed nodes in Azure Automation State Configuration.

You need to configure the managed computers in Pool7.

What should you do next?

- A. Modify the RefreshMode property of the Local Configuration Manager (LCM).
- B. Run the Register-AzureRmAutomationDscNode Azure Powershell cmdlet.
- C. Modify the ConfigurationMode property of the Local Configuration Manager (LCM).
- D. Install PowerShell Core.

**Correct Answer: B**

**Section:**

**Explanation:**

The Register-AzureRmAutomationDscNode cmdlet registers an Azure virtual machine as an APS Desired State Configuration (DSC) node in an Azure Automation account.

Scenario: The Azure DevOps organization includes:

The Docker extension

A deployment pool named Pool7 that contains 10 Azure virtual machines that run Windows Server 2019

Project 7	Project7 will contain a target deployment group named Group7 that maps to Pool7. Project7 will use Azure Automation State Configuration to maintain the desired state of the computers in Group7.
-----------	---

Reference:

<https://docs.microsoft.com/en-us/powershell/module/azurermsautomation/register-azurermsautomationdscnode>



### QUESTION 2

DRAG DROP

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 question on 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 sections 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 on 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

Contoso, Ltd. is a manufacturing company that has a main office in Chicago.

Contoso plans to improve its IT development and operations processes by implementing Azure DevOps principles. Contoso has an Azure subscription and creates an Azure DevOps organization.

The Azure DevOps organization includes:

The Docker extension

A deployment pool named Pool7 that contains 10 Azure virtual machines that run Windows Server 2019

The Azure subscription contains an Azure Automation account.

Contoso plans to create projects in Azure DevOps as shown in the following table.

Project name	Project details
Project 1	Project1 will provide support for incremental builds and third-party SDK components
Project 2	Project2 will use an automatic build policy. A small team of developers named Team2 will work independently on changes to the project. The Team2 members will not have permissions to Project2.
Project 3	Project3 will be integrated with SonarQube
Project 4	Project4 will provide support for a build pipeline that creates a Docker image and pushes the image to the Azure Container Registry. Project4 will use an existing Dockerfile.
Project 5	Project5 will contain a Git repository in Azure Reports and a continuous integration trigger that will initiate a build in response to any change except for changes within /folder1 of the repository.
Project 6	Project6 will provide support for build and deployment pipelines. Deployment will be allowed only if the number of current work items representing active software bugs is 0.
Project 7	Project7 will contain a target deployment group named Group7 that maps to Pool7. Project7 will use Azure Automation State Configuration to maintain the desired state of the computers in Group7.

**Technical requirements**

Contoso identifies the following technical requirements:

Implement build agents for Project1.

Whenever possible, use Azure resources.

Avoid using deprecated technologies.

Implement a code flow strategy for Project2 that will:

-Enable Team2 to submit pull requests for Project2.

-Enable Team2 to work independently on changes to a copy of Project2.

-Ensure that any intermediary changes performed by Team2 on a copy of Project2 will be subject to the same restrictions as the ones defined in the build policy of Project2. Whenever possible implement automation and minimize administrative effort.

Implement Project3, Project5, Project6, and Project7 based on the planned changes

Implement Project4 and configure the project to push Docker images to Azure Container Registry.

You need to implement the code flow strategy for Project2 in Azure DevOps.

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:

**Select and Place:**

**Actions**

**Answer Area**

Create a fork

Create a branch

Add a build policy for the fork

Add a build policy for the master branch

Create a repository

Add an application access policy.



Correct Answer:

### Actions

Create a branch

Add a build policy for the master branch

Add an application access policy.

### Answer Area

Create a repository

Create a fork

Add a build policy for the fork

### Section:

#### Explanation:

Step 1: Create a repository

A Get repository, or repo, is a folder that you've told Git to help you track file changes in. You can have any number of repos on your computer, each stored in their own folder.

Step 2: Create a fork

Step 3: Add a build policy for the fork

Build policies help teams protect their important branches of development. Policies enforce your team's code quality and change management standards.

Scenario:

Implement a code flow strategy for Project2 that will:

Enable Team2 to submit pull requests for Project2.

Enable Team2 to work independently on changes to a copy of Project2.

Ensure that any intermediary changes performed by Team2 on a copy of Project2 will be subject to the same restrictions as the ones defined in the build policy of Project2.

Project2 will use an automatic build policy. A small team of developers named Team2 will work independently on changes to the project. The Team2 members will not have permissions to Project2.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/manage-your-branches>

### QUESTION 3

#### DRAG DROP

#### 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 question on 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 sections 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 on 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

Contoso, Ltd. is a manufacturing company that has a main office in Chicago.

Contoso plans to improve its IT development and operations processes by implementing Azure DevOps principles. Contoso has an Azure subscription and creates an Azure DevOps organization.

The Azure DevOps organization includes:

The Docker extension

A deployment pool named Pool7 that contains 10 Azure virtual machines that run Windows Server 2019

The Azure subscription contains an Azure Automation account.

Contoso plans to create projects in Azure DevOps as shown in the following table.

Project name	Project details
Project 1	Project1 will provide support for incremental builds and third-party SDK components
Project 2	Project2 will use an automatic build policy. A small team of developers named Team2 will work independently on changes to the project. The Team2 members will not have permissions to Project2.
Project 3	Project3 will be integrated with SonarQube
Project 4	Project4 will provide support for a build pipeline that creates a Docker image and pushes the image to the Azure Container Registry. Project4 will use an existing Dockerfile.
Project 5	Project5 will contain a Git repository in Azure Reports and a continuous integration trigger that will initiate a build in response to any change except for changes within /folder1 of the repository.
Project 6	Project6 will provide support for build and deployment pipelines. Deployment will be allowed only if the number of current work items representing active software bugs is 0.
Project 7	Project7 will contain a target deployment group named Group7 that maps to Pool7. Project7 will use Azure Automation State Configuration to maintain the desired state of the computers in Group7.

## Technical requirements

Contoso identifies the following technical requirements:

Implement build agents for Project1.

Whenever possible, use Azure resources.

Avoid using deprecated technologies.

Implement a code flow strategy for Project2 that will:

-Enable Team2 to submit pull requests for Project2.

-Enable Team2 to work independently on changes to a copy of Project2.

-Ensure that any intermediary changes performed by Team2 on a copy of Project2 will be subject to the same restrictions as the ones defined in the build policy of Project2. Whenever possible implement automation and minimize administrative effort.

Implement Project3, Project5, Project6, and Project7 based on the planned changes

Implement Project4 and configure the project to push Docker images to Azure Container Registry.

You need to configure Azure Automation for the computers in Group7.

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:

## Select and Place:



**Actions**

Run the `Import-AzureRmAutomationDscConfiguration` Azure PowerShell cmdlet.


Create a Desired State Configuration (DSC) configuration file that has an extension of `.ps1`.

Run the `New-AzureRmResourceGroupDeployment` Azure PowerShell cmdlet.

Run the `Start-AzureRmAutomationDscCompilationJob` Azure PowerShell cmdlet.

Create an Azure Resource Manager template file that has an extension of `.json`.

**Answer Area**



**Correct Answer:**

**Actions**

Run the `New-AzureRmResourceGroupDeployment` Azure PowerShell cmdlet.

Create an Azure Resource Manager template file that has an extension of `.json`.

**Answer Area**

Create a Desired State Configuration (DSC) configuration file that has an extension of `.ps1`.

Run the `Import-AzureRmAutomationDscConfiguration` Azure PowerShell cmdlet.

Run the `Start-AzureRmAutomationDscCompilationJob` Azure PowerShell cmdlet.

**Section:**

**Explanation:**

Step 1: Create a Desired State Configuration (DSC) configuration file that has an extension of `.ps1`.

Step 2: Run the `Import-AzureRmAutomationDscConfiguration` Azure Powershell cmdlet

The `Import-AzureRmAutomationDscConfiguration` cmdlet imports an APS Desired State Configuration (DSC) configuration into Azure Automation. Specify the path of an APS script that contains a single DSC configuration.

Example:

```
PS C:\>Import-AzureRmAutomationDscConfiguration -AutomationAccountName "Contoso17"-ResourceGroupName "ResourceGroup01" -SourcePath "C:\DSC\client.ps1" -Force
```

This command imports the DSC configuration in the file named `client.ps1` into the Automation account named `Contoso17`. The command specifies the `Force` parameter. If there is an existing DSC configuration, this command replaces it.

Step 3: Run the Start-AzureRmAutomationDscCompilationJob Azure Powershell cmdlet

The Start-AzureRmAutomationDscCompilationJob cmdlet compiles an APS Desired State Configuration (DSC) configuration in Azure Automation.

References:

<https://docs.microsoft.com/en-us/powershell/module/azurermsautomation/import-azurermsautomationdscconfiguration>

<https://docs.microsoft.com/en-us/powershell/module/azurermsautomation/start-azurermsautomationdsc compilationjob>

## 01 - Develop a security and compliance plan

### QUESTION 1

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 plan to update the Azure DevOps strategy of your company.

You need to identify the following issues as they occur during the company's development process:

Licensing violations Prohibited libraries

Solution: You implement continuous integration.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: A**

**Section:**

**Explanation:**

WhiteSource is the leader in continuous open source software security and compliance management. WhiteSource integrates into your build process, irrespective of your programming languages, build tools, or development environments. It works automatically, continuously, and silently in the background, checking the security, licensing, and quality of your open source components against WhiteSource constantly-updated definitive database of open source repositories.

Reference:

<https://azuredevopslabs.com/labs/vstsextend/whitesource/>

### QUESTION 2

You have an Azure Resource Manager template that deploys a multi-tier application.

You need to prevent the user who performs the deployment from viewing the account credentials and connection strings used by the application. What should you use?

A. Azure Key Vault

B. a Web.config file

C. an Appsettings.json file

D. an Azure Storage table

E. an Azure Resource Manager parameter file

**Correct Answer: A**

**Section:**

**Explanation:**

When you need to pass a secure value (like a password) as a parameter during deployment, you can retrieve the value from an Azure Key Vault. You retrieve the value by referencing the key vault and secret in your parameter file. The value is never exposed because you only reference its key vault ID. The key vault can exist in a different subscription than the resource group you are deploying to.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-manager-keyvault-parameter>

### QUESTION 3

SIMULATION

Your company plans to implement a new compliance strategy that will require all Azure web apps to be backed up every five hours. You need to back up an Azure web app named az400-11566895-main every five hours to an Azure Storage account in your resource group. To complete this task, sign in to the Microsoft Azure portal.

A. See solution below.

**Correct Answer: A**

**Section:**

**Explanation:**

With the storage account ready, you can configure backs up in the web app or App Service.

1. Open the App Service az400-11566895-main, which you want to protect, in the Azure Portal and browse to Settings > Backups. Click Configure and a Backup Configuration blade should appear.
2. Select the storage account.
3. Click + to create a private container. You could name this container after the web app or App Service.
4. Select the container.
5. If you want to schedule backups, then set Scheduled Backup to On and configure a schedule: every five hours
6. Select your retention. Note that 0 means never delete backups.
7. Decide if at least one backup should always be retained.
8. Choose if any connected databases should be included in the web app backup.
9. Click Save to finalize the backup configuration.

The screenshot shows the 'Backup Configuration' blade in the Azure portal. It is divided into three main sections: 'Backup Storage', 'Backup Schedule', and 'Backup Database'.  
1. **Backup Storage:** Shows 'Storage Settings' for a container named 'petri' within the storage account 'petriaspbackup.blob.core.windows.net'.  
2. **Backup Schedule:** The 'Scheduled backup' toggle is set to 'On'. The 'Backup Every' field is set to '1' hour. The 'Start backup schedule from' is set to '2018-01-20 16:31:38' in 'UTC - Coordinated Universal Time'. The 'Retention (Days)' is set to '3655'. The 'Keep at least one backup' toggle is set to 'Yes'.  
3. **Backup Database:** A table lists databases to be included in the backup. The table has columns for 'INCLUDE IN BACKUP', 'CONNECTION STRING NAME', and 'DATABASE TYPE'. One entry is shown: 'Included' (checked), 'defaultConnection', and 'Sql Database'.

Reference:

<https://petri.com/backing-azure-app-service>

**QUESTION 4**

**SIMULATION**

You need to configure a virtual machine named VM1 to securely access stored secrets in an Azure Key Vault named az400-11566895-kv. To complete this task, sign in to the Microsoft Azure portal.



A. See solution below.

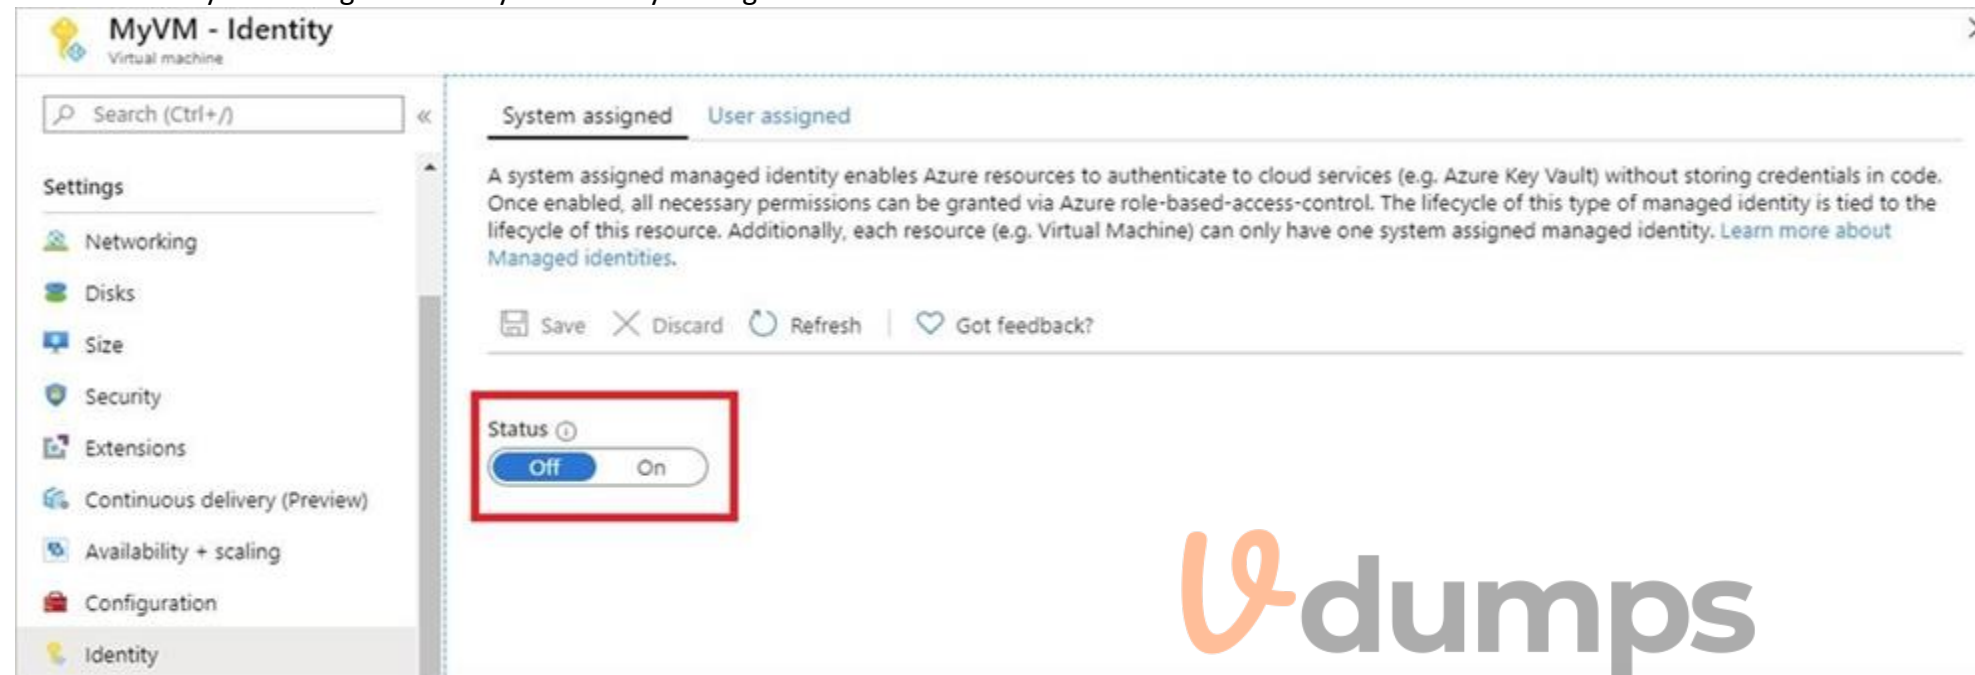
**Correct Answer: A**

**Section:**

**Explanation:**

You can use a system-assigned managed identity for a Windows virtual machine (VM) to access Azure Key Vault.

1. Sign in to Azure portal
2. Locate virtual machine VM1.
3. Select Identity
4. Enable the system-assigned identity for VM1 by setting the Status to On.



Note: Enabling a system-assigned managed identity is a one-click experience. You can either enable it during the creation of a VM or in the properties of an existing VM.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-nonaad>

#### QUESTION 5

You are deploying a server application that will run on a Server Core installation of Windows Server 2019. You create an Azure key vault and a secret.

You need to use the key vault to secure API secrets for third-party integrations.

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

NOTE: Each correct selection is worth one point.

- A. Configure RBAC for the key vault.
- B. Modify the application to access the key vault.
- C. Configure a Key Vault access policy.
- D. Deploy an Azure Desired State Configuration (DSC) extension.
- E. Deploy a virtual machine that uses a system-assigned managed identity.

**Correct Answer: B, C, E**

**Section:**

**Explanation:**

BE: An app deployed to Azure can take advantage of Managed identities for Azure resources, which allows the app to authenticate with Azure Key Vault using Azure AD authentication without credentials (Application ID and Password/Client Secret) stored in the app.

C:

1. Select Add Access Policy.
2. Open Secret permissions and provide the app with Get and List permissions.
3. Select Select principal and select the registered app by name. Select the Select button.
4. Select OK.
5. Select Save.
6. Deploy the app.

Reference:

<https://docs.microsoft.com/en-us/aspnet/core/security/key-vault-configuration>

#### QUESTION 6

You have an Azure DevOps organization named Contoso that contains a project named Project1.

You provision an Azure key vault named Keyvault1.

You need to reference Keyvault1 secrets in a build pipeline of Project1.

What should you do first?

- A. Add a secure file to Project1.
- B. Create an XAML build service.
- C. Create a variable group in Project1.
- D. Configure the security policy of Contoso.

**Correct Answer: A**

**Section:**

**Explanation:**

Before this will work, the build needs permission to access the Azure Key Vault. This can be added in the Azure Portal. Open the Access Policies in the Key Vault and add a new one. Choose the principle used in the DevOps build.

Reference:

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

#### QUESTION 7

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 plan to update the Azure DevOps strategy of your company.

You need to identify the following issues as they occur during the company's development process:

Licensing violations Prohibited libraries

Solution: You implement pre-deployment gates.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use implement continuous integration.

Note: WhiteSource is the leader in continuous open source software security and compliance management. WhiteSource integrates into your build process, irrespective of your programming languages, build tools, or development environments. It works automatically, continuously, and silently in the background, checking the security, licensing, and quality of your open source components against WhiteSource constantly-updated definitive database of open source repositories.

Reference:

<https://azuredevopslabs.com/labs/vstsextend/whitesource/>

### QUESTION 8

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 plan to update the Azure DevOps strategy of your company.

You need to identify the following issues as they occur during the company's development process:

Licensing violations Prohibited libraries

Solution: You implement automated security testing.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use implement continuous integration.

Note: WhiteSource is the leader in continuous open source software security and compliance management. WhiteSource integrates into your build process, irrespective of your programming languages, build tools, or development environments. It works automatically, continuously, and silently in the background, checking the security, licensing, and quality of your open source components against WhiteSource constantly-updated definitive database of open source repositories.

Reference:

<https://azuredevopslabs.com/labs/vstsextend/whitesource/>

### QUESTION 9

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 plan to update the Azure DevOps strategy of your company.

You need to identify the following issues as they occur during the company's development process:

Licensing violations Prohibited libraries

Solution: You implement continuous deployment.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead implement continuous integration.

Note: WhiteSource is the leader in continuous open source software security and compliance management. WhiteSource integrates into your build process, irrespective of your programming languages, build tools, or development environments. It works automatically, continuously, and silently in the background, checking the security, licensing, and quality of your open source components against WhiteSource constantly-updated definitive database of open source repositories.

Reference:

<https://azuredevopslabs.com/labs/vstsextend/whitesource/>

### QUESTION 10

SIMULATION

You manage a website that uses an Azure SQL Database named db1 in a resource group named RG1lod11566895. You need to modify the SQL database to protect against SQL injection.

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

A. See solution below.

**Correct Answer: A**

**Section:**

**Explanation:**

Set up Advanced Threat Protection in the Azure portal

1. Sign into the Azure portal.
2. Navigate to the configuration page of the server you want to protect. In the security settings, select Advanced Data Security.
3. On the Advanced Data Security configuration page:

The screenshot displays the 'vanazuresqlserver - Advanced Data Security' configuration page in the Azure portal. The left-hand navigation pane includes sections for 'Overview', 'Activity log', 'Access control (IAM)', 'Tags', 'Diagnose and solve problems', 'Settings', and 'Security'. The 'Advanced Data Security' option under the 'Security' section is highlighted with a red box. The main content area shows the 'ADVANCED DATA SECURITY' toggle set to 'ON'. Below this, the 'VULNERABILITY ASSESSMENT SETTINGS' section includes fields for 'Subscription' (SQL DB Content) and 'Storage account'. The 'Periodic recurring scans' toggle is set to 'OFF'. The 'Send scan reports to' field is empty. A checkbox for 'Also send email notification to admins and subscription owners' is unchecked. The 'ADVANCED THREAT PROTECTION SETTINGS' section is highlighted with a red box and contains a 'Send alerts to' dropdown menu set to 'Email addresses' (marked with a green checkmark) and a checked checkbox for 'Also send email notification to admins and subscription owners'. The 'Advanced Threat Protection types' dropdown is set to 'All'.

4. Enable Advanced Data Security on the server.

Note: Advanced Threat Protection for Azure SQL Database detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. Advanced Threat Protection can identify Potential SQL injection, Access from unusual location or data center, Access from unfamiliar principal or potentially harmful application, and Brute force SQL credentials

Reference:

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

<https://docs.microsoft.com/en-us/azure/azure-sql/database/threat-detection-configure>

### QUESTION 11

Your company is concerned that when developers introduce open source libraries, it creates licensing compliance issues. You need to add an automated process to the build pipeline to detect when common open source libraries are added to the code base. What should you use?

- A. OWASP ZAP
- B. Jenkins
- C. Code Style
- D. WhiteSource Bolt

**Correct Answer: D**

**Section:**

**Explanation:**

WhiteSource provides WhiteSource Bolt, a lightweight open source security and management solution developed specifically for integration with Azure DevOps and Azure DevOps Server.

Note: WhiteSource is the leader in continuous open source software security and compliance management. WhiteSource integrates into your build process, irrespective of your programming languages, build tools, or development environments. It works automatically, continuously, and silently in the background, checking the security, licensing, and quality of your open source components against WhiteSource constantly-updated definitive database of open source repositories.

Note:

There are several versions of this question in the exam. The question has two possible correct answers:

1. Black Duck
2. WhiteSource Bolt

Other incorrect answer options you may see on the exam include the following:

1. Microsoft Visual SourceSafe
2. PDM
3. SourceGear

Reference:

<https://www.azuredevopslabs.com/labs/vstsextend/whitesource/>



### QUESTION 12

You plan to use a NuGet package in a project in Azure DevOps. The NuGet package is in a feed that requires authentication. You need to ensure that the project can restore the NuGet package automatically. What should the project use to automate the authentication?

- A. an Azure Automation account
- B. an Azure Artifacts Credential Provider
- C. an Azure Active Directory (Azure AD) account that has multi-factor authentication (MFA) enabled
- D. an Azure Active Directory (Azure AD) service principal

**Correct Answer: B**

**Section:**

**Explanation:**

The Azure Artifacts Credential Provider automates the acquisition of credentials needed to restore NuGet packages as part of your .NET development workflow. It integrates with MSBuild, dotnet, and NuGet(.exe) and works on Windows, Mac, and Linux. Any time you want to use packages from an Azure Artifacts feed, the Credential Provider will automatically acquire and securely store a token on behalf of the NuGet client you're using.

Reference:

<https://github.com/Microsoft/artifacts-credprovider>

### QUESTION 13

You use Azure Pipelines to manage project builds and deployments.

You plan to use Azure Pipelines for Microsoft Teams to notify the legal team when a new build is ready for release. You need to configure the Organization Settings in Azure DevOps to support Azure Pipelines for Microsoft Teams. What should you turn on?

- A. Third-party application access via OAuth
- B. Azure Active Directory Conditional Access Policy Validation
- C. Alternate authentication credentials
- D. SSH authentication

**Correct Answer: A**

**Section:**

**Explanation:**

The Azure Pipelines app uses the OAuth authentication protocol, and requires Third-party application access via OAuth for the organization to be enabled. To enable this setting, navigate to Organization Settings > Security > Policies, and set the Third-party application access via OAuth for the organization setting to On.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/integrations/microsoft-teams>

#### QUESTION 14

You have an existing project in Azure DevOps.

You plan to integrate GitHub as the repository for the project.

You need to ensure that Azure Pipelines runs under the Azure Pipelines identity.

Which authentication mechanism should you use?

- A. personal access token (PAT)
- B. GitHub App
- C. Azure Active Directory (Azure AD)
- D. OAuth

**Correct Answer: B**

**Section:**

**Explanation:**

GitHub App uses the Azure Pipelines identity.

Incorrect Answers:

A: Personal access token and OAuth use your personal GitHub identity.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/repos/github>

#### QUESTION 15

You plan to provision a self-hosted Linux agent.

Which authentication mechanism should you use to register the self-hosted agent?

- A. personal access token (PAT)
- B. SSH key
- C. Alternate credentials
- D. certificate

**Correct Answer: A**

**Section:**

**Explanation:**

Note: PAT Supported only on Azure Pipelines and TFS 2017 and newer. After you choose PAT, paste the PAT token you created into the command prompt window. Use a personal access token (PAT) if your Azure DevOps Server or TFS instance and the agent machine are not in a trusted domain. PAT authentication is handled by your Azure DevOps Server or TFS instance instead of the domain controller.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/v2-linux>

#### QUESTION 16



You are building a Microsoft ASP.NET application that requires authentication.  
You need to authenticate users by using Azure Active Directory (Azure AD).  
What should you do first?

- A. Assign an enterprise application to users and groups
- B. Create an app registration in Azure AD
- C. Configure the application to use a SAML endpoint
- D. Create a new OAuth token from the application
- E. Create a membership database in an Azure SQL database

**Correct Answer: B**

**Section:**

**Explanation:**

Register your application to use Azure Active Directory. Registering the application means that your developers can use Azure AD to authenticate users and request access to user resources such as email, calendar, and documents.

Reference: <https://docs.microsoft.com/en-us/azure/active-directory/manage-apps/developer-guidance-for-integrating-applications>

#### QUESTION 17

You have an Azure DevOps organization named Contoso.  
You need to recommend an authentication mechanism that meets the following requirements:  
Supports authentication from Git  
Minimizes the need to provide credentials during authentication  
What should you recommend?

- A. personal access tokens (PATs) in Azure DevOps
- B. Alternate credentials in Azure DevOps
- C. user accounts in Azure Active Directory (Azure AD)
- D. managed identities in Azure Active Directory (Azure AD)

**Correct Answer: A**

**Section:**

**Explanation:**

Personal access tokens (PATs) give you access to Azure DevOps and Team Foundation Server (TFS), without using your username and password directly. These tokens have an expiration date from when they're created. You can restrict the scope of the data they can access. Use PATs to authenticate if you don't already have SSH keys set up on your system or if you need to restrict the permissions that are granted by the credential. Incorrect Answers:

B: Azure DevOps no longer supports Alternate Credentials authentication since the beginning of March 2, 2020. If you're still using Alternate Credentials, we [Microsoft] strongly encourage you to switch to a more secure authentication method (for example, personal access tokens).

Reference: <https://docs.microsoft.com/en-us/azure/devops/repos/git/auth-overview>

#### QUESTION 18

You have an application that consists of several Azure App Service web apps and Azure functions.  
You need to assess the security of the web apps and the functions.  
Which Azure feature can you use to provide a recommendation for the security of the application?

- A. Security & Compliance in Azure Log Analytics
- B. Resource health in Azure Service Health
- C. Smart Detection in Azure Application Insights
- D. Compute & apps in Azure Security Center



**Correct Answer: D**

**Section:**

**Explanation:**

Monitor compute and app services: Compute & apps include the App Services tab, which App services: list of your App service environments and current security state of each. Recommendations This section has a set of recommendations for each VM and computer, web and worker roles, Azure App Service Web Apps, and Azure App Service Environment that Security Center monitors. The first column lists the recommendation. The second column shows the total number of resources that are affected by that recommendation. The third column shows the severity of the issue. Incorrect Answers:

C: Smart Detection automatically warns you of potential performance problems, not security problems in your web application.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/proactive-diagnostics>

#### QUESTION 19

Your company has a project in Azure DevOps for a new web application.

The company identifies security as one of the highest priorities.

You need to recommend a solution to minimize the likelihood that infrastructure credentials will be leaked. What should you recommend?

- A. Add a Run Inline Azure PowerShell task to the pipeline.
- B. Add a PowerShell task to the pipeline and run Set-AzureKeyVaultSecret.
- C. Add an Azure Key Vault task to the pipeline.
- D. Add Azure Key Vault references to Azure Resource Manager templates.

**Correct Answer: B**

**Section:**

**Explanation:**

Azure Key Vault provides a way to securely store credentials and other keys and secrets.

The Set-AzureKeyVaultSecret cmdlet creates or updates a secret in a key vault in Azure Key Vault.

Reference:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.keyvault/set-azurekeyvaultsecret>



#### QUESTION 20

SIMULATION

You need to ensure that an Azure web app named az400-9940427-main can retrieve secrets from an Azure key vault named az400-9940427-kv1 by using a system managed identity. The solution must use the principle of least privilege.

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

- A. See solution below.

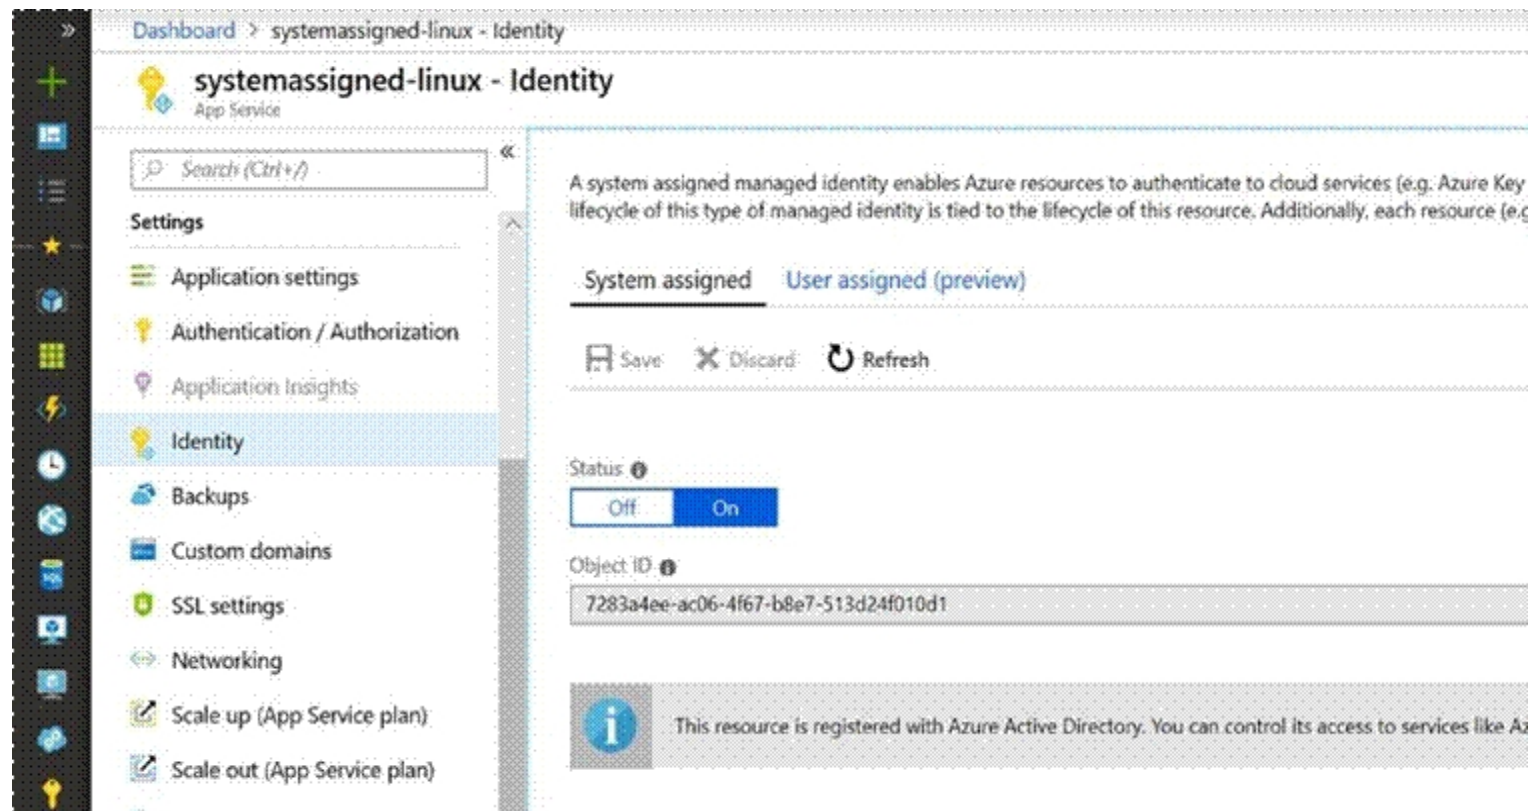
**Correct Answer: A**

**Section:**

**Explanation:**

1. In Azure portal navigate to the az400-9940427-main app.
2. Scroll down to the Settings group in the left navigation.
3. Select Managed identity.
4. Within the System assigned tab, switch Status to On. Click Save.





Reference:

<https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity>

#### QUESTION 21

You use WhiteSource Bolt to scan a Node.js application.

The WhiteSource Bolt scan identifies numerous libraries that have invalid licenses. The libraries are used only during development and are not part of a production deployment. You need to ensure that WhiteSource Bolt only scans production dependencies.

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

NOTE: Each correct selection is worth one point.

- A. Run `npm install` and specify the `--production` flag.
- B. Modify the WhiteSource Bolt policy and set the action for the licenses used by the development tools to Reassign.
- C. Modify the `devDependencies` section of the project's `Package.json` file.
- D. Configure WhiteSource Bolt to scan the `node_modules` directory only.

**Correct Answer: A, C**

**Section:**

**Explanation:**

A: To resolve NPM dependencies, you should first run "`npm install`" command on the relevant folders before executing the plugin. C: All npm packages contain a file, usually in the project root, called `package.json` - this file holds various metadata relevant to the project. This file is used to give information to npm that allows it to identify the project as well as handle the project's dependencies. It can also contain other metadata such as a project description, the version of the project in a particular distribution, license information, even configuration data - all of which can be vital to both npm and to the end users of the package.

Reference: <https://whitesource.atlassian.net/wiki/spaces/WD/pages/34209870/NPM+Plugin> <https://nodejs.org/en/knowledge/getting-started/npm/what-is-the-file-package-json>

#### QUESTION 22

You need to configure GitHub to use Azure Active Directory (Azure AD) for authentication.

What should you do first?

- A. Create a conditional access policy in Azure AD.

- B. Register GitHub in Azure AD.
- C. Create an Azure Active Directory B2C (Azure AD B2C) tenant.
- D. Modify the Security settings of the GitHub organization.

**Correct Answer: B**

**Section:**

**Explanation:**

When you connect to a Get repository from your Get client for the first time, the credential manager prompts for credentials. Provide your Microsoft account or Azure AD credentials. Note: Git Credential Managers simplify authentication with your Azure Repos Git repositories. Credential managers let you use the same credentials that you use for the Azure DevOps Services web portal. Credential managers support multi-factor authentication through Microsoft account or Azure Active Directory (Azure AD). Besides supporting multi-factor authentication with Azure Repos, credential managers also support two-factor authentication with GitHub repositories.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/set-up-credential-managers>

#### QUESTION 23

You have an Azure DevOps project named Project1 and an Azure subscription named Sub1.

You need to prevent releases from being deployed unless the releases comply with the Azure Policy rules assigned to Sub1. What should you do in the release pipeline of Project1?

- A. Add a deployment gate.
- B. Modify the Deployment queue settings.
- C. Configure a deployment trigger.
- D. Create a pipeline variable.

**Correct Answer: A**

**Section:**

**Explanation:**

You can check policy compliance with gates.

You can extend the approval process for the release by adding a gate. Gates allow you to configure automated calls to external services, where the results are used to approve or reject a deployment. You can use gates to ensure that the release meets a wide range of criteria, without requiring user intervention.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/release/deploy-using-approvals>



#### QUESTION 24

You have an Azure DevOps project that contains a build pipeline. The build pipeline uses approximately 50 open source libraries. You need to ensure that all the open source libraries comply with your company's licensing standards. Which service should you use?

- A. Ansible
- B. Maven
- C. WhiteSource Bolt
- D. Helm

**Correct Answer: C**

**Section:**

**Explanation:**

WhiteSource provides WhiteSource Bolt, a lightweight open source security and management solution developed specifically for integration with Azure DevOps and Azure DevOps Server. Note: WhiteSource is the leader in continuous open source software security and compliance management. WhiteSource integrates into your build process, irrespective of your programming languages, build tools, or development environments. It works automatically, continuously, and silently in the background, checking the security, licensing, and quality of your open source components against WhiteSource constantly-updated definitive database of open source repositories.

Note: Blackduck would also be a good answer, but it is not an option here.

Reference: <https://www.azuredevopslabs.com/labs/vstsextend/whitesource/>

### QUESTION 25

You are designing the security validation strategy for a project in Azure DevOps.

You need to identify package dependencies that have known security issues and can be resolved by an update. What should you use?

- A. Octopus Deploy
- B. Jenkins
- C. Gradle
- D. SonarQube

**Correct Answer: A**

**Section:**

**Explanation:**

Incorrect Answers:

B: Jenkins is a popular open-source automation server used to set up continuous integration and delivery (CI/CD) for your software projects. D: SonarQube is a set of static analyzers that can be used to identify areas of improvement in your code. It allows you to analyze the technical debt in your project and keep track of it in the future.

Reference:

<https://octopus.com/docs/packaging-applications>

### QUESTION 26

You administer an Azure DevOps project that includes package feeds.

You need to ensure that developers can unlist and deprecate packages. The solution must use the principle of least privilege. Which access level should you grant to the developers?

- A. Collaborator
- B. Contributor
- C. Owner



**Correct Answer: B**

**Section:**

**Explanation:**

Feeds have four levels of access: Owners, Contributors, Collaborators, and Readers. Owners can add any type of identity-individuals, teams, and groups-to any access level.

Permission	Reader	Collaborator	Contributor	Owner
List and restore/install packages	✓	✓	✓	✓
Save packages from upstream sources		✓	✓	✓
Push packages			✓	✓
Unlist/deprecate packages			✓	✓
Promote a package to a view			✓	✓
Delete/unpublish package				✓
Edit feed permissions				✓

Reference:

<https://docs.microsoft.com/en-us/azure/devops/artifacts/feeds/feed-permissions>

### QUESTION 27

Your company is concerned that when developers introduce open source libraries, it creates licensing compliance issues. You need to add an automated process to the build pipeline to detect when common open source libraries are added to the code base. What should you use?

- A. Microsoft Visual SourceSafe
- B. Code Style
- C. Black Duck
- D. Jenkins

**Correct Answer: C**

**Section:**

**Explanation:**

Secure and Manage Open Source Software

Black Duck helps organizations identify and mitigate open source security, license compliance and code-quality risks across application and container portfolios. Black Duck Hub and its plugin for Team Foundation Server (TFS) allows you to automatically find and fix open source security vulnerabilities during the build process, so you can proactively manage risk. The integration allows you to receive alerts and fail builds when any Black Duck Hub policy violations are met.

Note:

There are several versions of this question in the exam. The question has two possible correct answers:

1. Black Duck
2. WhiteSource Bolt

Other incorrect answer options you may see on the exam include the following:

1. OWASP ZAP
2. PDM
3. SourceGear

Reference:

<https://marketplace.visualstudio.com/items?itemName=black-duck-software.hub-tfs>



### QUESTION 28

You have an Azure DevOps project that contains a build pipeline. The build pipeline uses approximately 50 open source libraries. You need to ensure that all the open source libraries comply with your company's licensing standards. Which service should you use?

- A. NuGet
- B. Maven
- C. Black Duck
- D. Helm

**Correct Answer: C**

**Section:**

**Explanation:**

Secure and Manage Open Source Software

Black Duck helps organizations identify and mitigate open source security, license compliance and code-quality risks across application and container portfolios. Black Duck Hub and its plugin for Team Foundation Server (TFS) allows you to automatically find and fix open source security vulnerabilities during the build process, so you can proactively manage risk. The integration allows you to receive alerts and fail builds when any Black Duck Hub policy violations are met.

Note: WhiteSource would also be a good answer, but it is not an option here.

Reference: <https://marketplace.visualstudio.com/items?itemName=black-duck-software.hub-tfs>

### QUESTION 29

Your company develops an app for iOS. All users of the app have devices that are members of a private distribution group in Microsoft Visual Studio App Center. You plan to distribute a new release of the app.

You need to identify which certificate file you require to distribute the new release from App Center. Which file type should you upload to App Center?

- A. .cer
- B. .pfx
- C. .p12
- D. .pvk

**Correct Answer: C**

**Section:**

**Explanation:**

A successful IOS device build will produce an ipa file. In order to install the build on a device, it needs to be signed with a valid provisioning profile and certificate. To sign the builds produced from a branch, enable code signing in the configuration pane and upload a provisioning profile (.mobileprovision) and a valid certificate (.p12), along with the password for the certificate.

Reference:

<https://docs.microsoft.com/en-us/appcenter/build/xamarin/ios/>

**QUESTION 30**

**HOTSPOT**

Your company is creating a suite of three mobile applications.

You need to control access to the application builds. The solution must be managed at the organization level.

What 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**

Groups to control the build access:

- Active Directory groups
- Azure Active Directory groups
- Microsoft Visual Studio App Center distribution groups

Group type:

- Private
- Public
- Shared

**Answer Area:**

**Answer Area**

Groups to control the build access:

- Active Directory groups
- Azure Active Directory groups
- Microsoft Visual Studio App Center distribution groups

Group type:

- Private
- Public
- Shared

**Section:**

**Explanation:**

Box 1: Microsoft Visual Studio App Center distribution Groups

Distribution Groups are used to control access to releases. A Distribution Group represents a set of users that can be managed jointly and can have common access to releases. Example of Distribution Groups can be teams of users, like the QA Team or External Beta Testers or can represent stages or rings of releases, such as Staging.

Box 2: Shared

Shared distribution groups are private or public distribution groups that are shared across multiple apps in a single organization. Shared distribution groups eliminate the need to replicate distribution groups across multiple apps.

Note: With the Deploy with App Center Task in Visual Studio Team Services, you can deploy your apps from Azure DevOps (formerly known as VSTS) to App Center. By deploying to App Center, you will be able to distribute your builds to your users.

References: <https://docs.microsoft.com/en-us/appcenter/distribution/groups>

**QUESTION 31**

DRAG DROP

You use GitHub Enterprise Server as a source code repository.

You create an Azure DevOps organization named Contoso.

In the Contoso organization, you create a project named Project1.

You need to link GitHub commits, pull requests, and issues to the work items of Project1. The solution must use OAuth-based authentication.

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:**

**Actions**

From Developer settings in GitHub Enterprise Server, register a new OAuth app.

From Project Settings in Azure DevOps, create a service hook subscription.

From Organization settings in Azure DevOps, connect to Azure Active Directory (Azure AD).

From Project Settings in Azure DevOps, add a GitHub connection.

From Organization settings in Azure DevOps, add an OAuth configuration.

From Developer settings in GitHub Enterprise Server, generate a private key.

**Answer Area**



Correct Answer:

**Actions**

From Project Settings in Azure DevOps, create a service hook subscription.

From Organization settings in Azure DevOps, connect to Azure Active Directory (Azure AD).

From Developer settings in GitHub Enterprise Server, generate a private key.

**Answer Area**

From Developer settings in GitHub Enterprise Server, register a new OAuth app.

From Organization settings in Azure DevOps, add an OAuth configuration.

From Project Settings in Azure DevOps, add a GitHub connection.



**Section:**

**Explanation:**

Step 1: From Developer settings in GitHub Enterprise Server, register a new OAuth app.

If you plan to use OAuth to connect Azure DevOps Services or Azure DevOps Server with your GitHub Enterprise Server, you first need to register the application as an OAuth App

Step 2: Organization settings in Azure DevOps, add an OAuth configuration

Register your OAuth configuration in Azure DevOps Services.

Note:

1. Sign into the web portal for Azure DevOps Services.
2. Add the GitHub Enterprise Oauth configuration to your organization.
3. Open Organization settings>Oauth configurations, and choose Add Oauth configuration.
4. Fill in the form that appears, and then choose Create.

Step 3: From Project Settings in Azure DevOps, add a GitHub connection.

Connect Azure DevOps Services to GitHub Enterprise Server

Choose the Azure DevOps logo to open Projects, and then choose the Azure Boards project you want to configure to connect to your GitHub Enterprise repositories.

Choose (1) Project Settings, choose (2) GitHub connections and then (3) Click here to connect to your GitHub Enterprise organization.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/boards/github/connect-to-github>

### QUESTION 32

DRAG DROP

You are configuring an Azure DevOps deployment pipeline. The deployed application will authenticate to a web service by using a secret stored in an Azure key vault.

You need to use the secret in the deployment pipeline.

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:

Actions

Create a service principal in Azure Active Directory (Azure AD).

Add an app registration in Azure Active Directory (Azure AD).

Configure an access policy in the key vault.


Generate a self-signed certificate.

Add an Azure Resource Manager service connection to the pipeline.

Export a certificate from the key vault.

Answer Area

The logo for 'Vdumps' features a stylized orange 'V' followed by the word 'dumps' in a grey, lowercase, sans-serif font.



Correct Answer:



### Actions

- 
- Add an app registration in Azure Active Directory (Azure AD).
- 
- Generate a self-signed certificate.
- 
- Export a certificate from the key vault.

### Answer Area

- Create a service principal in Azure Active Directory (Azure AD)
- Configure an access policy in the key vault.
- Add an Azure Resource Manager service connection to the pipeline.



### Section:

#### Explanation:

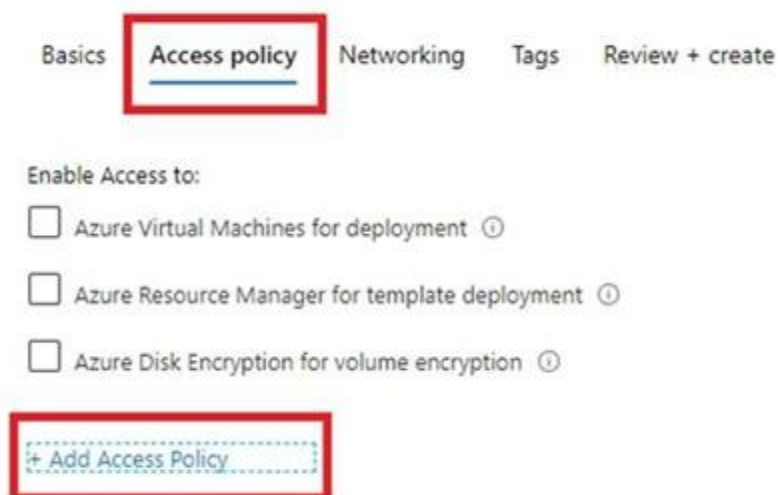
Step 1: Create a service principal in Azure Active Directory (Azure AD).

You will need a service principal to deploy an app to an Azure resource from Azure Pipelines.

Step 2: Configure an access policy in the key vault.

You need to secure access to your key vaults by allowing only authorized applications and users. To access the data from the vault, you will need to provide read (Get) permissions to the service principal that you will be using for authentication in the pipeline.

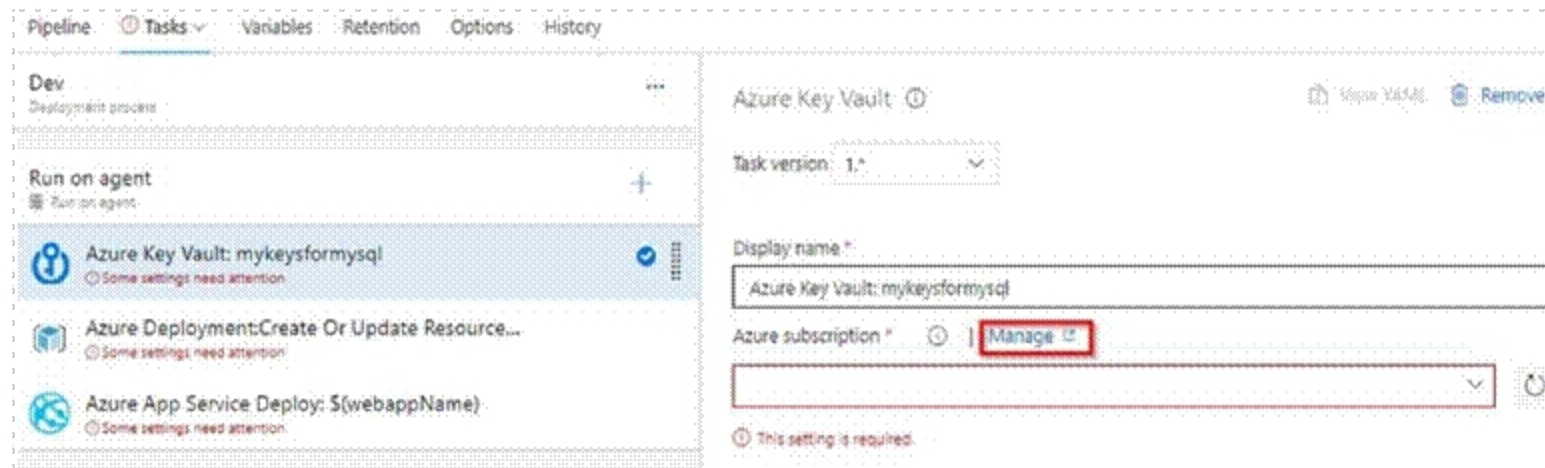
Select Access policy and then select + Add Access Policy to setup a new policy.



Step 3: Add an Azure Resource Manager service connection to the pipeline

You need to authorize the pipeline to deploy to Azure:

1. Select Pipelines | Pipelines,
2. Go to Releases under Pipelines and then select and Edit your pipeline.
3. Under Tasks, notice the release definition for Dev stage has a Azure Key Vault task. This task downloads Secrets from an Azure Key Vault. You will need to point to the subscription and the Azure Key Vault resource.
4. Click Manage, this will redirect to the Service connections page.



5. Click on New Service connection -> Azure Resource Manager -> Service Principal (manual). Fill the information from previously created service principal.

Reference:

<https://azuredevopslabs.com/labs/vstsextend/azurekeyvault/>

### QUESTION 33

DRAG DROP

You have a private project in Azure DevOps and two users named User1 and User2.

You need to add User1 and User2 to groups to meet the following requirements:

User1 must be able to create a code wiki.

User2 must be able to edit wiki pages.

The solution must use the principle of least privilege.

To which group should you add each user? To answer, drag the appropriate groups to the correct users. Each group 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:

#### Groups

- Build Administrators
- Contributors
- Project Administrators
- Project Valid Users
- Stakeholders

#### Answer Area

- User1:
- User2:

Correct Answer:

## Groups

Build Administrators
Project Valid Users
Stakeholders

## Answer Area

User1:	Project Administrators
User2:	Contributors

### Section:

### Explanation:

User1: Project Administrators

You must have the permission Create Repository to publish code as wiki. By default, this permission is set for members of the Project Administrators group.

User2: Contributors

Anyone who is a member of the Contributors security group can add or edit wiki pages.

Anyone with access to the team project, including stakeholders, can view the wiki.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/project/wiki/wiki-create-repo>

### QUESTION 34

#### HOTSPOT

Your company has an Azure subscription.

The company requires that all resource group in the subscription have a tag named organization set to a value of Contoso.

You need to implement a policy to meet the tagging requirement.

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

NOTE: Each correct selection is worth one point.

### Hot Area:



```
{
  "policyRule": {
    "if": {
      "allOf": [
        {
          "field": "type",
          "equals":
            ,
          {
            "MicrosoftResources/deployments"
            "MicrosoftResources/subscriptions"
            "MicrosoftResources/subscriptions/resourceGroups"
          }
          "not": {
            "field": "tags['organization']",
            "equals": "Contoso"
          }
        }
      ]
    },
    "then": {
      "effect":
        ,
      "details": [
        {
          "field": "tags['organization']",
          "value": "Contoso"
        }
      ]
    }
  }
}
```

Dropdown menu with options: "MicrosoftResources/deployments", "MicrosoftResources/subscriptions", "MicrosoftResources/subscriptions/resourceGroups"

Dropdown menu with options: "Append", "Deny", "DeployIfNotExists"



Answer Area:

```

{
  "policyRule": {
    "if": {
      "allOf": [
        {
          "field": "type",
          "equals":
          ,
          {
            "MicrosoftResources/deployments"
            "MicrosoftResources/subscriptions"
            "MicrosoftResources/subscriptions/resourceGroups"
          }
        }
      ]
    },
    "not": {
      "field": "tags['organization']",
      "equals": "Contoso"
    }
  }
},
"then": {
  "effect":
  ,
  "details": [
    {
      "field": "tags['organization']",
      "value": "Contoso"
    }
  ]
}
}

```



**Section:**

**Explanation:**

Box 1: " Microsoft.Resources/subscriptions/resourceGroups"

Box 2: "Deny",

Sample - Enforce tag and its value on resource groups

```

},
"policyRule": {
  "if": {
    "allOf": [
      {
        "field": "type",
        "equals": "Microsoft.Resources/subscriptions/resourceGroups"
      },
      {
        "not": {
          "field": "[concat('tags[',parameters('tagName'), ']')]",

```

```
"equals": "[parameters('tagValue')]"
}
}
],
},
"then": {
"effect": "deny"
}
}
}
}
```

References:

<https://docs.microsoft.com/en-us/azure/governance/policy/samples/enforce-tag-on-resource-groups>

### QUESTION 35

#### SIMULATION

You need to prepare a network security group (NSG) named az400-9940427-nsg1 to host an Azure DevOps pipeline agent. The solution must allow only the required outbound port for Azure DevOps and deny all other inbound and outbound access to the Internet.

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


A. See solution below.

**Correct Answer: A**

**Section:**

**Explanation:**

1. Open Microsoft Azure Portal and Log into your Azure account.
2. Select network security group (NSG) named az400-9940427-nsg1
3. Select Settings, Outbound security rules, and click Add
4. Click Advanced



5. Change the following settings:

Destination Port range: 8080

Protocol: TCP

Action: Allow

Note: By default, Azure DevOps Server uses TCP Port 8080.

Reference:

<https://robertsmit.wordpress.com/2017/09/11/step-by-step-azure-network-security-groups-nsg-security-center-azure-nsg-network/>


<https://docs.microsoft.com/en-us/azure/devops/server/architecture/required-ports?view=azure-devops>

### QUESTION 36

#### HOTSPOT


You have a project in Azure DevOps that has three teams as shown in the Teams exhibit. (Click the Teams tab.)





You create a new dashboard named Dash1.

You configure the dashboard permissions for the Control project as shown in the Permissions exhibit. (Click the Permissions tab.)



All other permissions have the default values set.

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
Web Team can delete Dash1.	<input type="radio"/>	<input type="radio"/>
Contoso Team can view Dash1.	<input type="radio"/>	<input type="radio"/>
Project administrators can create new dashboards.	<input type="radio"/>	<input type="radio"/>

Answer Area:

## Answer Area

Statements	Yes	No
Web Team can delete Dash1.	<input type="radio"/>	<input checked="" type="radio"/>
Contoso Team can view Dash1.	<input checked="" type="radio"/>	<input type="radio"/>
Project administrators can create new dashboards.	<input checked="" type="radio"/>	<input type="radio"/>

**Section:**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/devops/report/dashboards/charts-dashboard-permissions-access>

### QUESTION 37

DRAG DROP

Your company has a project in Azure DevOps.

You plan to create a release pipeline that will deploy resources by using Azure Resource Manager templates. The templates will reference secrets stored in Azure Key Vault.

You need to recommend a solution for accessing the secrets stored in the key vault during deployments. The solution must use the principle of least privilege.

What should you include in the recommendation? To answer, drag the appropriate configurations to the correct targets. Each configuration 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:**



### Configurations

### Answer Area

A Key Vault access policy

Enable key vaults for template deployment by using:

A Key Vault advanced access policy

Restrict access to the secrets in Key Vault by using:

RBAC

Correct Answer:

### Configurations

### Answer Area

A Key Vault advanced access policy

Enable key vaults for template deployment by using:

RBAC

Restrict access to the secrets in Key Vault by using:

A Key Vault access policy



Section:

Explanation:

<https://docs.microsoft.com/en-us/azure/key-vault/general/secure-your-key-vault>

#### QUESTION 38

DRAG DROP

You need to configure access to Azure DevOps agent pools to meet the following requirements:

Use a project agent pool when authoring build or release pipelines.

View the agent pool and agents of the organization.

Use the principle of least privilege.

Which role memberships are required for the Azure DevOps organization and the project? To answer, drag the appropriate role memberships to the correct targets. Each role membership 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:

**Roles**

- Administrator
- Reader
- Service Account
- User

**Answer Area**

Organization:

Project:

**Correct Answer:****Roles**

- Administrator
- 
- Service Account
- 

**Answer Area**

Organization:

Project:

**Section:****Explanation:**

References: <https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/pools-queues?view=azure-devops&tabs=yaml%2Cbrowser>

**QUESTION 39**

You create a Microsoft ASP.NET Core application.

You plan to use Azure Key Vault to provide secrets to the application as configuration data.

You need to create a Key Vault access policy to assign secret permissions to the application. The solution must use the principle of least privilege. Which secret permissions should you use?

- A. List only
- B. Get only
- C. Get and List

**Correct Answer: B****Section:****Explanation:**

Application data plane permissions:

Keys: sign

Secrets: get

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/key-vault-secure-your-key-vault>

**QUESTION 40**

You have a branch policy in a project in Azure DevOps. The policy requires that code always builds successfully. You need to ensure that a specific user can always merge changes to the master branch, even if the code fails to

compile. The solution must use the principle of least privilege. What should you do?

- A. Add the user to the Build Administrators group.
- B. Add the user to the Project Administrators group.
- C. From the Security settings of the repository, modify the access control for the user.
- D. From the Security settings of the branch, modify the access control for the user.

**Correct Answer: D**

**Section:**

**Explanation:**

In some cases, you need to bypass policy requirements so you can push changes to the branch directly or complete a pull request even if branch policies are not satisfied. For these situations, grant the desired permission from the previous list to a user or group. You can scope this permission to an entire project, a repo, or a single branch. Manage this permission along the with other Git permissions.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/branch-policies>

#### QUESTION 41

Your company uses Azure DevOps.

Only users who have accounts in Azure Active Directory can access the Azure DevOps environment.

You need to ensure that only devices that are connected to the on-premises network can access the Azure DevOps environment. What should you do?

- A. Assign the Stakeholder access level to all users.
- B. In Azure Active Directory, configure risky sign-ins.
- C. In Azure DevOps, configure Security in Project Settings.
- D. In Azure Active Directory, configure conditional access.

**Correct Answer: D**

**Section:**

**Explanation:**

Conditional Access is a capability of Azure Active Directory. With Conditional Access, you can implement automated access control decisions for accessing your cloud apps that are based on conditions. Conditional Access policies are enforced after the first-factor authentication has been completed.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/overview>

#### QUESTION 42

You have the following Azure policy.



```

if: {
  allof: [
    {
      "field": "type",
      "equals": "Microsoft.Storage/storageAccounts"
    },
    {
      "field": "Microsoft.Storage/storageAccounts/supportsHttpsTrafficOnly",
      "notEquals": "true"
    }
  ]
},
then: {
  effect: "deny"
}

```

You assign the policy to the Tenant root group.  
What is the effect of the policy?

- A. prevents all HTTP traffic to existing Azure Storage accounts
- B. ensures that all traffic to new Azure Storage accounts is encrypted
- C. prevents HTTPS traffic to new Azure Storage accounts when the accounts are accessed over the Internet
- D. ensures that all data for new Azure Storage accounts is encrypted at rest

**Correct Answer: B**

**Section:**

**Explanation:**

Denies non HTTPS traffic.



#### QUESTION 43

You have an Azure DevOps organization named Contoso, an Azure DevOps project named Project1, an Azure subscription named Sub1, and an Azure key vault named vault1. You need to ensure that you can reference the values of the secrets stored in vault1 in all the pipelines of Project1. The solution must prevent the values from being stored in the pipelines. What should you do?

- A. Create a variable group in Project1.
- B. Add a secure file to Project1.
- C. Modify the security settings of the pipelines.
- D. Configure the security policy of Contoso.

**Correct Answer: A**

**Section:**

**Explanation:**

Use a variable group to store values that you want to control and make available across multiple pipelines.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/library/variable-groups>

#### QUESTION 44

DRAG DROP

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.

Select and Place:

### Answer Area

#### Objects

a system-assigned managed identity

a cluster

an application registration in contoso.com

an RBAC binding

Correct Answer:

### Answer Area

#### Objects

an application registration in contoso.com

a cluster

a system-assigned managed identity

an RBAC binding

#### Section:

#### Explanation:

Step 1: Create an AKS cluster

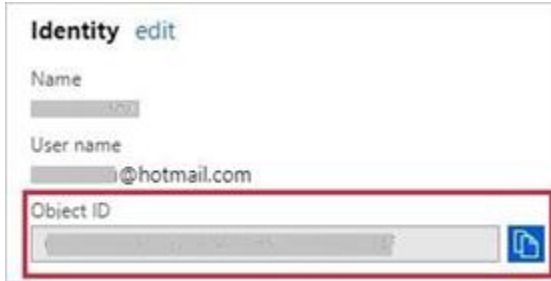
Step 2: a system-assigned managed identity

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.



The screenshot shows the 'Identity' section of an Azure DevOps account. It includes fields for 'Name', 'User name' (with a value ending in '@hotmail.com'), and 'Object ID'. The 'Object ID' field is highlighted with a red rectangular box.

Step 3: a RBAC binding

Reference:

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

#### QUESTION 45

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:

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.

Hot Area:

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 Area:

## 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	

### Section:

#### 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>



## 02 - Develop a security and compliance plan

### 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.

### Overview

Litware, Inc. is an independent software vendor (ISV). Litware has a main office and five branch offices.

### Existing Environment

#### Application Architecture

The company's primary application is a single monolithic retirement fund management system based on ASP.NET web forms that use logic written in VB.NET. Some new sections of the application are written in C#.

Variations of the application are created for individual customers. Currently, there are more than 80 live code branches in the application's code base.

The application was developed by using Microsoft Visual Studio. Source code is stored in Team Foundation Server (TFS) in the main office. The branch offices access the source code by using TFS proxy servers.

#### Architectural Issues

Litware focuses on writing new code for customers. No resources are provided to refactor or remove existing code. Changes to the code base take a long time, as dependencies are not obvious to individual developers.

Merge operations of the code often take months and involve many developers. Code merging frequently introduces bugs that are difficult to locate and resolve.

Customers report that ownership costs of the retirement fund management system increase continually. The need to merge unrelated code makes even minor code changes expensive.

Customers report that bug reporting is overly complex.

## Requirements

### Planned Changes

Litware plans to develop a new suite of applications for investment planning. The investment planning applications will require only minor integration with the existing retirement fund management system.

The investment planning applications suite will include one multi-tier web application and two iOS mobile applications. One mobile application will be used by employees; the other will be used by customers.

Litware plans to move to a more agile development methodology. Shared code will be extracted into a series of packages.

Litware has started an internal cloud transformation process and plans to use cloud-based services whenever suitable.

Litware wants to become proactive in detecting failures, rather than always waiting for customer bug reports.

### Technical Requirements

The company's investment planning applications suite must meet the following requirements:

New incoming connections through the firewall must be minimized.

Members of a group named Developers must be able to install packages.

The principle of least privilege must be used for all permission assignments.

A branching strategy that supports developing new functionality in isolation must be used.

Members of a group named Team Leaders must be able to create new packages and edit the permissions of package feeds. Visual Studio App Center must be used to centralize the reporting of mobile application crashes and device types in use. By default, all releases must remain available for 30 days, except for production releases, which must be kept for 60 days. Code quality and release quality are critical. During release, deployments must not proceed between stages if any active bugs are logged against the release. The mobile applications must be able to call the share pricing service of the existing retirement fund management system. Until the system is upgraded, the service will only support basic authentication over HTTPS. The required operating system configuration for the test servers changes weekly. Azure Automation State Configuration must be used to ensure that the operating system on each test server is configured the same way when the servers are created and checked periodically.

### Current Technical Issue

The test servers are configured correctly when first deployed, but they experience configuration drift over time. Azure Automation State Configuration fails to correct the configurations.

Azure Automation State Configuration nodes are registered by using the following command.

```
Register-AzureRmAutomationDscNode
-ResourceGroupName 'TestResourceGroup'
-AutomationAccountName 'LitwareAutomationAccount'
-AzureVMName $vmname
-ConfigurationMode 'ApplyOnly'
```



## QUESTION 1

### DRAG DROP

Which package feed access levels should be assigned to the Developers and Team Leaders groups for the investment planning applications suite? To answer, drag the appropriate access levels to the correct groups. Each access 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:



**Access Levels**

- Collaborator
- Contributor
- Owner
- Reader

**Answer Area**

Developers:

Team Leaders:

**Correct Answer:**

**Access Levels**

- Collaborator
- Contributor
- 
- 

**Answer Area**

Developers:

Team Leaders:

**Section:**

**Explanation:**

Box 1: Reader

Members of a group named Developers must be able to install packages.

Feeds have four levels of access: Owners, Contributors, Collaborators, and Readers. Owners can add any type of identity-individuals, teams, and groups-to any access level.

Box 2: Owner

Members of a group named Team Leaders must be able to create new packages and edit the permissions of package feeds.

Permission	Reader	Collaborator	Contributor	Owner
List and restore/install packages	✓	✓	✓	✓
Save packages from upstream sources		✓	✓	✓
Push packages			✓	✓
Unlist/deprecate packages			✓	✓
Delete/unpublish package				✓
Edit feed permissions				✓
Rename and delete feed				✓

**QUESTION 2**

**HOTSPOT**

How should you configure the release retention policy for the investment planning depletions suite? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**



Required secrets:  ▼

Certificate
Personal access token
Shared Access Authorization token
Username and password

Storage location:  ▼

Azure Data Lake
Azure Key Vault
Azure Storage with HTTPS access
Azure Storage with HTTP access

Answer Area:

## Answer Area

Required secrets:

Certificate
Personal access token
Shared Access Authorization token
Username and password

Storage location:

Azure Data Lake
Azure Key Vault
Azure Storage with HTTPS access
Azure Storage with HTTP access

### Section:

#### Explanation:

Box 1: Shared Access Authorization token

Every request made against a storage service must be authorized, unless the request is for a blob or container resource that has been made available for public or signed access. One option for authorizing a request is by using Shared Key.

Box 2: Azure Storage with HTTPS access

Scenario: The mobile applications must be able to call the share pricing service of the existing retirement fund management system. Until the system is upgraded, the service will only support basic authentication over HTTPS. The investment planning application suite will include one multi-tier web application and two iOS mobile application. One mobile application will be used by employees; the other will be used by customers.

Reference:

<https://docs.microsoft.com/en-us/rest/api/storageservices/authorize-with-shared-key>

### QUESTION 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 question on 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 sections 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 on 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.

##### Application Architecture

The company's primary application is a single monolithic retirement fund management system based on ASP.NET web forms that use logic written in VB.NET. Some new sections of the application are written in C#.

Variations of the application are created for individual customers. Currently, there are more than 80 live code branches in the application's code base.

The application was developed by using Microsoft Visual Studio. Source code is stored in Team Foundation Server (TFS) in the main office. The branch offices access the source code by using TFS proxy servers.

##### Architectural Issues

Litware focuses on writing new code for customers. No resources are provided to refactor or remove existing code. Changes to the code base take a long time, as dependencies are not obvious to individual developers.

Merge operations of the code often take months and involve many developers. Code merging frequently introduces bugs that are difficult to locate and resolve.

Customers report that ownership costs of the retirement fund management system increase continually. The need to merge unrelated code makes even minor code changes expensive.

Customers report that bug reporting is overly complex.

#### Planned changes

Litware plans to develop a new suite of applications for investment planning. The investment planning applications will require only minor integration with the existing retirement fund management system.

The investment planning applications suite will include one multi-tier web application and two iOS mobile application. One mobile application will be used by employees; the other will be used by customers.

Litware plans to move to a more agile development methodology. Shared code will be extracted into a series of packages.

Litware has started an internal cloud transformation process and plans to use cloud-based services whenever suitable.

Litware wants to become proactive in detecting failures, rather than always waiting for customer bug reports.

#### Technical requirements

The company's investment planning applications suite must meet the following requirements:

New incoming connections through the firewall must be minimized.

Members of a group named Developers must be able to install packages.

The principle of least privilege must be used for all permission assignments.

A branching strategy that supports developing new functionality in isolation must be used.

Members of a group named Team Leaders must be able to create new packages and edit the permissions of package feeds. Visual Studio App Center must be used to centralize the reporting of mobile application crashes and device types in use. By default, all releases must remain available for 30 days, except for production releases, which must be kept for 60 days. Code quality and release quality are critical. During release, deployments must not proceed between stages if any active bugs are logged against the release. The mobile applications must be able to call the share pricing service of the existing retirement fund management system. Until the system is upgraded, the service will only support basic authentication over HTTPS. The required operating system configuration for the test servers changes weekly. Azure Automation State Configuration must be used to ensure that the operating system on each test server is configured the same way when the servers are created and checked periodically.

#### Current Technical Issue

The test servers are configured correctly when first deployed, but they experience configuration drift over time. Azure Automation State Configuration fails to correct the configurations.

Azure Automation State Configuration nodes are registered by using the following command.

```
Register-AzureRmAutomationDscNode  
-ResourceGroupName 'TestResourceGroup'  
-AutomationAccountName 'LitwareAutomationAccount'  
-AzureVMName $vmname  
-ConfigurationMode 'ApplyOnly'
```

You need to configure a cloud service to store the secrets required by the mobile applications to call the share pricing service. What should you include in the solution? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

#### Answer Area

Required secrets:	<ul style="list-style-type: none"><li>Certificate</li><li>Personal access token</li><li>Shared Access Authorization token</li><li>Username and password</li></ul>
Storage location:	<ul style="list-style-type: none"><li>Azure Data Lake</li><li>Azure Key Vault</li><li>Azure Storage with HTTP access</li><li>Azure Storage with HTTPS access</li></ul>

Hot Area:

Answer:

**Answer Area**

Required secrets:  ▼

Certificate
Personal access token
Shared Access Authorization token
Username and password

Storage location:  ▼

Azure Data Lake
Azure Key Vault
Azure Storage with HTTP access
Azure Storage with HTTPS access

**Hot Area:**

**Answer Area**

Required secrets:  ▼

Certificate
Personal access token
Shared Access Authorization token
Username and password

Storage location:  ▼

Azure Data Lake
Azure Key Vault
Azure Storage with HTTP access
Azure Storage with HTTPS access

**Answer Area:**



## Answer Area

Required secrets:

Certificate
Personal access token
Shared Access Authorization token
Username and password

Storage location:

Azure Data Lake
Azure Key Vault
Azure Storage with HTTP access
Azure Storage with HTTPS access

### Section:

#### Explanation:

Every request made against a storage service must be authorized, unless the request is for a blob or container resource that has been made available for public or signed access. One option for authorizing a request is by using Shared Key. Scenario: The mobile applications must be able to call the share pricing service of the existing retirement fund management system. Until the system is upgraded, the service will only support basic authentication over HTTPS. The investment planning applications suite will include one multi-tier web application and two iOS mobile application. One mobile application will be used by employees; the other will be used by customers.

References: <https://docs.microsoft.com/en-us/rest/api/storageservices/authorize-with-shared-key>

### 03 - Develop a security and compliance plan

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.

Overview

General Overview

Woodgrove Bank is a financial services company that has a main office in the United Kingdom.

Technical Requirements and Planned Changes

Planned Changes

Woodgrove Bank plans to implement the following project management changes:

Implement Azure DevOps for project tracking.

Centralize source code control in private GitHub repositories.

Implement Azure Pipelines for build pipelines and release pipelines.

Woodgrove Bank plans to implement the following changes to the identity environment:

Deploy an Azure AD tenant named woodgrovebank.com.

Sync the Active Directory domain to Azure AD.

Configure App1 to use a service principal.

Integrate GitHub with Azure AD.

Woodgrove Bank plans to implement the following changes to the core apps:

Migrate App1 to ASP.NET Core.  
 Integrate Azure Pipelines and the third-party build tool used to develop App2.  
 Woodgrove Bank plans to implement the following changes to the DevOps environment:  
 Deploy App1 to Azure App Service.  
 Implement source control for the DB1 schema.  
 Migrate all the source code from TFS1 to GitHub.  
 Deploy App2 to an Azure virtual machine named VM1.  
 Merge the POC branch into the GitHub default branch.  
 Implement an Azure DevOps dashboard for stakeholders to monitor development progress.

**Technical Requirements**

Woodgrove Bank identifies the following technical requirements:  
 The initial databases for new environments must contain both schema and reference data.  
 An Azure Monitor alert for VM1 must be configured to meet the following requirements:  
 - Be triggered when average CPU usage exceeds 80 percent for 15 minutes.  
 - Calculate CPU usage averages once every minute.  
 • The commit history of the POC branch must replace the history of the default branch.  
 The commit history of the POC branch must replace the history of the default branch.  
 The Azure DevOps dashboard must display the metrics shown in the following table.

Number	Required data
1	A comparison between the work the development team planned to deliver and what was delivered
2	The status of the environments in a release definition
3	The total number of results from a work item query

Access to Azure DevOps must be restricted to specific IP addresses.  
 Page load times for App1 must be captured and monitored.  
 Administrative effort must be minimized.

**QUESTION 1**

DRAG DROP

You need to configure authentication for App1. The solution must support the planned changes.

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:**



**Actions Commands Cmdlets Statements**

**Answer Area**

Create an app.

Add a secret.

Create a credential.

Configure the ID and secret for App1.

Create a managed service identity.



Correct Answer:

**Actions Commands Cmdlets Statements**

**Answer Area**

Add a secret.

Create a credential.



Create an app.

Create a managed service identity.

Configure the ID and secret for App1.



**Section:**

**Explanation:**

Woodgrove Bank plans to implement the following changes to the identity environment:

Configure App1 to use a service principal.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal>

**QUESTION 2**

You need to meet the technical requirements for controlling access to Azure DevOps.

What should you use?



- A. Azure Multi-Factor Authentication (MFA)
- B. on-premises firewall rules
- C. conditional access policies in Azure AD
- D. Azure role-based access control (Azure RBAC)

**Correct Answer: C**

**Section:**

**Explanation:**

Scenario: Access to Azure DevOps must be restricted to specific IP addresses.

Azure DevOps is authenticated through Azure Active Directory. You can use Azure AD's conditional access to prevent logins from certain geographies and address ranges.

Reference:

<https://www.rebeladmin.com/2018/08/step-step-guide-configure-location-based-conditional-access-policies/>

### QUESTION 3

You need to configure Azure Pipelines to control App2 builds.

Which authentication method should you use?

- A. Windows NTLM
- B. certificate
- C. SAML
- D. personal access token (PAT)

**Correct Answer: D**

**Section:**

**Explanation:**

Scenario: Deploy App2 to an Azure virtual machine named VM1.

A personal access token (PAT) is used as an alternate password to authenticate into Azure DevOps.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/organizations/accounts/use-personal-access-tokens-to-authenticate>



## 01 - Define and implement continuous integration

### QUESTION 1

I need more hosted build resources. What can I do?

- A. The Azure Pipelines pool provides all Azure DevOps organizations
- B. Host your own agents on infrastructure that you manage.
- C. Buy additional parallel jobs.
- D.

**Correct Answer: A**

**Section:**

**Explanation:**

with cloud-hosted build agents and free build minutes each month. If you need more Microsoft- hosted build resources, or need to run more jobs in parallel, then you can either:

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/pools-queues>

### QUESTION 2

DRAG DROP

Your company has four projects. The version control requirements for each project are shown in the following table.

Project	Requirement
Project 1	Project leads must be able to restrict access to individual files and folders in the repository.
Project 2	The version control system must enforce the following rules before merging any changes to the main branch: <ul style="list-style-type: none"> <li>• Changes must be reviewed by at least two project members.</li> <li>• Changes must be associated to at least one work team.</li> </ul>
Project 3	The project members must be able to work in Azure Repos directly from Xcode.
Project 4	The release branch must only be viewable or editable by the project leads.

You plan to use Azure Repos for all the projects.

Which version control system should you use for each project? To answer, drag the appropriate version control systems to the correct projects. Each version control system 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:**

Version Control Systems	Answer Area
<input type="text" value="Git"/>	Project 1: <input type="text"/>
<input type="text" value="Perforce"/>	Project 2: <input type="text"/>
<input type="text" value="Subversion"/>	Project 3: <input type="text"/>
<input type="text" value="Team Foundation Version Control"/>	Project 4: <input type="text"/>



**Correct Answer:**

Version Control Systems	Answer Area
<input type="text" value="Git"/>	Project 1: <input type="text" value="Team Foundation Version Control"/>
<input type="text" value="Perforce"/>	Project 2: <input type="text" value="Git"/>
<input type="text" value="Subversion"/>	Project 3: <input type="text" value="Subversion"/>
<input type="text" value="Team Foundation Version Control"/>	Project 4: <input type="text" value="Git"/>

**Section:**

**Explanation:**

Box 1: Team Foundation Version Control

TFVC lets you apply granular permissions and restrict access down to a file level.

Box 2: Git

get is the default version control provider for new projects. You should use get for version control in your projects unless you have a specific need for centralized version control features in TFVC.

Box 3: Subversion

Note: Xcode is an integrated development environment (IDE) for macOS containing a suite of software development tools developed by Apple

Box 4: Git

Note: Perforce: Due to its multitenant nature, many groups can work on versioned files. The server tracks changes in a central database of MD5 hashes of file content, along with descriptive meta data and separately retains a master repository of file versions that can be verified through the hashes.

References:

<https://searchitoperations.techtarget.com/definition/Perforce-Software>

<https://docs.microsoft.com/en-us/azure/devops/repos/git/share-your-code-in-git-xcode>

<https://docs.microsoft.com/en-us/azure/devops/repos/tfvc/overview>

**QUESTION 3**

**HOTSPOT**

You 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 displays the Azure DevOps web portal interface. On the left is a navigation sidebar with options like Overview, Boards, Repos, Pipelines, Environments, Releases, Library, Task groups, Deployment groups, WhiteSource Bolt, Test Plans, and Artifacts. The main area shows the results for a pipeline run titled 'Jobs in run #20191120.1' under the 'Fast Track' pipeline. The pipeline consists of several stages: 'build vm', 'deploy\_to\_dev', 'deploy\_to\_uat', and 'Finalize build'. Each stage contains one or more tasks, all of which are marked as successful with green checkmarks. A detailed view of the 'initial\_build' task is shown in a dark overlay, listing its 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.

**Hot Area:**

**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 Area:

**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	



Section:

Explanation:

Reference:

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

QUESTION 4

DRAG DROP

You are configuring Azure DevOps build pipelines.

You plan to use hosted build agents.

Which build agent pool should you use to compile each application type? To answer, drag the appropriate build agent pools to the correct application types. Each build agent pool 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:

**Build Agent Pools**

Hosted Windows Container

Hosted Ubuntu 1604

Hosted macOS

Hosted

Default

**Answer Area**

An application that runs on iOS:

An Internet Information Services (IIS) web application that runs in Docker:

Correct Answer:

**Build Agent Pools**

Hosted Windows Container

Hosted Ubuntu 1604

Default

**Answer Area**

An application that runs on iOS:

An Internet Information Services (IIS) web application that runs in Docker:



**Section:**

**Explanation:**

- Hosted macOS
- Hosted Windows container

<https://devblogs.microsoft.com/devops/removing-older-images-in-azure-pipelines-hosted-pools/>

<https://docs.microsoft.com/en-us/azure/devops/release-notes/2019/sprint-154-update#single-hosted-pool> <https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/pools-queues?view=azure-devops&tabs=yaml%2Cbrowser>

**QUESTION 5**

You store source code in a Get repository in Azure Repos. You use a third-party continuous integration (CI) tool to control builds. What will Azure DevOps use to authenticate with the tool?

- A. certificate authentication
- B. a personal access token (PAT)
- C. a Shared Access Signature (SAS) token
- D. NTLM authentication

**Correct Answer: B**

**Section:**

**Explanation:**

Personal access tokens (PATs) give you access to Azure DevOps and Team Foundation Server (TFS), without using your username and password directly.

Reference: <https://docs.microsoft.com/en-us/azure/devops/repos/git/auth-overview>

#### 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 need to recommend an integration strategy for the build process of a Java application. The solution must meet the following requirements:

The builds must access an on-premises dependency management system.

The build outputs must be stored as Server artifacts in Azure DevOps.

The source code must be stored in a Get repository in Azure DevOps.

Solution: Configure an Octopus Tentacle on an on-premises machine. Use the Package Application task in the build pipeline. Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section:**

**Explanation:**

Octopus Deploy is an automated deployment server that makes it easy to automate deployment of ASP.NET web applications, Java applications, NodeJS application and custom scripts to multiple environments. Octopus can be installed on various platforms including Windows, Mac and Linux. It can also be integrated with most version control tools including VSTS and GIT. When you deploy software to Windows servers, you need to install Tentacle, a lightweight agent service, on your Windows servers so they can communicate with the Octopus server. When defining your deployment process, the most common step type will be a package step. This step deploys your packaged application onto one or more deployment targets. When deploying a package you will need to select the machine role that the package will be deployed to.

Reference:

<https://octopus.com/docs/deployment-examples/package-deployments> <https://explore.emtecinc.com/blog/octopus-for-automated-deployment-in-devops-models>

#### QUESTION 7

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 need to recommend an integration strategy for the build process of a Java application. The solution must meet the following requirements:

The builds must access an on-premises dependency management system.

The build outputs must be stored as Server artifacts in Azure DevOps.

The source code must be stored in a Get repository in Azure DevOps.

Solution: Install and configure a self-hosted build agent on an on-premises machine. Configure the build pipeline to use the Default agent pool. Include the Java Tool Installer task in the build pipeline.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**



**Section:**

**Explanation:**

#### QUESTION 8

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 need to recommend an integration strategy for the build process of a Java application. The solution must meet the following requirements:

The builds must access an on-premises dependency management system.

The build outputs must be stored as Server artifacts in Azure DevOps.

The source code must be stored in a Get repository in Azure DevOps.

Solution: Configure the build pipeline to use a Hosted VS 2019 agent pool. Include the Java Tool Installer task in the build pipeline.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use Octopus Tentacle.

Reference:

<https://explore.emtecinc.com/blog/octopus-for-automated-deployment-in-devops-models>

#### QUESTION 9

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 need to recommend an integration strategy for the build process of a Java application. The solution must meet the following requirements:

The builds must access an on-premises dependency management system.

The build outputs must be stored as Server artifacts in Azure DevOps.

The source code must be stored in a Get repository in Azure DevOps.

Solution: Configure the build pipeline to use a Hosted Ubuntu agent pool. Include the Java Tool Installer task in the build pipeline.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use Octopus Tentacle.

Reference: <https://explore.emtecinc.com/blog/octopus-for-automated-deployment-in-devops-models>

#### QUESTION 10

Your company uses a Get repository in Azure Repos to manage the source code of a web application. The master branch is protected from direct updates. Developers work on new features in the topic branches.

Because of the high volume of requested features, it is difficult to follow the history of the changes to the master branch.

You need to enforce a pull request merge strategy. The strategy must meet the following requirements:

Consolidate commit histories.

Merge the changes into a single commit.  
Which merge strategy should you use in the branch policy?

- A. squash merge
- B. fast-forward merge
- C. Get fetch
- D. no-fast-forward merge

**Correct Answer: A**

**Section:**

**Explanation:**

Squash merging is a merge option that allows you to condense the Get history of topic branches when you complete a pull request. Instead of each commit on the topic branch being added to the history of the default branch, a squash merge takes all the file changes and adds them to a single new commit on the default branch.

A simple way to think about this is that squash merge gives you just the file changes, and a regular merge gives you the file changes and the commit history. Note: Squash merging keeps your default branch histories clean and easy to follow without demanding any workflow changes on your team. Contributors to the topic branch work how they want in the topic branch, and the default branches keep a linear history through the use of squash merges. The commit history of a master branch updated with squash merges will have one commit for each merged branch. You can step through this history commit by commit to find out exactly when work was done.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/merging-with-squash>

#### QUESTION 11

Your company uses cloud-hosted Jenkins for builds.

You need to ensure that Jenkins can retrieve source code from Azure Repos.

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

NOTE: Each correct selection is worth one point.

- A. Create a webhook in Jenkins.
- B. Add the Team Foundation Server (TFS) plug-in to Jenkins.
- C. Add a personal access token to your Jenkins account.
- D. Create a personal access token (PAT) in your Azure DevOps account.
- E. Create a service hook in Azure DevOps.

**Correct Answer: B, C, D**

**Section:**

**Explanation:**

B: Jenkins requires a plug-in to connect to TFS and check for updates to a project.

Jenkins' built-in Get Plugin or Team Foundation Server Plugin can poll a Team Services repository every few minutes and queue a job when changes are detected.

C: Use Azure DevOps/ Visual Studio Team Services to create a Personal access token.

D: After you have generated credentials using Visual Studio Team Services, you need to use those credentials in Jenkins.

Reference:

<http://www.aisoftwarellc.com/blog/post/how-to-setup-automated-builds-using-jenkins-and-visual-studio-team-foundation-server/2044>

#### QUESTION 12

You are automating the build process for a Java-based application by using Azure DevOps.

You need to add code coverage testing and publish the outcomes to the pipeline.

What should you use?

- A. Bullseye Coverage
- B. JUnit





- C. JaCoCo
- D. MSTest

**Correct Answer: C**

**Section:**

**Explanation:**

Use Publish Code Coverage Results task in a build pipeline to publish code coverage results to Azure Pipelines or TFS, which were produced by a build in Cobertura or JaCoCo format. Incorrect Answers:

A: Bullseye Coverage is used for C++ code, and not for Java.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/tasks/test/publish-code-coverage-results>

### QUESTION 13

Your company has an Azure DevOps project,

The source code for the project is stored in an on-premises repository and uses on an on-premises build server. You plan to use Azure DevOps to control the build process on the build server by using a self-hosted agent. You need to implement the self-hosted agent.

You download and install the agent on the build server.

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

- A. From Azure, create a shared access signature (SAS).
- B. From the build server, create a certificate, and then upload the certificate to Azure Storage.
- C. From the build server, create a certificate, and then upload the certificate to Azure Key Vault.
- D. From DevOps, create a personal access token (PAT).
- E. From the build server, run config.cmd.

**Correct Answer: B, E**

**Section:**

**Explanation:**

B: Make sure you install your self-signed ssl server certificate into the OS certificate store.

E: When you have a self-signed SSL certificate for your on-premises TFS server, make sure to configure the Get we shipped to allow that self-signed SSL certificate. Enable Get to use SChannel during configure with 2.129.0 or higher version agent Pass --gituseschannel during agent configuration ./config.cmd --gituseschannel

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/certificate>

### QUESTION 14

Explanation:

Create an Azure Resource Manager service connection with an existing service principal

AB: Enter the information about your service principal into the Azure subscription dialog textboxes:

Tenant ID

Subscription ID

Subscription name

Service principal ID

Either the service principal client key or, if you have selected Certificate, enter the contents of both the certificate and private key sections of the \*.pem file.

D: To deploy to a specific Azure resource, the task will need additional data about that resource.

If you're using the classic editor, select data you need. For example, the App service name.

If you're using YAML, then go to the resource in the Azure portal, and then copy the data into your code. For example, to deploy a web app, you would copy the name of the App Service into the WebAppName value.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/library/connect-to-azure>

- A. the tenant ID



- B. the subscription ID
- C. the client secret
- D. the app ID
- E.

**Correct Answer: A, B, D**

**Section:**

**Explanation:**

Create an Azure Resource Manager service connection with an existing service principal

AB: Enter the information about your service principal into the Azure subscription dialog textboxes:

Tenant ID

Subscription ID

Subscription name

Service principal ID

Either the service principal client key or, if you have selected Certificate, enter the contents of both the certificate and private key sections of the \*.pem file.

D: To deploy to a specific Azure resource, the task will need additional data about that resource.

If you're using the classic editor, select data you need. For example, the App service name.

If you're using YAML, then go to the resource in the Azure portal, and then copy the data into your code. For example, to deploy a web app, you would copy the name of the App Service into the WebAppName value.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/library/connect-to-azure>

A. the tenant ID

B. the subscription ID

C. the client secret

D. the app ID

E. the object ID

Answer: ABD

Explanation:

Create an Azure Resource Manager service connection with an existing service principal

AB: Enter the information about your service principal into the Azure subscription dialog textboxes:

Tenant ID

Subscription ID

Subscription name

Service principal ID

Either the service principal client key or, if you have selected Certificate, enter the contents of both the certificate and private key sections of the \*.pem file. D: To deploy to a specific Azure resource, the task will need additional data about that resource.

If you're using the classic editor, select data you need. For example, the App service name.

If you're using YAML, then go to the resource in the Azure portal, and then copy the data into your code. For example, to deploy a web app, you would copy the name of the App Service into the WebAppName value.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/library/connect-to-azure>

#### QUESTION 15

You need to execute inline testing of an Azure DevOps pipeline that uses a Docker deployment model. The solution must prevent the results from being published to the pipeline. What should you use for the inline testing?

- A. a single stage Dockerfile
- B. an Azure Kubernetes Service (AKS) pod
- C. a multi-stage Dockerfile
- D. a Docker Compose file

**Correct Answer: D**

**Section:**



**Explanation:**

Use Docker when running integration tests with Azure Pipelines.

Reference: <https://crossprogramming.com/2019/12/27/use-docker-when-running-integration-tests-with-azure-pipelines.html>

**QUESTION 16**

You are designing an Azure DevOps strategy for your company's development team.

You suspect that the team's productivity is low due to accumulate technical debt.

You need to recommend a metric to assess the amount of the team's technical debt.

What should you recommend?

- A. the number of code modules in an application
- B. the number of unit test failures
- C. the percentage of unit test failures
- D. the percentage of overall time spent on rework

**Correct Answer: D**

**Section:**

**Explanation:**

Technical Debt is the estimated cost to fix code elements issues.

Technical Debt Ratio: Ratio between the cost to develop the software and the cost to fix it. The Technical Debt Ratio formula is:

Remediation cost / Development cost

Which can be restated as:

Remediation cost / (Cost to develop 1 line of code \* Number of lines of code)

Reference: <http://www.azure365.co.in/devops/3PDevOps-4>

**QUESTION 17**

You are developing an open source solution that uses a GitHub repository.

You create a new public project in Azure DevOps.

You plan to use Azure Pipelines for continuous build. The solution will use the GitHub Checks API.

Which authentication type should you use?

- A. OpenID
- B. GitHub App
- C. a personal access token (PAT)
- D. SAML

**Correct Answer: B**

**Section:**

**Explanation:**

Write permission for the Checks API is only available to GitHub Apps.

Note: Authenticating as a GitHub App lets you do a couple of things:

You can retrieve high-level management information about your GitHub App.

You can request access tokens for an installation of the app.

Reference:

<https://docs.github.com/en/rest/guides/getting-started-with-the-checks-api>

**QUESTION 18**

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.

Your company has a project in Azure DevOps for a new web application.

You need to ensure that when code is checked in, a build runs automatically.

Solution: From the Continuous deployment trigger settings of the release pipeline, you enable the Pull request trigger setting.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

In Visual Designer you enable continuous integration (CI) by:

1. Select the Triggers tab.
2. Enable Continuous integration.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/get-started-designer>

#### QUESTION 19

You are automating the build process for a Java-based application by using Azure DevOps.

You need to add code coverage testing and publish the outcomes to the pipeline.

What should you use?

- A. Cobertura
- B. Bullseye Coverage
- C. MSTest
- D. Coverlet
- E. NUnit
- F. Coverage.py

**Correct Answer: A**

**Section:**

**Explanation:**

Use Publish Code Coverage Results task in a build pipeline to publish code coverage results to Azure Pipelines or TFS, which were produced by a build in Cobertura or JaCoCo format.

Incorrect Answers:

B: Bullseye Coverage is used for C++ code, and not for Java.

D: If you're building on Linux or macOS, you can use Coverlet or a similar tool to collect code coverage metrics. Code coverage results can be published to the server by using the Publish Code Coverage Results task. To leverage this functionality, the coverage tool must be configured to generate results in Cobertura or JaCoCo coverage format.

F: Coverage.py is used for Python, not for Java.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/tasks/test/publish-code-coverage-results>

#### QUESTION 20

You have an existing build pipeline in Azure Pipelines.

You need to use incremental builds without purging the environment between pipeline executions.

What should you use?

- A. a self-hosted agent



- B. Microsoft-hosted parallel jobs
- C. a File Transform task

**Correct Answer: A**

**Section:**

**Explanation:**

When you run a pipeline on a self-hosted agent, by default, none of the subdirectories are cleaned in between two consecutive runs. As a result, you can do incremental builds and deployments, provided that tasks are implemented to make use of that. You can override this behavior using the workspace setting on the job.

Incorrect Answers:

B: The workspace clean options are applicable only for self-hosted agents. When using Microsoft-hosted agents job are always run on a new agent.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/process/phases>

#### QUESTION 21

You have a private project in Azure DevOps.

You need to ensure that a project manager can create custom work item queries to report on the project's progress. The solution must use the principle of least privilege. To which security group should you add the project manager?

- A. Reader
- B. Project Collection Administrators
- C. Project Administrators
- D. Contributor

**Correct Answer: D**

**Section:**

**Explanation:**

Contributors have permissions to contribute fully to the project code base and work item tracking. The main permissions they don't have or those that manage or administer resources.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/organizations/security/permissions>

#### QUESTION 22

Your company has a project in Azure DevOps for a new application. The application will be deployed to several Azure virtual machines that run Windows Server 2019.

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

Ensure that the virtual machines maintain a consistent configuration.

Minimize administrative effort to configure the virtual machines.

What should you include in the recommendation?

- A. Azure Resource Manager templates and the PowerShell Desired State Configuration (DSC) extension for Windows
- B. Deployment YAML and Azure pipeline deployment groups
- C. Azure Resource Manager templates and the Custom Script Extension for Windows
- D. Deployment YAML and Azure pipeline stage templates

**Correct Answer: C**

**Section:**

**Explanation:**

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.



Incorrect Answers:

B: YAML doesn't work with Azure pipeline deployment groups.

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/custom-script-windows>

### QUESTION 23

You have an Azure DevOps project that uses many package feeds.

You need to simplify the project by using a single feed that stores packages produced by your company and packages consumed from remote feeds. The solution must support public feeds and authenticated feeds. What should you enable in DevOps?

- A. Universal Packages
- B. upstream sources
- C. views in Azure Artifacts
- D. a symbol server

**Correct Answer: C**

**Section:**

### QUESTION 24

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 plan to create a release pipeline that will deploy Azure resources by using Azure Resource Manager templates. The release pipeline will create the following resources:

Two resource groups

Four Azure virtual machines in one resource group

Two Azure SQL databases in other resource group

You need to recommend a solution to deploy the resources.

Solution: Create a main template that has two linked templates, each of which will deploy the resources in its respective group.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section:**

**Explanation:**

To deploy your solution, you can use either a single template or a main template with many related templates. The related template can be either a separate file that is linked to from the main template, or a template that is nested within the main template.

Reference: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-linked-templates>

### QUESTION 25

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. Your company uses Azure DevOps to manage the build and release processes for applications.

You use a Get repository for applications source control.

You need to implement a pull request strategy that reduces the history volume in the master branch. Solution: You implement a pull request strategy that uses fast-forward merges.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section:**

**Explanation:**

No fast-forward merge - This option merges the commit history of the source branch when the pull request closes and creates a merge commit in the target branch.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/branch-policies>

#### QUESTION 26

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. Your company uses Azure DevOps to manage the build and release processes for applications.

You use a Get repository for applications source control.

You need to implement a pull request strategy that reduces the history volume in the master branch. Solution: You implement a pull request strategy that uses squash merges.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use fast-forward merge.

Note:

Squash merge - Complete all pull requests with a squash merge, creating a single commit in the target branch with the changes from the source branch. No fast-forward merge - This option merges the commit history of the source branch when the pull request closes and creates a merge commit in the target branch.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/branch-policies>

#### QUESTION 27

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. Your company uses Azure DevOps to manage the build and release processes for applications.

You use a Get repository for applications source control.

You need to implement a pull request strategy that reduces the history volume in the master branch. Solution: You implement a pull request strategy that uses an explicit merge.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use fast-forward merge.

Note:



No fast-forward merge - This option merges the commit history of the source branch when the pull request closes and creates a merge commit in the target branch.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/branch-policies>

#### QUESTION 28

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.

Your company uses Azure DevOps to manage the build and release processes for applications.

You use a Get repository for applications source control.

You need to implement a pull request strategy that reduces the history volume in the master branch.

Solution: You implement a pull request strategy that uses a three-way merge.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use fast-forward merge.

Note:

No fast-forward merge - This option merges the commit history of the source branch when the pull request closes and creates a merge commit in the target branch.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/branch-policies>



#### QUESTION 29

You need to recommend a Docker container build strategy that meets the following requirements:

Minimizes image sizes

Minimizes the security surface area of the final image

What should you include in the recommendation?

A. multi-stage builds

B. PowerShell Desired State Configuration (DSC)

C. Docker Swarm

D. single-stage builds

**Correct Answer: A**

**Section:**

**Explanation:**

Multi-stage builds are a new feature requiring Docker 17.05 or higher on the daemon and client. Multistage builds are useful to anyone who has struggled to optimize Dockerfiles while keeping them easy to read and maintain. Incorrect Answers:

C: A swarm consists of multiple Docker hosts which run in swarm mode and act as managers (to manage membership and delegation) and workers (which run swarm services).

Reference:

<https://docs.docker.com/develop/develop-images/multistage-build/>

#### QUESTION 30

You plan to create an image that will contain a .NET Core application.

You have a Dockerfile file that contains the following code. (Line numbers are included for reference only.)



```
01 FROM microsoft/dotnet: 3.1-sdk
02 COPY . /
03 RUN dotnet publish -c Release -o out
04 FROM microsoft/dotnet: 3.1-sdk
05 COPY --from=0 /out /
06 WORKDIR /
07 ENTRYPOINT ["dotnet", "app1.dll"]
```

You need to ensure that the image is as small as possible when the image is built.  
Which line should you modify in the file?

- A. 1
- B. 3
- C. 4
- D. 7

**Correct Answer: C**

**Section:**

**Explanation:**

#### QUESTION 31

You manage build pipelines and deployment pipelines by using Azure DevOps.  
Your company has a team of 500 developers. New members are added continually to the team.  
You need to automate the management of users and licenses whenever possible.  
Which task must you perform manually?

- A. modifying group memberships
- B. adding users
- C. assigning entitlements
- D. procuring licenses

**Correct Answer: D**

**Section:**

**Explanation:**

Incorrect Answers:

A: You can seamlessly replace existing solutions with group-based licensing to more easily manage licenses in Azure DevOps. You can use Group rules.

C: Member Entitlement Management APIs allow managing Entitlements that include -

License

Extensions

Project/Team memberships

Reference:

<https://docs.microsoft.com/en-us/azure/devops/organizations/accounts/migrate-to-group-based-resource-management?view=vsts&tabs=new-nav>

<https://docs.microsoft.com/en-us/rest/api/azure/devops/memberentitlementmanagement/?view=azure-devops-rest-5.0>

#### QUESTION 32

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. Your company has a project in Azure DevOps for a new web application.

You need to ensure that when code is checked in, a build runs automatically.

Solution: From the Triggers tab of the build pipeline, you select Batch changes while a build is in progress. Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead, In Visual Designer you enable continuous integration (CI) by:

1. Select the Triggers tab.
2. Enable Continuous integration.

Note: Batch changes

Select this check box if you have many team members uploading changes often and you want to reduce the number of builds you are running. If you select this option, when a build is running, the system waits until the build is completed and then queues another build of all changes that have not yet been built.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/get-started-designer>

### QUESTION 33

You have 50 Node.js-based projects that you scan by using WhiteSource. Each project includes Package.json, Package-lock.json, and Npm-shrinkwrap.json files. You need to minimize the number of libraries reports by WhiteSource to only the libraries that you explicitly reference. What should you do?

- A. Configure the File System Agent plug-in.
- B. Add a devDependencies section to Package-lock.json.
- C. Configure the Artifactory plug-in.
- D. Delete Package-lock.json.



**Correct Answer: B**

**Section:**

**Explanation:**

Separate Your Dependencies

Within your package.json file be sure you split out your npm dependencies between devDependencies and (production) dependencies. The key part is that you must then make use of the --production flag when installing the npm packages. The --production flag will exclude all packages defined in the devDependencies section.

Reference: <https://blogs.msdn.microsoft.com/visualstudioalmrangers/2017/06/08/manage-your-open-source-usage-and-security-as-reported-by-your-cicd-pipeline/>

### QUESTION 34

Your company deploys applications in Docker containers.

You want to detect known exploits in the Docker images used to provision the Docker containers.

You need to integrate image scanning into the application lifecycle. The solution must expose the exploits as early as possible during the application lifecycle. What should you configure?

- A. a task executed in the continuous integration pipeline and a scheduled task that analyzes the image registry
- B. manual tasks performed during the planning phase and the deployment phase
- C. a task executed in the continuous deployment pipeline and a scheduled task against a running production container
- D. a task executed in the continuous integration pipeline and a scheduled task that analyzes the production container

**Correct Answer: A**

**Section:**

**Explanation:**

You can use the Docker task to sign into ACR and then use a subsequent script to pull an image and scan the container image for vulnerabilities. Use the docker task in a build or release pipeline. This task can be used with Docker or Azure Container registry.

Incorrect Answers:

C: We should not wait until deployment. We want to detect the exploits as early as possible.

D: We should wait until the image is in the product container. We want to detect the exploits as early as possible.

Reference: <https://docs.microsoft.com/en-us/azure/devops/articles/security-validation-cicd-pipeline?view=vsts>

### QUESTION 35

Your company has a hybrid cloud between Azure and Azure Stack.

The company uses Azure DevOps for its full CI/CD pipelines. Some applications are built by using Erlang and Hack. You need to ensure that Erlang and Hack are supported as part of the build strategy across the hybrid cloud. The solution must minimize management overhead. What should you use to execute the build pipeline?

- A. a Microsoft-hosted agent
- B. Azure DevOps self-hosted agents on Azure DevTest Labs virtual machines.
- C. Azure DevOps self-hosted agents on Hyper-V virtual machines
- D. Azure DevOps self-hosted agents on virtual machines that run on Azure Stack

**Correct Answer: D**

**Section:**

**Explanation:**

Azure Stack offers virtual machines (VMs) as one type of an on-demand, scalable computing resource. You can choose a VM when you need more control over the computing environment.

Reference: <https://docs.microsoft.com/en-us/azure/azure-stack/user/azure-stack-compute-overview>

### QUESTION 36

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. Your company has a project in Azure DevOps for a new web application. You need to ensure that when code is checked in, a build runs automatically.

Solution: From the Pre-deployment conditions settings of the release pipeline, you select After stage. Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead, In Visual Designer you enable continuous integration (CI) by:

1. Select the Triggers tab.
2. Enable Continuous integration.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/get-started-designer>

### QUESTION 37

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. Your company has a project in Azure DevOps for a new web application. You need to ensure that when code is checked in, a build runs automatically.

Solution: From the Pre-deployment conditions settings of the release pipeline, you select Batch changes while a build is in progress. Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead, In Visual Designer you enable continuous integration (CI) by:

1. Select the Triggers tab.
2. Enable Continuous integration.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/get-started-designer>

### QUESTION 38

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: D**

**Section:**

**Explanation:**

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 39

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. Your company has a project in Azure DevOps for a new web application.

You need to ensure that when code is checked in, a build runs automatically.

Solution: From the Triggers tab of the build pipeline, you select Enable continuous integration.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section:**

**Explanation:**

In Visual Designer you enable continuous integration (CI) by:

1. Select the Triggers tab.
2. Enable Continuous integration.

A continuous integration trigger on a build pipeline indicates that the system should automatically queue a new build whenever a code change is committed.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/get-started-designer>

#### QUESTION 40

You have an Azure DevOps organization named Contoso and an Azure DevOps project named Project1.

You plan to use Microsoft-hosted agents to build container images that will host full Microsoft .NET Framework apps in a YAML pipeline in Project1. What are two possible virtual machine images that you can use for the Microsoft-hosted agent pool? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. vs2017-win2016
- B. ubuntu-16.04
- C. win1803
- D. macOS-10.13
- E. vs.2015-win2012r2

**Correct Answer: B, C**

**Section:**

**Explanation:**

The Microsoft-hosted agent pool provides 7 virtual machine images to choose from:

Ubuntu 16.04 (ubuntu-16.04)

Windows Server 1803 (win1803) - for running Windows containers

Visual Studio 2019 Preview on Windows Server 2019 (windows-2019)

Visual Studio 2017 on Windows Server 2016 (vs2017-win2016)

Visual Studio 2015 on Windows Server 2012R2 (vs2015-win2012r2)

macOS X Mojave 10.14 (macOS-10.14)

macOS X High Sierra 10.13 (macOS-10.13)

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/hosted?view=azure-devops>



#### QUESTION 41

You plan to share packages that you wrote, tested, validated, and deployed by using Azure Artifacts. You need to release multiple builds of each package by using a single feed. The solution must limit the release of packages that are in development. What should you use?

- A. local symbols
- B. views
- C. global symbols
- D. upstream sources

**Correct Answer: D**

**Section:**

**Explanation:**

Upstream sources enable you to manage all of your product's dependencies in a single feed. We recommend publishing all of the packages for a given product to that product's feed, and managing that product's dependencies from remote feeds in the same feed, via upstream sources. This setup has a few benefits:

Simplicity: your NuGet.config, .npmrc, or settings.xml contains exactly one feed (your feed).

Determinism: your feed resolves package requests in order, so rebuilding the same codebase at the same commit or changeset uses the same set of packages  
Provenance: your feed knows the provenance of packages it saved via upstream sources, so you can verify that you're using the original package, not a custom or malicious copy published to your feed  
Peace of mind: packages used via upstream sources are guaranteed to be saved in the feed on first use; if the upstream source is disabled/removed, or the remote feed goes down or deletes a package you depend on, you can continue to develop and build

Reference:

<https://docs.microsoft.com/en-us/azure/devops/artifacts/concepts/upstream-sources?view=vsts>

#### QUESTION 42

You have a project in Azure DevOps named Project1. Project1 contains a build pipeline named Pipe1 that builds an application named App1. You have an agent pool named Pool1 that contains a Windows Server 2019-based self-hosted agent. Pipe1 uses Pool1. You plan to implement another project named Project2. Project2 will have a build pipeline named Pipe2 that builds an application named App2. App1 and App2 have conflicting dependencies.

You need to minimize the possibility that the two build pipelines will conflict with each other. The solution must minimize infrastructure costs. What should you do?

- A. Add another self-hosted agent.
- B. Add a Docker Compose task to the build pipelines.
- C. Change the self-hosted agent to use Red Hat Enterprise Linux (RHEL) 8.
- D. Create two container jobs.

**Correct Answer: D**

**Section:**

**Explanation:**

To get more control over software dependencies and operating system, you can use Container jobs. Note that the decisions whether to run your pipeline inside a container and whether to use a self-hosted agent are independent. You can directly run your pipeline on a self-hosted agent, or inside a container. You can also execute your pipeline in a container on a Microsoft-hosted agent or on a self-hosted agent. Incorrect Answers:

A: For additional control over hardware, you can use a self-hosted build agent.

Reference:

<http://thewindowsupdate.com/2019/09/09/resolving-complex-software-and-hardware-dependencies-in-azure-devops-pipelines/>

#### QUESTION 43

You are developing an application. The application source has multiple branches.

You make several changes to a branch used for experimentation.

You need to update the main branch to capture the changes made to the experimentation branch and override the history of the Get repository. Which Get option should you use?

- A. Rebase
- B. Fetch
- C. Merge
- D. Push

**Correct Answer: C**

**Section:**

**Explanation:**

Create pull requests to review and merge code in a Get project. Pull requests let your team review code and give feedback on changes before merging it into the master branch.

Incorrect Answers:

A: Use rebase to address the problem of updating your branch with the latest changes from the main branch. Rebase takes the changes made in the commits in your current branch and replays them on the history of another branch. The commit history of your current branch will be rewritten so that it starts from the most recent commit in the target branch of the rebase. Rebasing your changes in your feature branch off the latest changes in the main branch lets you test your changes on the most recent version in the main branch while keeping a clean get history.

D: Share changes made in commits and branches using the push command. Push your branches to the remote repository. get adds your commits to an existing branch on the remote or creates a new branch with the same commits as your local branch.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/pull-requests>

#### QUESTION 44

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 use Azure Pipelines to build and test a React.js application.

You have a pipeline that has a single job.

You discover that installing JavaScript packages from npm takes approximately five minutes each time you run the pipeline. You need to recommend a solution to reduce the pipeline execution time.

Solution: You recommend defining a container job that uses a custom container that has the JavaScript packages preinstalled. Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead enable pipeline caching.

Note:

npm-cache is a command line utility that caches dependencies installed via npm, bower, jspm and composer. It is useful for build processes that run [npm|bower|composer|jspm] install every time as part of their build process. Since dependencies don't change often, this often means slower build times. npm-cache helps alleviate this problem by caching previously installed dependencies on the build machine.

Reference: <https://www.npmjs.com/package/npm-cache>

**QUESTION 45**

HOTSPOT

You currently use JIRA, Jenkins, and Octopus as part of your DevOps processes.

You plan to use Azure DevOps to replace these tools.

Which Azure DevOps service should you use to replace each tool? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Hot Area:**

**Answer Area**

JIRA:  ▼

- Boards
- Build pipelines
- Release pipelines
- Repos

Jenkins:  ▼

- Boards
- Build pipelines
- Release pipelines
- Repos

Octopus:  ▼

- Boards
- Build pipelines
- Release pipelines
- Repos



**Answer Area:**

Answer Area

JIRA:  ▼  
Boards  
Build pipelines  
Release pipelines  
Repos

Jenkins:  ▼  
Boards  
Build pipelines  
Release pipelines  
Repos

Octopus:  ▼  
Boards  
Build pipelines  
Release pipelines  
Repos



Section:

Explanation:

Reference: <https://octopus.com/blog/octopus-jira-integration> <https://www.azuredevopslabs.com/labs/vstsextend/jenkins/>

QUESTION 46

DRAG DROP

You have a project in Azure DevOps that uses packages from multiple public feeds. Some of the feeds are unreliable.

You need to consolidate the packages into a single feed.

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:



**Actions**

- Modify the configuration files to reference the Azure Artifacts feed.
- Run an initial package restore.
- Create a Microsoft Visual Studio project that includes all the packages.
- Create an Azure Artifacts feed that uses upstream sources.
- Create a NuGet package.
- Create an npm package.

**Answer Area**



**Correct Answer:**

**Actions**

- Modify the configuration files to reference the Azure Artifacts feed.
- Run an initial package restore.
- 
- 
- 
- Create an npm package.

**Answer Area**

- Create a NuGet package.
- Create an Azure Artifacts feed that uses upstream sources.
- Create a Microsoft Visual Studio project that includes all the packages.



**Section:**

**Explanation:**

Step 1: Create a NuGet package.

NuGet and Maven are public package managers that support multiple feeds.

Step 2: Create an Azure Artifacts feed that uses upstream sources

If you want to use packages from multiple feeds, use upstream sources to bring packages from multiple feeds together into a single feed.

Step 3: Create a Microsoft Visual Studio project that includes all the packages

Consume NuGet packages from upstream sources: Now you can open Visual Studio and install packages from the upstream sources you just configured.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/artifacts/how-to/set-up-upstream-sources>

**QUESTION 47**

## HOTSPOT

You have the Azure DevOps pipeline shown in the following exhibit.

The screenshot shows the Azure DevOps pipeline configuration interface. The pipeline is named "Build pipeline" and is currently in a "Save & queue" state. The pipeline consists of five tasks: "Get sources" (Build pipeline), "Cloud Agent" (Run on agent), "NuGet restore" (NuGet Installer), "Compile Application" (.NET Core), "Copy Files" (Copy Files), and "Publish Artifact" (Publish Artifacts). The "Cloud Agent" task is highlighted in blue.

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.

### Hot Area:

**Answer Area**

The pipeline has  job(s).

0
1
4

The pipeline has  task(s).

0
1
4

Answer Area:



### Answer Area

The pipeline has  job(s).

0
1
4

The pipeline has  task(s).

0
1
4

**Section:**

**Explanation:**

Box 1: 1

The Cloud agent job only.

Box 2: 4

The pipeline has the four tasks: NuGet restore, Compile Application, Copy Files, and Publish Artifact.

Reference:

<https://azuredevopslabs.com/labs/azuredevops/continuousintegration/>

**QUESTION 48**

DRAG DROP

You need to use Azure Automation State Configuration to manage the ongoing consistency of virtual machine configurations.

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.

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:**



**Actions**

Onboard the virtual machines to Azure Automation State Configuration.

Check the compliance status of the node.

Create a management group.

Assign the node configuration.

Compile a configuration into a node configuration.

Upload a configuration to Azure Automation State Configuration.

Assign tags to the virtual machines.

**Answer Area**



**Correct Answer:**

**Actions**

Create a management group.

Assign tags to the virtual machines.

**Answer Area**


Upload a configuration to Azure Automation State Configuration.

Compile a configuration into a node configuration.

Onboard the virtual machines to Azure Automation State Configuration.

Assign the node configuration.

Check the compliance status of the node.



**Section:**

**Explanation:**

**QUESTION 49**

DRAG DROP

You are building an application that has the following assets:

Source code

Logs from automated tests and builds

Large and frequently updated binary assets

A common library used by multiple applications

Where should you store each asset? To answer, drag the appropriate Azure services to the correct assets. Each service may be used once. 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:**

Azure Services	Answer Area
Azure Artifacts	Source code: <input type="text"/>
Azure Pipelines	A common library used by multiple applications: <input type="text"/>
Azure Repos	Logs from automated tests and builds: <input type="text"/>
Azure Storage	Large and frequently updated binary assets: <input type="text"/>
Azure Test Plans	

**Correct Answer:**

Azure Services	Answer Area
<input type="text"/>	Source code: <input type="text" value="Azure Repos"/>
<input type="text"/>	A common library used by multiple applications: <input type="text" value="Azure Artifacts"/>
<input type="text"/>	Logs from automated tests and builds: <input type="text" value="Azure Pipelines"/>
<input type="text"/>	Large and frequently updated binary assets: <input type="text" value="Azure Storage"/>
Azure Test Plans	

**Section:**

**Explanation:**

Box 1: Azure Repos

Box 2: Azure Artifacts

Use Azure Artifacts to create, host, and share packages with your team.

Box 3: Azure Pipelines

In the pipeline view you can see all the stages and associated tests. The view provides a summary of the test results

Box 4: Azure Storage

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/get-started/what-is-repos>

<https://azure.microsoft.com/en-us/services/devops/artifacts/>

<https://docs.microsoft.com/en-us/azure/devops/pipelines/test/review-continuous-test-results-after-build>

#### QUESTION 50

You have a build pipeline in Azure Pipelines that uses different jobs to compile an application for 10 different architectures.

The build pipeline takes approximately one day to complete.

You need to reduce the time it takes to execute the build pipeline

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

NOTE: Each correct selection is worth one point

- A. Move to a blue/green deployment pattern.
- B. Create an agent pool.
- C. Create a deployment group.
- D. Reduce the size of the repository.
- E. Increase the number of parallel jobs.

**Correct Answer: B, E**

**Section:**

#### QUESTION 51

HOTSPOT

You are designing YAML-based Azure pipelines for the apps shown in the following table.

Name	Platform	Release requirements
App1	Azure virtual machine	Replace a fixed set of existing instances of the previous version of App1 with instances of the new version of the app in each iteration.
App2	Azure Kubernetes Service (AKS) cluster	Roll out a limited deployment of the new version of App2 to validate the functionality of the app. Once testing is successful, expand the rollout.

You need to configure the YAML strategy value for each app. The solution must minimize app downtime.

Which value should you configure for each app? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Hot Area:**



**Answer Area**

App1: 

	▼
canary	
rolling	
runonce	

App2: 

	▼
canary	
rolling	
runonce	

Answer Area:

**Answer Area**

App1: 

	▼
canary	
rolling	
runonce	

App2: 

	▼
canary	
rolling	
runonce	



**Section:**

**Explanation:**

App1: rolling

A rolling deployment replaces instances of the previous version of an application with instances of the new version of the application on a fixed set of virtual machines (rolling set) in each iteration.

App2: canary

Canary deployment strategy is an advanced deployment strategy that helps mitigate the risk involved in rolling out new versions of applications. By using this strategy, you can roll out the changes to a small subset of servers first. As you gain more confidence in the new version, you can release it to more servers in your infrastructure and route more traffic to it. Incorrect Answers:

runonce:

runOnce is the simplest deployment strategy wherein all the lifecycle hooks, namely preDeploy deploy, routeTraffic, and postRouteTraffic, are executed once. Then, either on: success or on: failure is executed.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/process/deployment-jobs>

#### QUESTION 52

DRAG DROP

Your company has two virtual machines that run Linux in a third-party public cloud.

You plan to use the company's Azure Automation State Configuration implementation to manage the two virtual machines and detect configuration drift.

You need to onboard the Linux virtual machines.

You install PowerShell Desired State Configuration (DSC) on the virtual machines, and then run register.py.

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

Select and Place:

#### Actions

Create a DSC metaconfiguration

Copy the metaconfiguration to the virtual machines

Add the virtual machines as DSC nodes in Azure Automation

Install Windows Management Framework 5.1 on the virtual machines

From the virtual machines, run `setdsclocalconfigurationmanager.py`

#### Answer Area



vdumps

Correct Answer:



## Actions

## Answer Area

Create a DSC metaconfiguration

Copy the metaconfiguration to the virtual machines

Add the virtual machines as DSC nodes in Azure Automation



### Section:

#### Explanation:

Step 1: Create a DSC metaconfiguration

Load up the DSC Configuration into Azure Automation.

Step 2: Copy the metaconfiguration to the virtual machines.

Linking the Node Configuration to the Linux Host

Step 3: Add the virtual machines as DSC nodes in Azure Automation.

go to DSC Nodes, select your node, and then click Assign node configuration. This step assigns the DSC configuration to the Linux machine.

Next up will be to link the node configuration to the host. Go to the host and press the "Assign node..."-button. Next up you can select your node configuration.



### QUESTION 53

#### SIMULATION

You plan to store signed images in an Azure Container Registry instance named az4009940427acr1.

You need to modify the SKU for az4009940427acr1 to support the planned images. The solution must minimize costs.

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

A. See solution below.

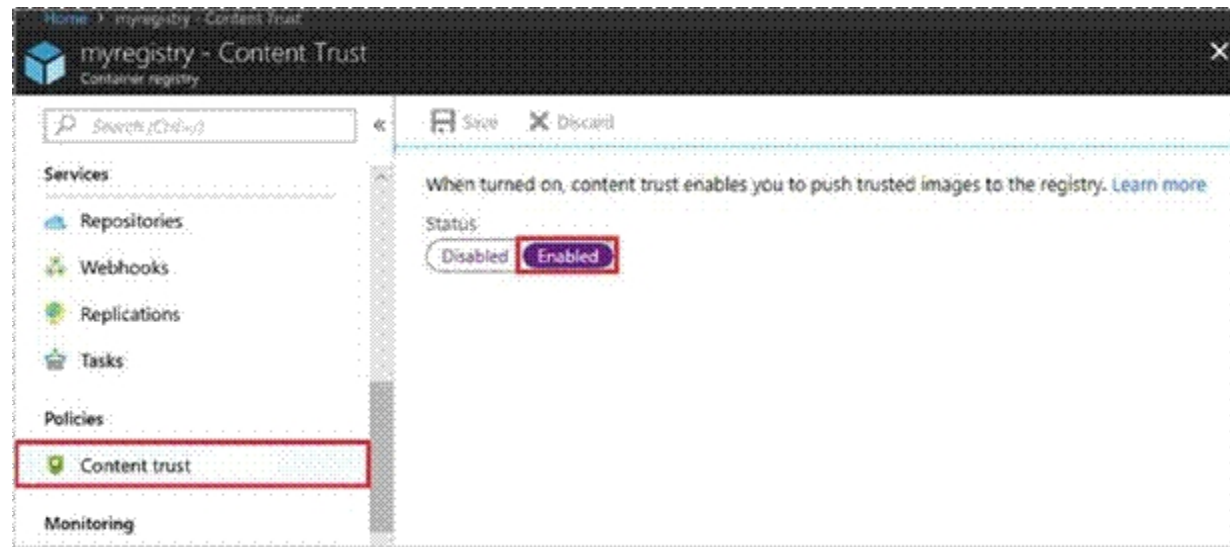
**Correct Answer: A**

#### Section:

#### Explanation:

1. Open Microsoft Azure Portal, and select the Azure Container Registry instance named az4009940427acr1.

2. Under Policies, select Content Trust > Enabled > Save.



Reference:

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-content-trust>

#### QUESTION 54

##### HOTSPOT

Your company uses Team Foundation Server 2013 (TFS 2013).

You plan to migrate to Azure DevOps.

You need to recommend a migration strategy that meets the following requirements:

Preserves the dates of Team Foundation Version Control changesets

Preserves the changes dates of work items revisions

Minimizes migration effort

Migrates all TFS artifacts

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



##### Hot Area:

##### Answer Area

On the TFS server:

- Install the TFS Java SDK.
- Upgrade TFS to the most recent RTW release.
- Upgrade to the most recent version of PowerShell Core.

To perform the migration:

- Copy the assets manually.
- Use public API-based tools.
- Use the TFS Database Import Service.
- Use the TFS Integration Platform.

##### Answer Area:

## Answer Area

On the TFS server:

- Install the TFS Java SDK.
- Upgrade TFS to the most recent RTW release.
- Upgrade to the most recent version of PowerShell Core.

To perform the migration:

- Copy the assets manually.
- Use public API-based tools.
- Use the TFS Database Import Service.
- Use the TFS Integration Platform.

### Section:

### Explanation:

Box 1: Upgrade TFS to the most recent RTM release.

One of the major prerequisites for migrating your Team Foundation Server database is to get your database schema version as close as possible to what is currently deployed in Azure DevOps Services.

Box 2: Use the TFS Database Import Service

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 55

#### DRAG DROP

You are configuring Azure Pipelines for three projects in Azure DevOps as shown in the following table.

Project name	Project Details
Project1	The project team provides preconfigured YAML files that it wants to use to manage future pipeline configuration changes.
Project2	The sensitivity of the project requires that the source code be hosted on the managed Windows server on your company's network.
Project3	The project team requires a centralized version control system to ensure that developers work with the most recent version.

Which version control system should you recommend for each project? To answer, drag the appropriate version control systems to the correct projects. Each version control system 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:

**Version Control Systems**

- Assembla Subversion
- Bitbucket Cloud
- Git in Azure Repos
- GitHub Enterprise

**Answer Area**

- Project1:
- Project2:
- Project3:

**Correct Answer:**

**Version Control Systems**

- 
- Bitbucket Cloud
- 
- 

**Answer Area**

- Project1: Git in Azure Repos
- Project2: GitHub Enterprise
- Project3: Assembla Subversion

**Section:**

**Explanation:**

**QUESTION 56**

**HOTSPOT**

You need to deploy Azure Kubernetes Service (AKS) to host an application. The solution must meet the following requirements:

Containers must only be published internally.

AKS clusters must be able to create and manage containers in Azure.

What should you use for each requirement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Hot Area:**

## Answer Area

Containers must only be published internally:

	▼
Azure Container Instances	
Azure Container Registry	
Dockerfile	

AKS clusters must be able to create and manage containers in Azure:

	▼
An Azure Active Directory (Azure AD) group	
An Azure Automation account	
An Azure service principal	

Answer Area:

## Answer Area

Containers must only be published internally:

	▼
Azure Container Instances	
Azure Container Registry	
Dockerfile	

AKS clusters must be able to create and manage containers in Azure:

	▼
An Azure Active Directory (Azure AD) group	
An Azure Automation account	
An Azure service principal	

**Section:**

**Explanation:**

Box 1: Azure Container Registry

Azure services like Azure Container Registry (ACR) and Azure Container Instances (ACI) can be used and connected from independent container orchestrators like kubernetes (k8s). You can set up a custom ACR and connect it to an existing k8s cluster to ensure images will be pulled from the private container registry instead of the public docker hub.

Box 2: An Azure service principal

When you're using Azure Container Registry (ACR) with Azure Kubernetes Service (AKS), an authentication mechanism needs to be established. You can set up AKS and ACR integration during the initial creation of your AKS cluster. To allow an AKS cluster to interact with ACR, an Azure Active Directory service principal is used.

References:

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

<https://docs.microsoft.com/en-us/azure/aks/cluster-container-registry-integration>

### QUESTION 57

DRAG DROP

You are deploying a new application that uses Azure virtual machines.

You plan to use the Desired State Configuration (DSC) extension on the virtual machines.

You need to ensure that the virtual machines always have the same Windows feature installed.

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:

### Actions

- Configure the DSC extension on the virtual machines.
- Create a YAML configuration file.
- Load the file to Azure Blob storage.
- Configure the Custom Script Extension on the virtual machines.
- Load the file to Azure Files.
- Create a PowerShell configuration file.

### Answer Area



Correct Answer:



## Actions

- Configure the DSC extension on the virtual machines.
- Create a YAML configuration file.
- 
- 
- Load the file to Azure Files.
- 

## Answer Area

- Create a PowerShell configuration file.
- Load the file to Azure Blob storage.
- Configure the Custom Script Extension on the virtual machines.



### Section:

### Explanation:

Step 1: Create a PowerShell configuration file

You create a simple PowerShell DSC configuration file.

Step 2: Load the file to Azure Blob storage

Package and publish the module to a publically accessible blob container URL

Step 3: Configure the Custom Script Extension on the virtual machines.

The Custom Script Extension downloads and executes scripts on Azure virtual machines.

Reference:

<https://docs.microsoft.com/en-us/azure/automation/automation-dsc-getting-started>

<https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/custom-script-windows>

### QUESTION 58

#### DRAG DROP

You have an Azure DevOps release pipeline as shown in the following exhibit.

- Create Resource Group (if not created)  
Azure CLI
- Create Storage Account (if not created)  
Azure CLI
- Create OWASP Container  
Azure CLI

You need to complete the pipeline to configure OWASP ZAP for security testing.

Which five Azure CLI tasks should you add in sequence? To answer, move the tasks from the list of tasks to the answer area and arrange them in the correct order.

Select and Place:

### Tasks

### Answer Area

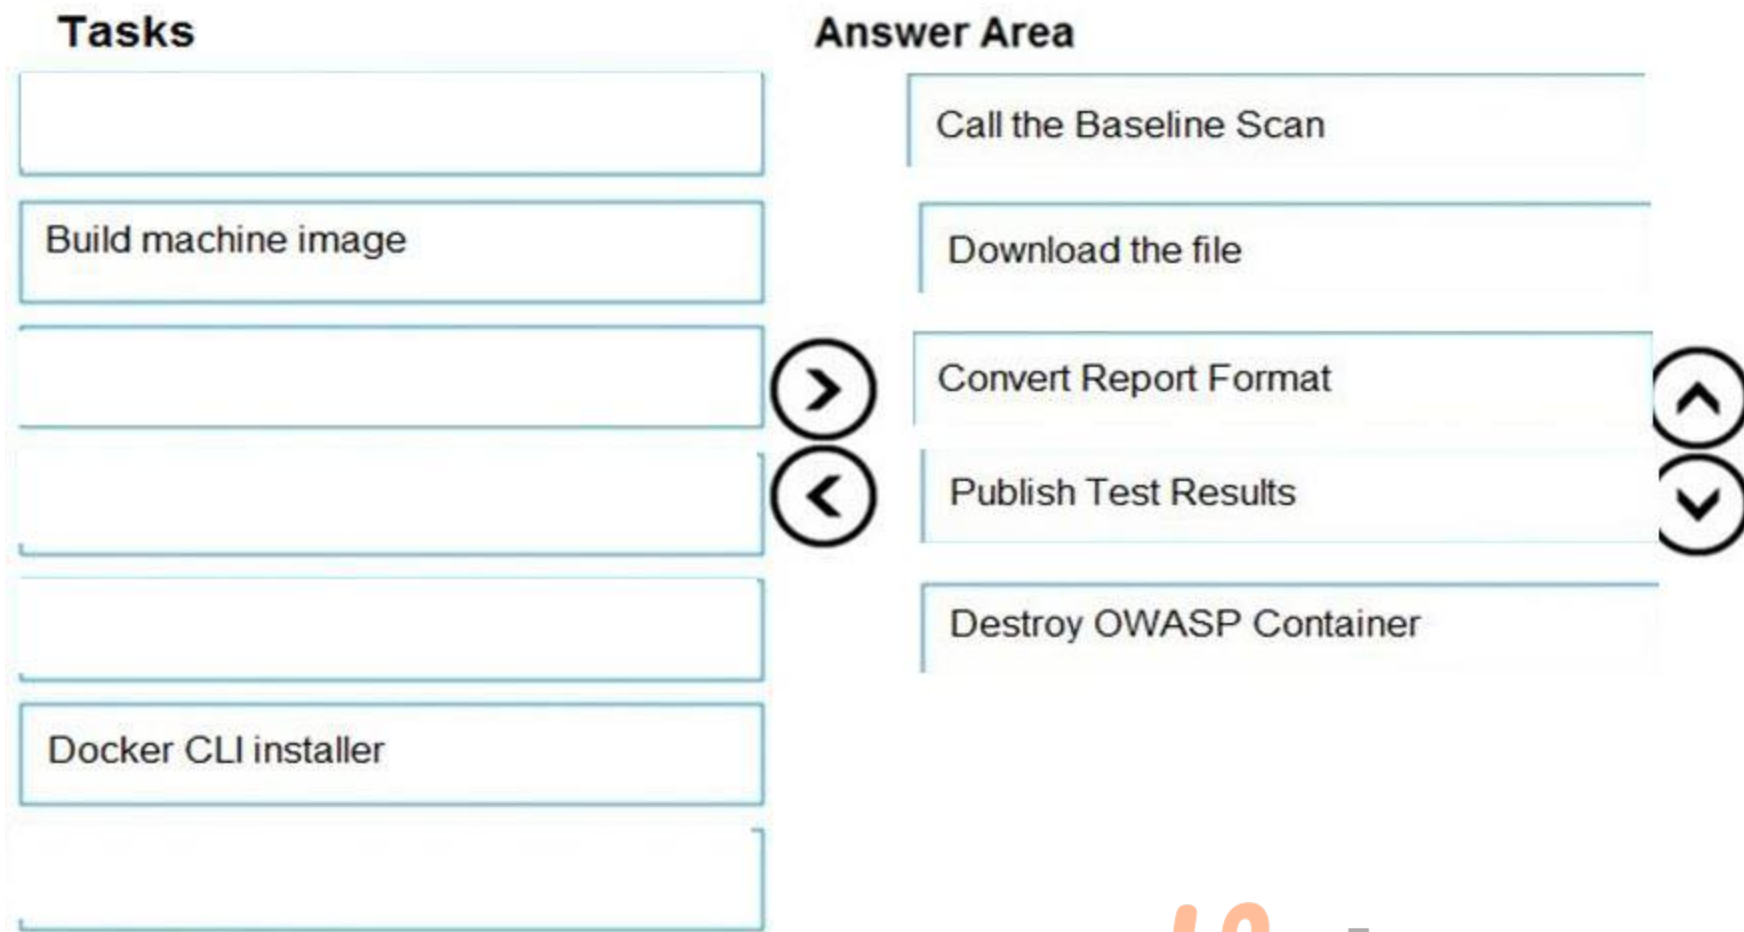
- Convert Report Format
- Build machine image
- Publish Test Results
- Destroy OWASP Container
- Call the Baseline Scan
- Docker CLI installer
- Download the file



Correct Answer:







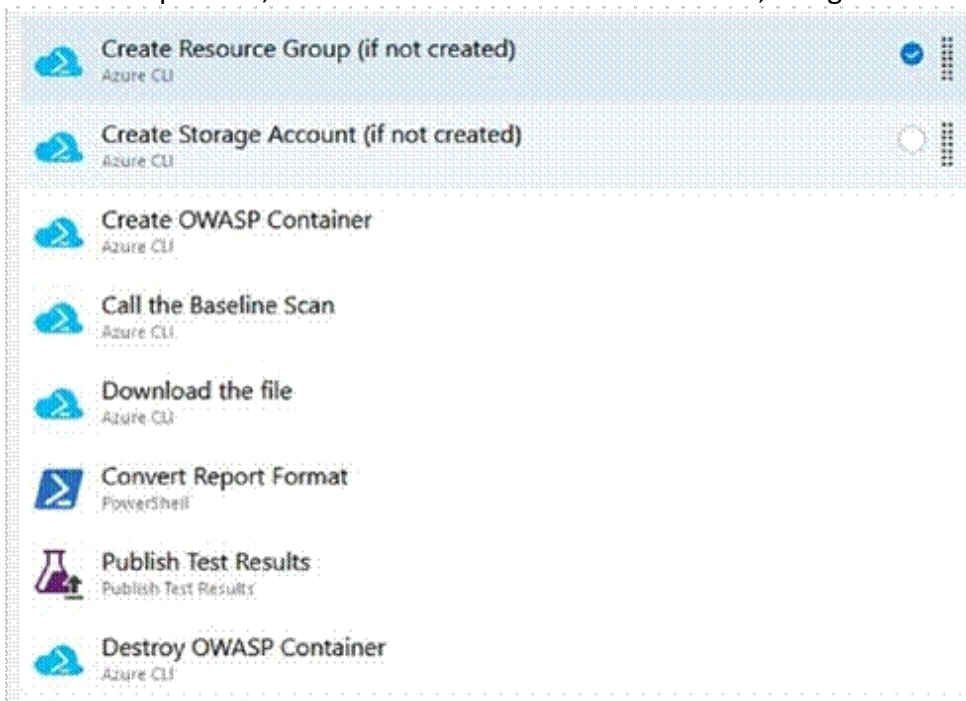
**Section:**

**Explanation:**

Defining the Release Pipeline

Once the application portion of the Release pipeline has been configured, the security scan portion can be defined. In our example, this consists of 8 tasks, primarily using the Azure CLI task to create and use the ACI instance (and supporting structures).

Otherwise specified, all the Azure CLI tasks are Inline tasks, using the default configuration options.



Reference:

<https://devblogs.microsoft.com/premier-developer/azure-devops-pipelines-leveraging-owasp-zap-in-the-release-pipeline/>

### QUESTION 59

#### SIMULATION

You plan to deploy a runbook that will create Azure AD user accounts.

You need to ensure that runbooks can run the Azure PowerShell cmdlets for Azure Active Directory.

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

A. See solution below.

**Correct Answer: A**

**Section:**

**Explanation:**

Azure Automation now ships with the Azure PowerShell module of version 0.8.6, which introduced the ability to non-interactively authenticate to Azure using OrgId (Azure Active Directory user) credential-based authentication. Using the steps below, you can set up Azure Automation to talk to Azure using this authentication type.


Step 1: Find the Azure Active Directory associated with the Azure subscription to manage:

1. Log in to the Azure portal as the service administrator for the Azure subscription you want to manage using Azure Automation. You can find this user by logging in to the Azure portal as any user with access to this Azure subscription, then clicking Settings, then Administrators.



The logo for 'Vdumps' features a stylized orange 'V' followed by the word 'dumps' in a grey, lowercase, sans-serif font.

2. Note the name of the directory associated with the Azure subscription you want to manage. You can find this directory by clicking Settings, then Subscriptions.



Step 2: Create an Azure Active Directory user in the directory associated with the Azure subscription to manage:

You can skip this step if you already have an Azure Active Directory user in this directory. and plan to use this OrgId to manage Azure.

1. In the Azure portal click on Active Directory service.



2. Click the directory name that is associated with this Azure subscription.
3. Click on the Users tab and then click the Add User button.
4. For type of user, select "New user in your organization." Enter a username for the user to create.
5. Fill out the user's profile. For role, pick "User." Don't enable multi-factor authentication. Multi-factor accounts cannot be used with Azure Automation.
6. Click Create.
7. Jot down the full username (including part after @ symbol) and temporary password.

Step 3: Allow this Azure Active Directory user to manage this Azure subscription.

1. Click on Settings (bottom Azure tab under StorSimple)



2. Click Administrators
3. Click the Add button. Type the full user name (including part after @ symbol) of the Azure Active Directory user you want to set up to manage Azure. For subscriptions, choose the Azure subscriptions you want this user to be able to manage. Click the check mark.

Step 4: Configure Azure Automation to use this Azure Active Directory user to manage this Azure subscription Create an Azure Automation credential asset containing the username and password of the Azure Active Directory user that you have just created. You can create a credential asset in Azure Automation by clicking into an Automation Account and then clicking the Assets tab, then the Add Setting button.



ADD CREDENTIAL

Define Credential

USER NAME  
someuser@onmicrosoft.com

PASSWORD  
\*\*\*\*\*

CONFIRM PASSWORD  
\*\*\*\*\*

1 2

Note: Once you have set up the Azure Active Directory credential in Azure and Azure Automation, you can now manage Azure from Azure Automation runbooks using this credential.

Reference:

<https://azure.microsoft.com/sv-se/blog/azure-automation-authenticating-to-azure-using-azure-active-directory/>

#### QUESTION 60

You use a Get repository in Azure Repos to manage the source code of a web application. Developers commit changes directly to the default branch.

You need to implement a change management procedure that meets the following requirements:

The default branch must be protected, and new changes must be built in the feature branches first.

Changes must be reviewed and approved by at least one release manager before each merge.

Changes must be brought into the default branch by using pull requests.

What should you configure in Azure Repos?

- A. branch policies of the default branch
- B. Services in Project Settings
- C. Deployment pools in Project Settings
- D. branch security of the default branch

**Correct Answer: A**

**Section:**

**Explanation:**

Branch policies help teams protect their important branches of development. Policies enforce your team's code quality and change management standards.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/branch-policies>

#### QUESTION 61

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 DevOps project.

Your build process creates several artifacts.

You need to deploy the artifacts to on-premises servers.

Solution: You deploy an Azure self-hosted agent to an on-premises server. You add a Copy and Publish Build Artifacts task to the deployment pipeline. Does this meet the goal?

- A. Yes

B. No

**Correct Answer: A**

**Section:**

**Explanation:**

To build your code or deploy your software using Azure Pipelines, you need at least one agent.

If your on-premises environments do not have connectivity to a Microsoft-hosted agent pool (which is typically the case due to intermediate firewalls), you'll need to manually configure a self-hosted agent on on-premises computer(s). The agents must have connectivity to the target on-premises environments, and access to the Internet to connect to Azure Pipelines or Team Foundation Server.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/agents?view=azure-devops>

#### QUESTION 62

You have a project in Azure DevOps named Project1. Project1 contains a pipeline that builds a container image named Image1 and pushes Image1 to an Azure container registry named ACR1. Image1 uses a base image stored in Docker

Hub.

You need to ensure that Image1 is updated automatically whenever the base image is updated.

What should you do?

- A. Enable the Azure Event Grid resource provider and subscribe to registry events.
- B. Add a Docker Hub service connection to Azure Pipelines.
- C. Create and run an Azure Container Registry task.
- D. Create a service hook in Project1.

**Correct Answer: C**

**Section:**

**Explanation:**

ACR Tasks supports automated container image builds when a container's base image is updated, such as when you patch the OS or application framework in one of your base images.

Reference:

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-tutorial-base-image-update>

#### QUESTION 63

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 DevOps project.

Your build process creates several artifacts.

You need to deploy the artifacts to on-premises servers.

Solution: You deploy an Octopus Deploy server. You deploy a polled Tentacle agent to an on-premises server. You add an Octopus task to the deployment pipeline. Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead you should deploy an Azure self-hosted agent to an on-premises server.

Note: To build your code or deploy your software using Azure Pipelines, you need at least one agent. If your on-premises environments do not have connectivity to a Microsoft-hosted agent pool (which is typically the case due to intermediate firewalls), you'll need to manually configure a self-hosted agent on on-premises computer(s).

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/agents?view=azure-devops>

**QUESTION 64**

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 plan to create a release pipeline that will deploy Azure resources by using Azure Resource Manager templates. The release pipeline will create the following resources:

Two resource groups

Four Azure virtual machines in one resource group

Two Azure SQL databases in other resource group

You need to recommend a solution to deploy the resources.

Solution: Create a main template that will deploy the resources in one resource group and a nested template that will deploy the resources in the other resource group.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Use two linked templates, instead of the nested template.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-linked-templates>

**QUESTION 65**

HOTSPOT

Your company uses a get source-code repository.

You plan to implement Gitflow as a workflow strategy.

You need to identify which branch types are used for production code and preproduction code in the strategy.

Which branch type should you identify for each code type? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Hot Area:**

**Answer Area**

Production code:	<table border="1"><tr><td></td><td>▼</td></tr><tr><td>Master</td><td></td></tr><tr><td>Feature</td><td></td></tr><tr><td>Develop</td><td></td></tr></table>		▼	Master		Feature		Develop	
	▼								
Master									
Feature									
Develop									
Preproduction code:	<table border="1"><tr><td></td><td>▼</td></tr><tr><td>Master</td><td></td></tr><tr><td>Feature</td><td></td></tr><tr><td>Develop</td><td></td></tr></table>		▼	Master		Feature		Develop	
	▼								
Master									
Feature									
Develop									

**Answer Area:**



## Answer Area

Production code:

Preproduction code:

### Section:

#### Explanation:

Box 1: Master

The Master branch contains production code. All development code is merged into master in sometime.

Box 2: Develop

The Develop branch contains pre-production code. When the features are finished then they are merged into develop.

Incorrect Answers:

During the development cycle, a variety of supporting branches are used:

- Feature branches are used to develop new features for the upcoming releases. May branch off from develop and must merge into develop.

Feature branches are used to develop new features for the upcoming releases. May branch off from develop and must merge into develop.

Reference:

<https://medium.com/@patrickporto/4-branching-workflows-for-git-30d0a8ee7bf>

### QUESTION 66

DRAG DROP

You have an Azure subscription that contains a resources group named RG1. RG1 contains the following resources:

Four Azure virtual machines that run Windows Server and have Internet Information Services (IIS) installed. SQL Server on an Azure virtual machine.

An Azure Load Balancer.

You need to deploy an application to the virtual machines in RG1 by using Azure Pipelines.

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:

**Actions** **Answer Area**

- Create an agent pool
- Add the Puppet Agent extension to the virtual machines
- Add and configure a deployment group job for the pipeline
- Add the Azure Pipelines Agent extension to the virtual machines
- Create a deployment group
- Execute the pipeline



**Correct Answer:**

**Actions** **Answer Area**

- 
- Add the Puppet Agent extension to the virtual machines
- 
- 
- 
- Execute the pipeline



- Create an agent pool
- Create a deployment group
- Add the Azure Pipelines Agent extension to the virtual machines
- Add and configure a deployment group job for the pipeline



**Section:**

**Explanation:**

Step 1: Create an agent pool

Azure Pipelines provides a pre-defined agent pool named Azure Pipelines with Microsoft-hosted agents.

Step 2: Create a deployment group

Deployment groups make it easy to define logical groups of target machines for deployment, and install the required agent on each machine.

Step 3: Add the Azure Pipelines Agent extension to the virtual machines

Install the Azure Pipelines Agent Azure VM extension

Step 4: Add and configure a deployment group job for the pipeline

Tasks that you define in a deployment group job run on some or all of the target servers, depending on the arguments you specify for the tasks and the job itself.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/release/deployment-groups/howto-provision-deployment-group-agents>

**QUESTION 67**

Your company develops a client banking application that processes a large volume of data.



Code quality is an ongoing issue for the company. Recently, the code quality has deteriorated because of an increase in time pressure on the development team. You need to implement static code analysis. During which phase should you use static code analysis?

- A. integration testing
- B. staging
- C. production release
- D. build

**Correct Answer: D**

**Section:**

**Explanation:**

#### QUESTION 68

SIMULATION

You have an Azure function hosted in an App Service plan named az400-9940427-func1.

You need to configure az400-9940427-func1 to upgrade the functions automatically whenever new code is committed to the master branch of <https://github.com/Azure-Samples/functions-quickstart>. To complete this task, sign in to the Microsoft Azure portal.

- A. See solution below.

**Correct Answer: A**

**Section:**

**Explanation:**

1. Open Microsoft Azure Portal
2. Log into your Azure account, select App Services in the Azure portal left navigation, and then select configure az400-9940427-func1.
3. On the app page, select Deployment Center in the left menu.
4. On the Build provider page, select Azure Pipelines (Preview), and then select Continue.
5. On the Configure page, in the Code section:  
For GitHub, drop down and select the Organization, Repository, and Branch you want to deploy continuously.
6. Select Continue.
7. On the Test page, choose whether to enable load tests, and then select Continue.
8. Depending on your App Service plan pricing tier, you may see a Deploy to staging page. Choose whether to enable deployment slots, and then select Continue.
9. After you configure the build provider, review the settings on the Summary page, and then select Finish.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-continuous-deployment>

#### QUESTION 69

DRAG DROP

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 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

```
dotnet publish -c Release -o out
```

```
dotnet restore
```

```
mcr.microsoft.com/dotnet/aspnet:5.0
```

```
mcr.microsoft.com/dotnet/sdk:5.0
```

### Answer Area

```
FROM [Value] AS build-env
```

```
COPY . /app/
```

```
WORKDIR /app
```

```
RUN [Value]
```

```
FROM [Value]
```

```
COPY --from=build-env /app/out /app
```

```
WORKDIR /app
```

```
ENTRYPOINT ["dotnet", "MvcMovie.dll"]
```

Correct Answer:

### Values

```
dotnet publish -c Release -o out
```

### Answer Area

```
FROM mcr.microsoft.com/dotnet/sdk:5.0 AS build-env
```

```
COPY . /app/
```

```
WORKDIR /app
```

```
RUN dotnet restore
```

```
FROM mcr.microsoft.com/dotnet/aspnet:5.0
```

```
COPY --from=build-env /app/out /app
```

```
WORKDIR /app
```

```
ENTRYPOINT ["dotnet", "MvcMovie.dll"]
```

Section:

Explanation:

Box 1: mcr.microsoft.com/dotnet/sdk:5.0

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:5.0 comes packaged with the .NET core 5.0 SDK installed, so it's up to the task of building ASP .NET core projects targeting version 5.0

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: mcr.microsoft.com/dotnet/aspnet:5.0

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/en-us/virtualization/windowscontainers/quick-start/building-sample-app>

### QUESTION 70

DRAG DROP

You are configuring the settings of a new get 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.

Select and Place:

### Answer Area



#### Policy Types

A build policy

A check-in policy

A status policy

Committed code must compile successfully:

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

Correct Answer:

## Answer Area

### Policy Types

Committed code must compile successfully:

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

### Section:

#### 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/>



## 02 - Define and implement continuous integration

### 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.

### Overview

Contoso, Ltd. is a manufacturing company that has a main office in Chicago.

### Existing Environment

Contoso plans to improve its IT development and operations processes by implementing Azure DevOps principles. Contoso has an Azure subscription and creates an Azure DevOps organization.

The Azure DevOps organization includes:

### The Docker extension

A deployment pool named Pool7 that contains 10 Azure virtual machines that run Windows Server 2019

The Azure subscription contains an Azure Automation account.

### Requirements

### Planned changes

Contoso plans to create projects in Azure DevOps as shown in the following table.

Project name	Project details
Project 1	Project1 will provide support for incremental builds and third-party SDK components
Project 2	Project2 will use an automatic build policy. A small team of developers named Team2 will work independently on changes to the project. The Team2 members will not have permissions to Project2.
Project 3	Project3 will be integrated with SonarQube
Project 4	Project4 will provide support for a build pipeline that creates a Docker image and pushes the image to the Azure Container Registry. Project4 will use an existing Dockerfile.
Project 5	Project5 will contain a Git repository in Azure Repos and a continuous integration trigger that will initiate a build in response to any change except for changes within /folder1 of the repository.
Project 6	Project6 will provide support for build and deployment pipelines. Deployment will be allowed only if the number of current work items representing active software bugs is 0.
Project 7	Project7 will contain a target deployment group named Group7 that maps to Pool7. Project7 will use Azure Automation State Configuration to maintain the desired state of the computers in Group7.

#### Technical requirements

Contoso identifies the following technical requirements:

Implement build agents for Project1.

Whenever possible, use Azure resources.

Avoid using deprecated technologies.

Implement a code flow strategy for Project2 that will:

- Enable Team2 to submit pull requests for Project2.

- Enable Team2 to work independently on changes to a copy of Project2.

- Ensure that any intermediary changes performed by Team2 on a copy of Project2 will be subject to the same restrictions as the ones defined in the build policy of Project2. Whenever possible, implement automation and minimize administrative effort.

Implement Project3, Project5, Project6, and Project7 based on the planned changes.

Implement Project4 and configure the project to push Docker images to Azure Container Registry.

#### QUESTION 1

DRAG DROP

You need to recommend a procedure to implement the build agent for Project1.

Which three 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:**

**Actions**

Sign in to Azure DevOps by using an account that is assigned the Administrator service connection security role.

Install the Azure Pipelines agent on on-premises virtual machine.

Create a personal access token in the Azure DevOps organization of Contoso.

Install and register the Azure Pipelines agent on an Azure virtual machine.

Sign in to Azure DevOps by using an account that is assigned the agent pool administrator role.

**Answer Area**

**Correct Answer:**

**Actions**

Install the Azure Pipelines agent on on-premises virtual machine.

Sign in to Azure DevOps by using an account that is assigned the agent pool administrator role.

**Answer Area**

Sign in to Azure DevOps by using an account that is assigned the Administrator service connection security role.

Create a personal access token in the Azure DevOps organization of Contoso.

Install and register the Azure Pipelines agent on an Azure virtual machine.



**Section:**

**Explanation:**

**Scenario:**

Project 1	Project1 will provide support for incremental builds and third-party SDK components
-----------	---

Step 1: Sign in to Azure DevOps by using an account that is assigned the Administrator service connection security role.

Note: Under Agent Phase, click Deploy Service Fabric Application. Click Docker Settings and then click Configure Docker settings. In Registry Credentials Source, select Azure Resource Manager Service Connection. Then select your Azure subscription.

Step 2: Create a personal access token..

A personal access token or PAT is required so that a machine can join the pool created with the Agent Pools (read, manage) scope.

Step 3: Install and register the Azure Pipelines agent on an Azure virtual machine.

By running a Azure Pipeline agent in the cluster, we make it possible to test any service, regardless of type.

References:

<https://docs.microsoft.com/en-us/azure/service-fabric/service-fabric-tutorial-deploy-container-app-with-cicd-vsts>

<https://mohitgoyal.co/2019/01/10/run-azure-devops-private-agents-in-kubernetes-clusters/>

## QUESTION 2

DRAG DROP

You need to implement Project6.

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:

### Actions

Open the release pipeline editor.

Disable the continuous integration trigger.

Enable Gates.

Add a manual intervention task.

Open the **Triggers** tab.

Add Query Work Items.

### Answer Area



 **vdumps**

Correct Answer:

### Actions

Disable the continuous integration trigger.

Add a manual intervention task.

Open the **Triggers** tab.

### Answer Area

Open the release pipeline editor.

Enable Gates.

Add Query Work Items.



#### Section:

#### Explanation:

Scenario: Implement Project3, Project5, Project6, and Project7 based on the planned changes



Project 6	Project6 will provide support for build and deployment pipelines. Deployment will be allowed only if the number of current work items representing active software bugs is 0.
-----------	---

Step 1: Open the release pipeline editor.

In the Releases tab of Azure Pipelines, select your release pipeline and choose Edit to open the pipeline editor.

Step 2: Enable Gates.

Choose the pre-deployment conditions icon for the Production stage to open the conditions panel. Enable gates by using the switch control in the Gates section.

Step 3: Add Query Work items.

Choose + Add and select the Query Work Items gate.

Configure the gate by selecting an existing work item query.



Deployment gates ⓘ + Add ▾

**Query Work Items** Enabled ⓘ

Query Work Items ⓘ

Task version: 0.\* ▾

Display name \*  
Query Work Items

Query \* ⓘ  
Active Bugs ▾

Upper threshold \* ⓘ  
0

Advanced ^

Lower threshold \* ⓘ  
0

Output Variables ^

Reference name ⓘ

Variables list  
There are no output variables associated with this task [more information](#)

Evaluation options ▾



Note: A case for release gate is:

Incident and issues management. Ensure the required status for work items, incidents, and issues. For example, ensure deployment occurs only if no priority zero bugs exist, and validation that there are no active incidents takes place after deployment.

References:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/release/deploy-using-approvals?view=azure-devops#configure-gate>

### QUESTION 3

In Azure DevOps, you create Project3.

You need to meet the requirements of the project.

What should you do first?

- A. From Azure DevOps, modify the build definition.
- B. From SonarQube, obtain an authentication token.
- C. From Azure DevOps, create a service endpoint.
- D. From SonarQube, create a project.

**Correct Answer: C**

**Section:**

**Explanation:**

The first thing to do is to declare your SonarQube server as a service endpoint in your VSTS/DevOps project settings.

Reference: <https://docs.sonarqube.org/display/SCAN/Analyzing+with+SonarQube+Extension+for+vsts-TFS>

#### QUESTION 4

You need to implement Project4.  
What should you do first?

- A. Add the FROM instruction in the Dockerfile file.
- B. Add a Copy and Publish Build Artifacts task to the build pipeline.
- C. Add a Docker task to the build pipeline.
- D. Add the MAINTAINER instruction in the Dockerfile file.

**Correct Answer: C**

**Section:**

**Explanation:**

Scenario: Implement Project4 and configure the project to push Docker images to Azure Container Registry.

Project 4	Project4 will provide support for a build pipeline that creates a Docker image and pushes the image to the Azure Container Registry. Project4 will use an existing Dockerfile.
-----------	--

You use Azure Container Registry Tasks commands to quickly build, push, and run a Docker container image natively within Azure, showing how to offload your "inner-loop" development cycle to the cloud. ACR Tasks is a suite of features within Azure Container Registry to help you manage and modify container images across the container lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-quickstart-task-cli>

#### QUESTION 5

HOTSPOT

How should you configure the filters for the Project5 trigger? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

The screenshot shows two configuration sections for a trigger. Each section has a dropdown menu labeled 'Set a' and a list of four filter options: 'branch filter to exclude', 'branch filter to include', 'path filter to exclude', and 'path filter to include'. The first section is associated with the path '/folder1.' and the second with '/.'.

Hot Area:

Answer:

Answer Area

Set a  ▼

- branch filter to exclude
- branch filter to include
- path filter to exclude
- path filter to include

Set a  ▼

- branch filter to exclude
- branch filter to include
- path filter to exclude
- path filter to include

Hot Area:

Answer Area

Set a  ▼

- branch filter to exclude
- branch filter to include
- path filter to exclude
- path filter to include

Set a  ▼

- branch filter to exclude
- branch filter to include
- path filter to exclude
- path filter to include



Answer Area:

Answer Area

Set a  ▼

- branch filter to exclude
- branch filter to include
- path filter to exclude
- path filter to include

Set a  ▼

- branch filter to exclude
- branch filter to include
- path filter to exclude
- path filter to include

Section:

**Explanation:**

<https://docs.microsoft.com/en-us/azure/devops/pipelines/repos/azure-repos-git?view=azure-devops&tabs=classic#ci-triggers>

**01 - Define and implement a continuous delivery and release management strategy**

**QUESTION 1**

**SIMULATION**

You need to ensure that the <https://contoso.com/statushook> webhook is called every time a repository named az40010480345acr1 receives a new version of an image named dotnetapp. To complete this task, sign in to the Microsoft Azure portal.

A. See solution below.

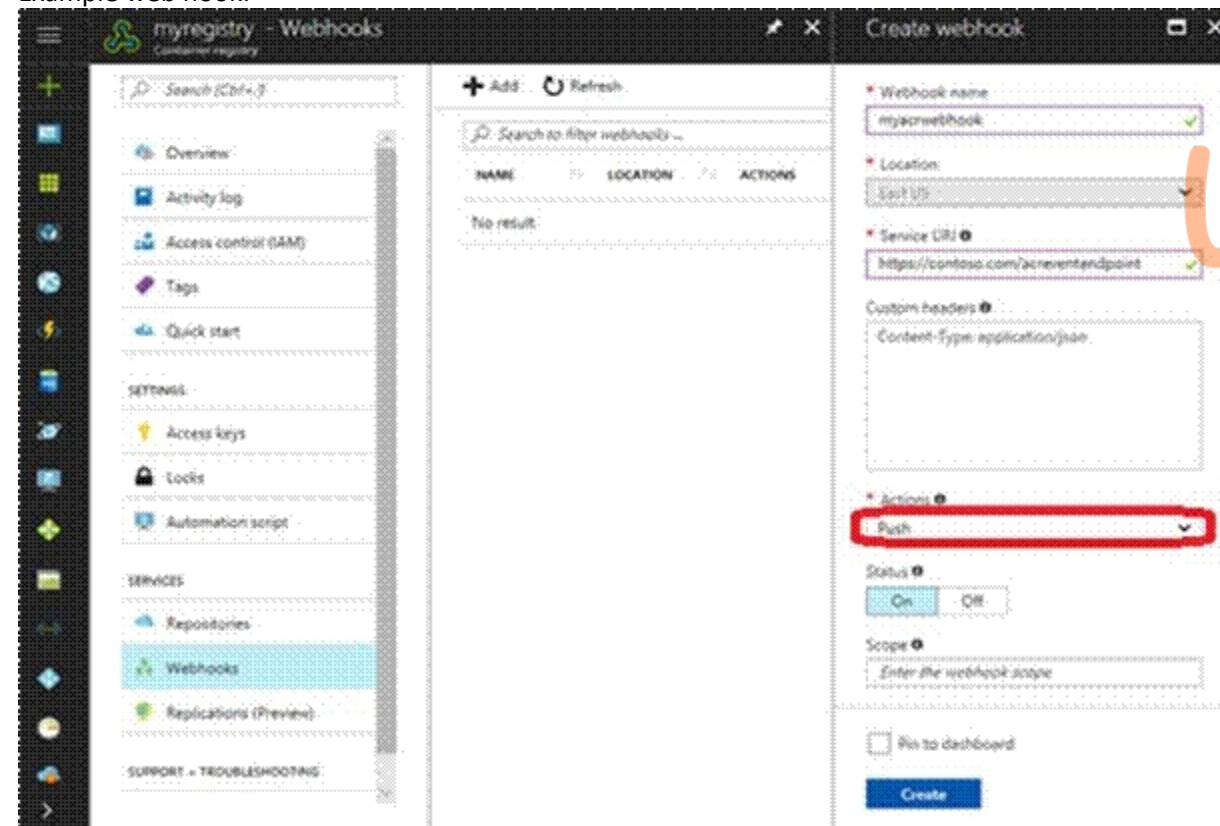
**Correct Answer: A**

**Section:**

**Explanation:**

1. Sign in to the Azure portal.
2. Navigate to the container registry az40010480345acr1.
3. Under Services, select Webhooks.
4. Select the existing webhook <https://contoso.com/statushook>, and double-click on it to get its properties. 5. For Trigger actions select image push

Example web hook:



Reference:

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-webhook>

**QUESTION 2**

Your company uses Azure DevOps to manage the build and release processes for applications.

You use a get repository for applications source control.

You plan to create a new branch from an existing pull request. Later, you plan to merge the new branch and the target branch of the pull request. You need to use a pull request action to create the new branch. The solution must ensure that the branch uses only a portion of the code in the pull request. Which pull request action should you use?

- A. Set as default branch
- B. Approve with suggestions
- C. Cherry-pick
- D. Reactivate
- E. Revert

**Correct Answer: C**

**Section:**

**Explanation:**

Cherry-pick a pull request

To copy changes made in a pull request to another branch in your repo, follow these steps:

1. In a completed pull request, select Cherry-pick, or for an active pull request, select Cherry-pick from the ... menu. Cherry-picking a pull request in this way creates a new branch with the copied changes. Merge into a target branch in a second pull request.
2. In Target branch, enter the branch you want to merge the copied changes.
3. In Topic branch name, enter a new branch to contain the copied changes, then select Cherry-pick.
4. Select Create pull request to merge the topic branch into the target branch to complete the cherry-pick.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/repos/git/pull-requests>

### QUESTION 3

You are designing a build pipeline in Azure Pipelines.

The pipeline requires a self-hosted agent. The build pipeline will run once daily and will take 30 minutes to complete.

You need to recommend a compute type for the agent. The solution must minimize costs.

What should you recommend?

- A. an Azure Kubernetes Service (AKS) cluster
- B. Azure Container Instances
- C. an Azure virtual machine scale set
- D. Azure virtual machines



**Correct Answer: B**

**Section:**

**Explanation:**

<https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/agents?view=azure-devops&tabs=browser#faq>

### QUESTION 4

You have a project in Azure DevOps. You have an Azure Resource Group deployment project in Microsoft Visual Studio that is checked in to the Azure DevOps project. You need to create a release pipeline that will deploy resources by using Azure Resource Manager templates. The solution must minimize administrative effort. Which task type should you include in the solution?

- A. Azure Cloud Service Deployment
- B. Azure RM Web App Deployment
- C. Azure PowerShell
- D. Azure App Service Manage

**Correct Answer: C**

**Section:**

**Explanation:**

There are two different ways to deploy templates to Azure DevOps Services. Both methods provide the same results, so choose the one that best fits your workflow. 1. Add a single step to your build pipeline that runs the PowerShell script that's included in the Azure Resource Group deployment project (Deploy-AzureResourceGroup.ps1). The script copies artifacts and then deploys the template. 2. Add multiple Azure DevOps Services build

steps, each one performing a stage task.

The first option has the advantage of using the same script used by developers in Visual Studio and providing consistency throughout the lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/vs-azure-tools-resource-groups-ci-in-vsts>

#### QUESTION 5

You have an Azure DevOps project that contains a release pipeline and a get repository.

When a new code revision is committed to the repository, a build and release is triggered.

You need to ensure that release information for the pipeline is added automatically to the work items associated to the Get commit. What should you do?

- A. Modify the Integrations options for the pipeline.
- B. Modify the post-deployment conditions for the last stage of the pipeline.
- C. Add an agentless job to the pipeline.
- D. Modify the service hooks for the project.

**Correct Answer: A**

**Section:**

**Explanation:**

#### QUESTION 6

SIMULATION

You plan to add a new web farm that will be published by using an IP address of 10.0.0.5.

You need to allow traffic from the web farm to an Azure Database for MySQL server named az400-11566895-mysql. To complete this task, sign in to the Microsoft Azure portal.

- A. See solution below.

**Correct Answer: A**


**Section:**

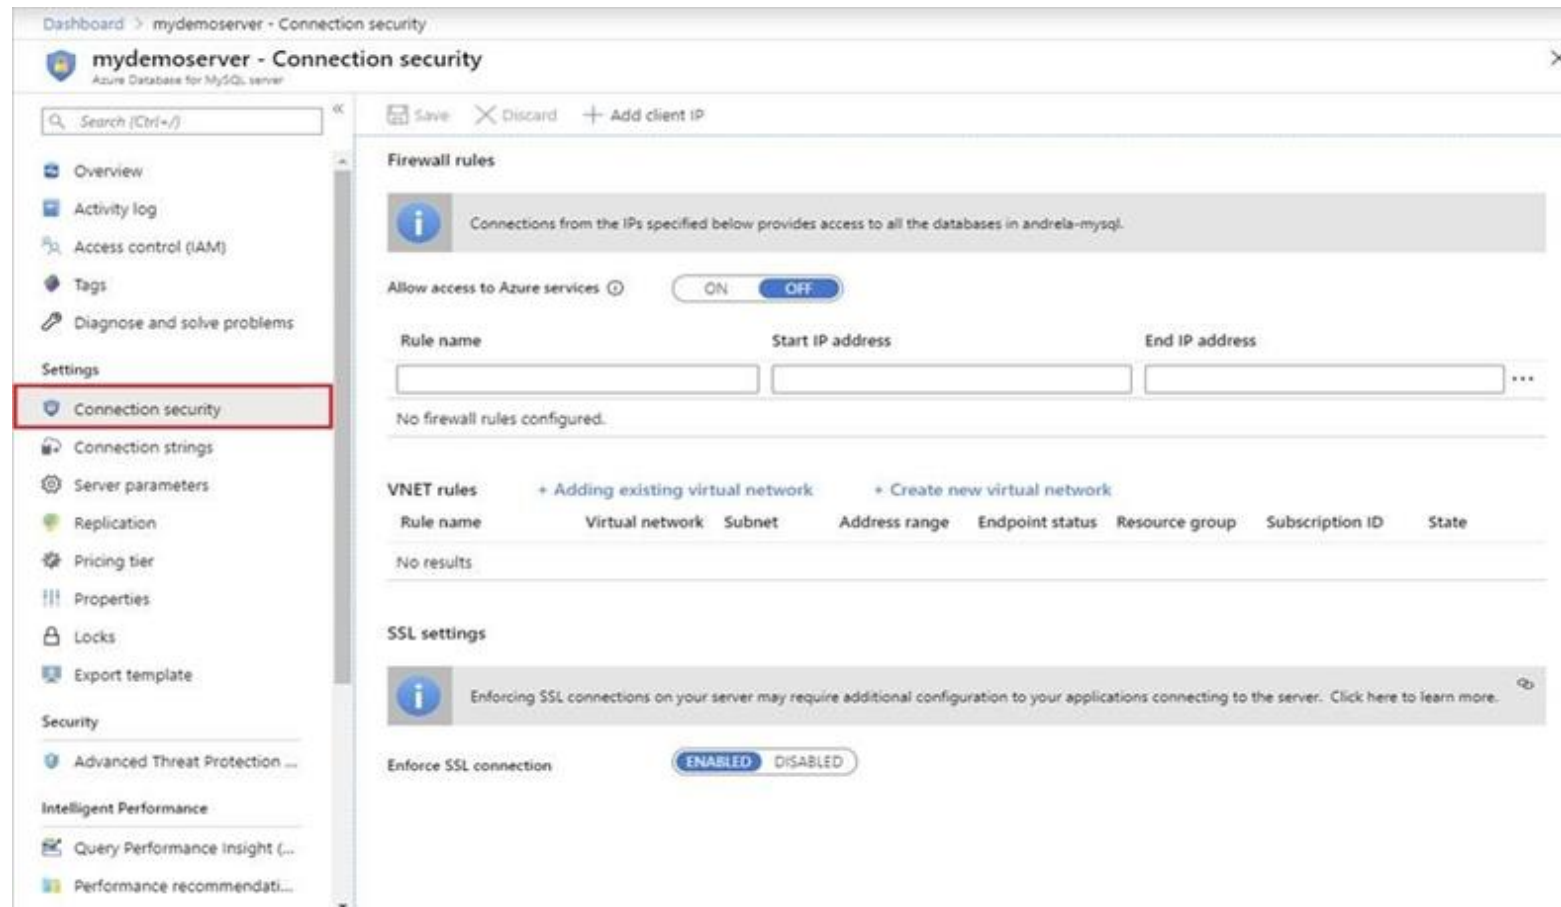
**Explanation:**

Server-level firewall rules can be used to manage access to an Azure Database for MySQL Server from a specified IP address or a range of IP addresses.

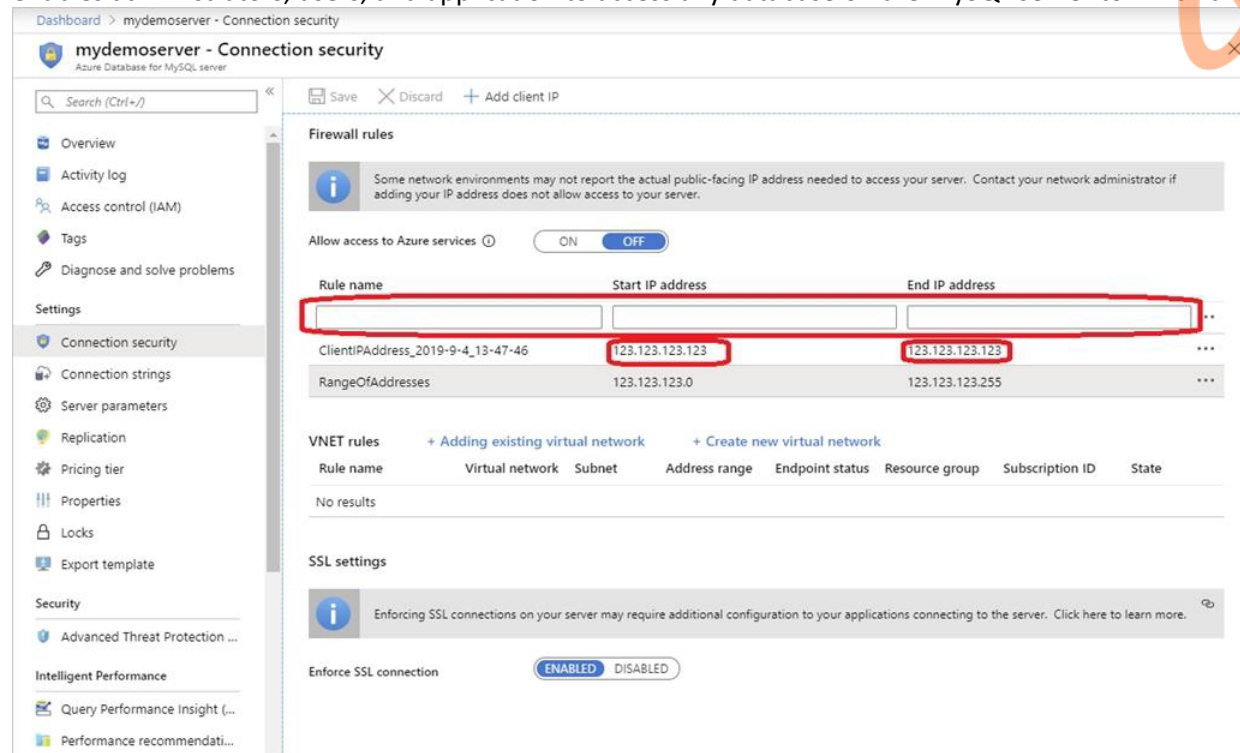
Create a server-level firewall rule in the Azure portal

1. On the MySQL server page, under Settings heading, click Connection Security to open the Connection Security page for the Azure Database for MySQL.





2. In the firewall rules for the Azure Database for MySQL, you can specify a single IP address or a range of addresses. If you want to limit the rule to a single IP address, type the same address in the Start IP and End IP fields. Opening the firewall enables administrators, users, and application to access any database on the MySQL server to which they have valid credentials.



3. Click Save on the toolbar to save this server-level firewall rule. Wait for the confirmation that the update to the firewall rules is successful.

Reference:

<https://docs.microsoft.com/en-us/azure/mysql/howto-manage-firewall-using-portal#create-a-server-level-firewall-rule-in-the-azure-portal>

## QUESTION 7

Your company has a release pipeline in an Azure DevOps project.

You plan to deploy to an Azure Kubernetes Services (AKS) cluster by using the Helm package and deploy task. You need to install a service in the AKS namespace for the planned deployment. Which service should you install?

- A. Azure Container Registry
- B. Chart
- C. Kubectl
- D. Tiller

**Correct Answer: D**

**Section:**

**Explanation:**

Before you can deploy Helm in an RBAC-enabled AKS cluster, you need a service account and role binding for the Tiller service. Incorrect Answers:

C: Kubectl is a command line interface for running commands against Kubernetes clusters.

Reference:

<https://docs.microsoft.com/en-us/azure/aks/kubernetes-helm>

**QUESTION 8**







You need to create a virtual machine template in an Azure DevTest Labs environment named az400-9940427-dtl1. The template must be based on Windows Server 2019 Datacenter. Virtual machines created from the template must include the selenium tool and the Google Chrome browser.



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

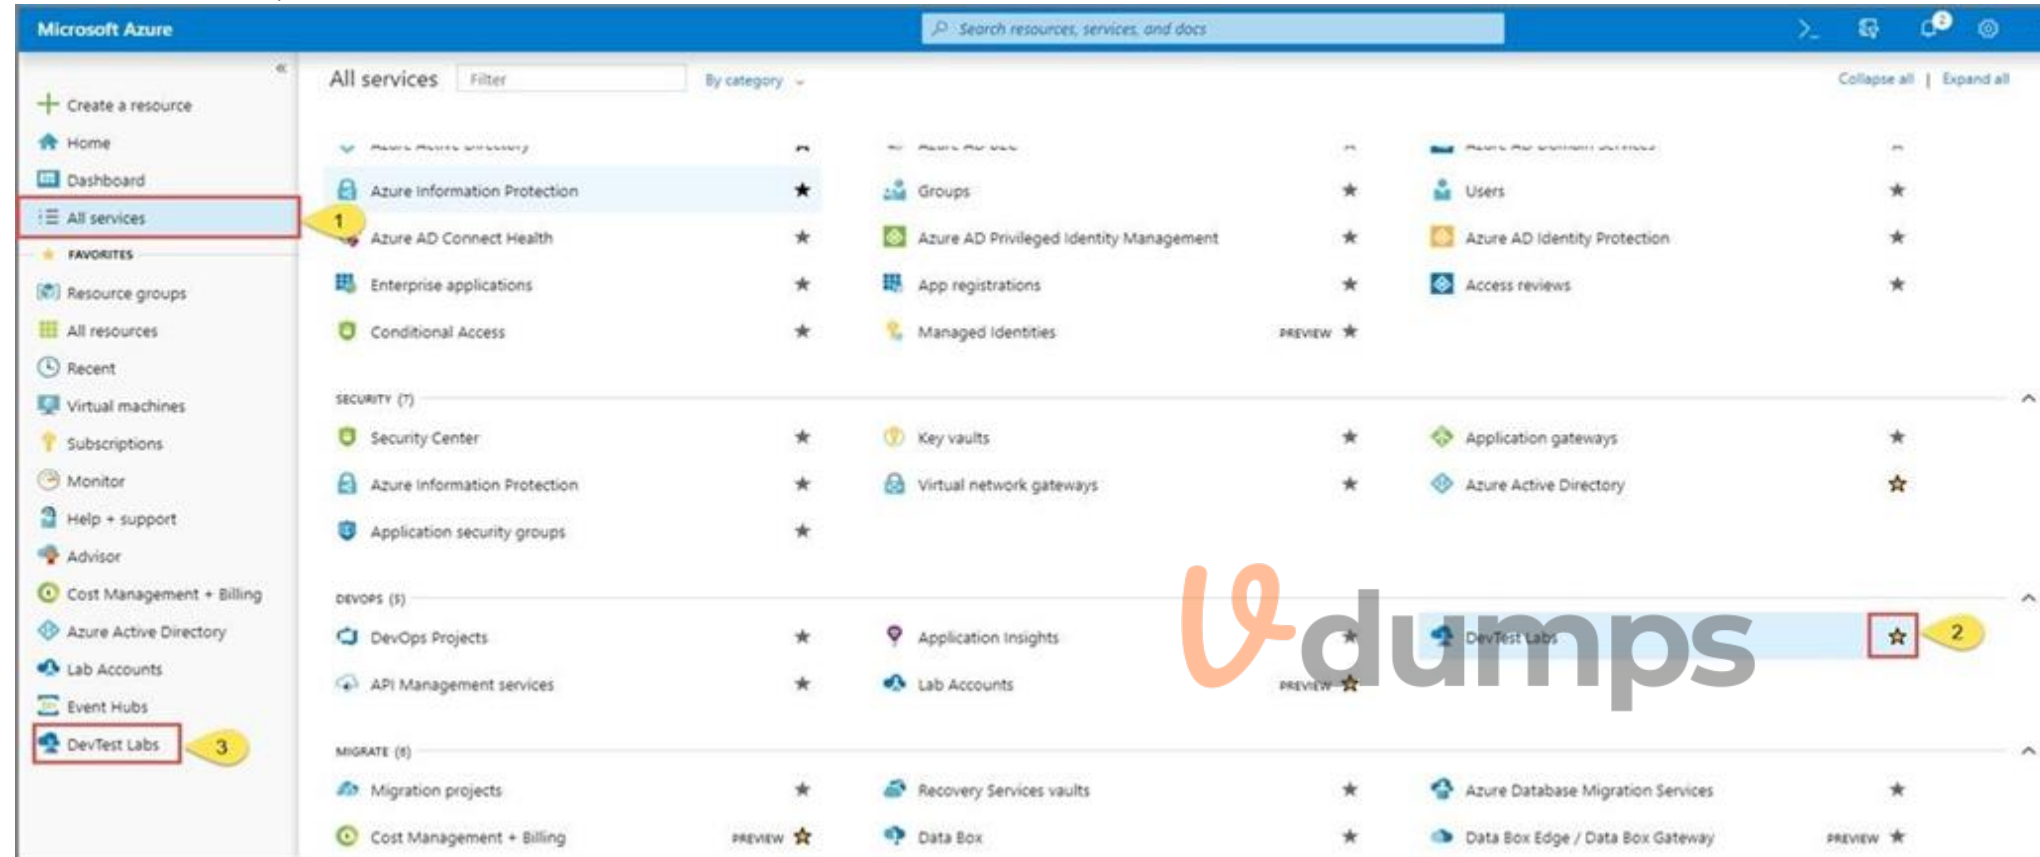
A. See solution below.

**Correct Answer: A**

**Section:**

**Explanation:**

1. Open Microsoft Azure Portal
2. Select All Services, and then select DevTest Labs in the DEVOPS section.



3. From the list of labs, select the az400-9940427-dt11 lab
4. On the home page for your lab, select + Add on the toolbar.
5. Select the Windows Server 2019 Datacenter base image for the VM.
6. Select automation options at the bottom of the page above the Submit button.
7. You see the Azure Resource Manager template for creating the virtual machine.
8. The JSON segment in the resources section has the definition for the image type you selected earlier.

Reference:

<https://docs.microsoft.com/bs-cyrl-ba/azure//lab-services/devtest-lab-vm-powershell>

### QUESTION 9

You plan to use Terraform to deploy an Azure resource group.

You need to install the required frameworks to support the planned deployment.

Which two frameworks should you install? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Vault
- B. Terratest
- C. Node.js
- D. Yeoman

E. Tiller

**Correct Answer: B, D**

**Section:**

**Explanation:**

You can use the combination of Terraform and Yeoman. Terraform is a tool for creating infrastructure on Azure. Yeoman makes it easy to create Terraform modules.

Terratest provides a collection of helper functions and patterns for common infrastructure testing tasks, like making HTTP requests and using SSH to access a specific virtual machine. The following list describes some of the major advantages of using Terratest:

Convenient helpers to check infrastructure - This feature is useful when you want to verify your real infrastructure in the real environment. Organized folder structure - Your test cases are organized clearly and follow the standard Terraform module folder structure. Test cases are written in Go - Many developers who use Terraform are Go developers. If you're a Go developer, you don't have to learn another programming language to use Terratest. Extensible infrastructure - You can extend additional functions on top of Terratest, including Azure-specific features.

Reference:

<https://docs.microsoft.com/en-us/azure/developer/terraform/create-base-template-using-yeoman>

<https://docs.microsoft.com/en-us/azure/developer/terraform/test-modules-using-terratest>

## QUESTION 10

### SIMULATION

You plan to implement a CI/CD strategy for an Azure Web App named az400-11566895-main.

You need to configure a staging environment for az400-11566895-main.

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

A. See solution below.

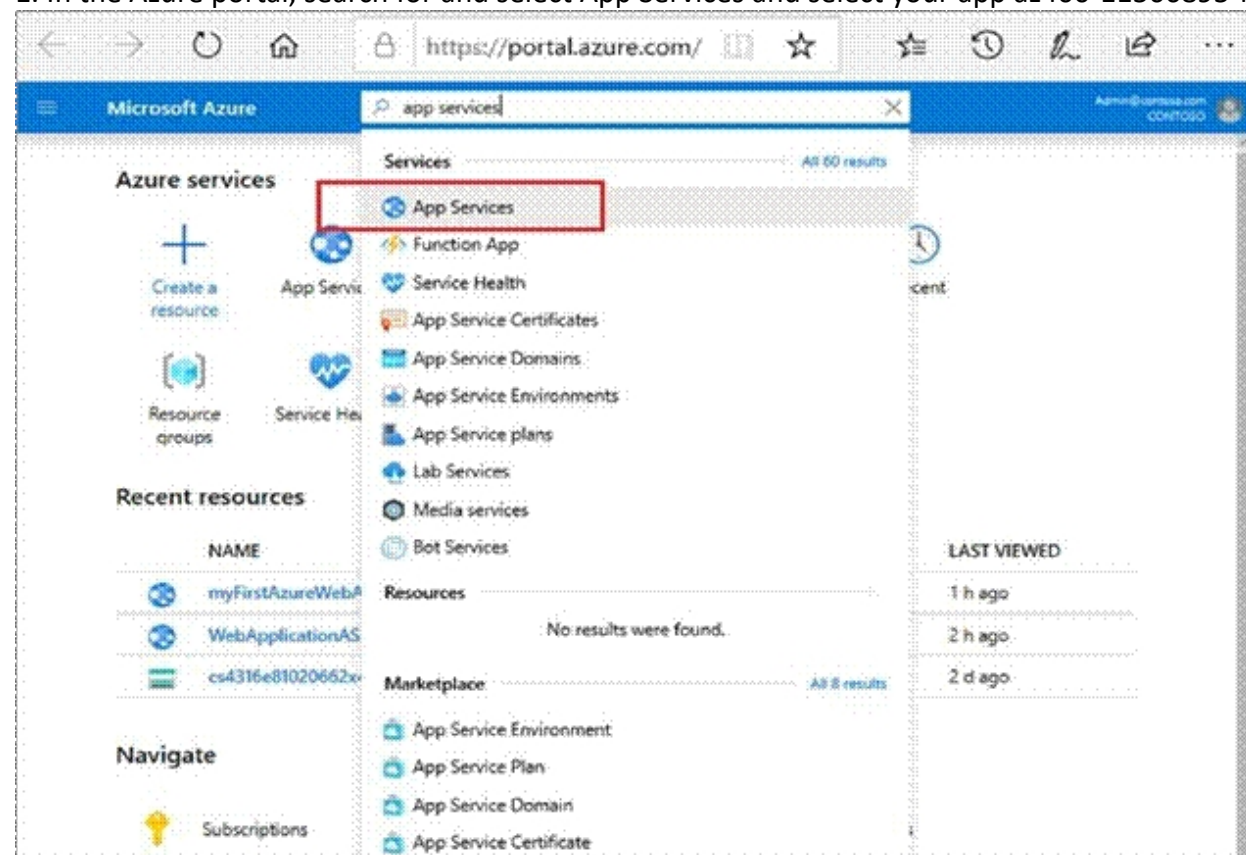
**Correct Answer: A**

**Section:**

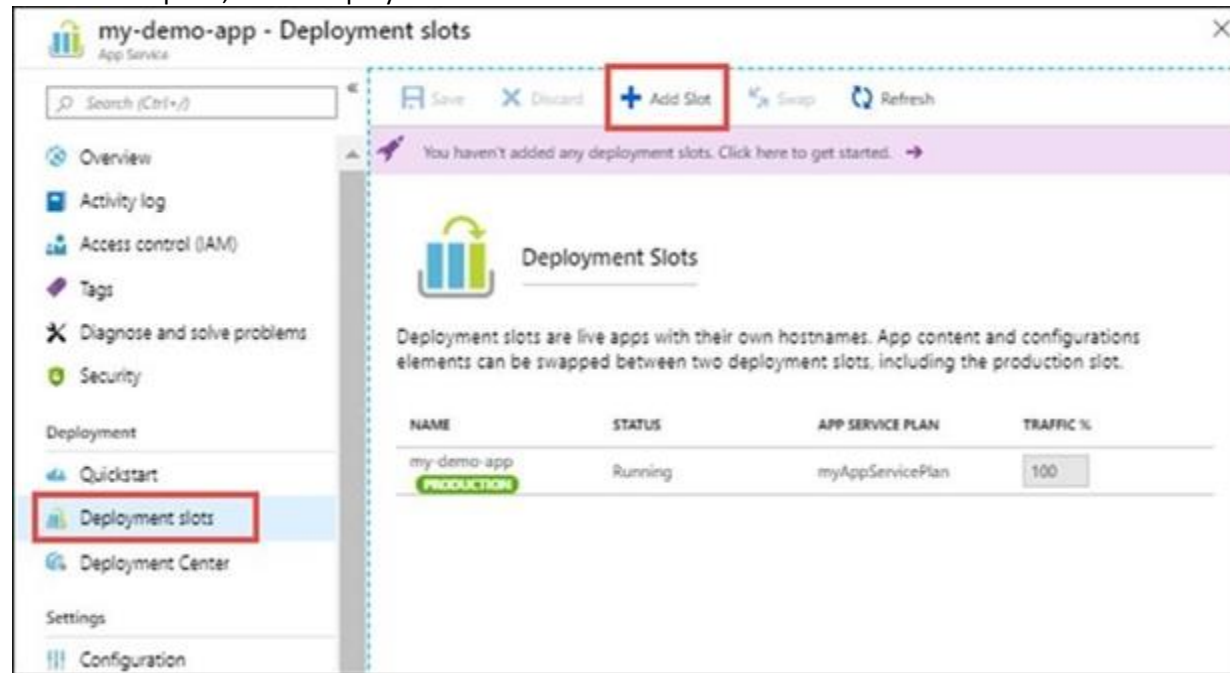
**Explanation:**

Add a slot

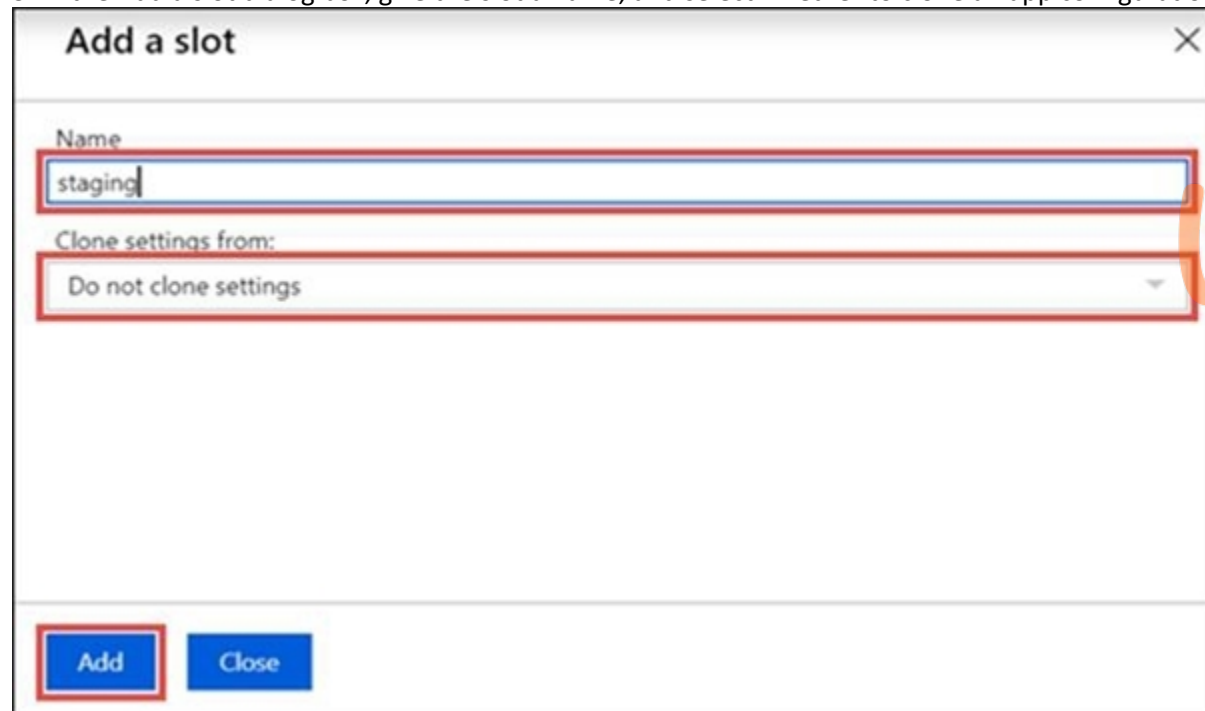
1. In the Azure portal, search for and select App Services and select your app az400-11566895-main.



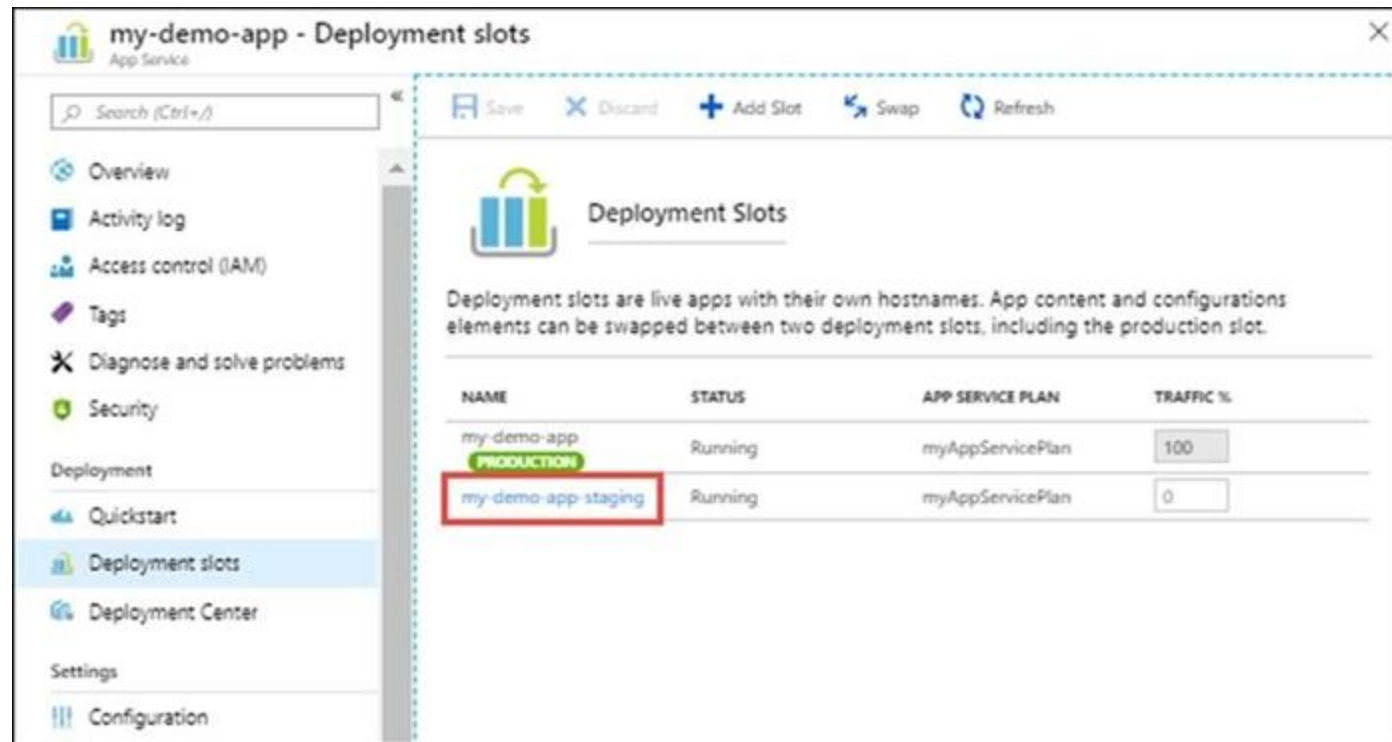
2. In the left pane, select Deployment slots > Add Slot.



3. In the Add a slot dialog box, give the slot a name, and select whether to clone an app configuration from another deployment slot. Select Add to continue.



4. After the slot is added, select Close to close the dialog box. The new slot is now shown on the Deployment slots page.



Reference:  
<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots>

#### QUESTION 11

##### SIMULATION

You have several apps that use an Azure SQL Database named db1.

You need to ensure that queries to db1 are tuned by Azure over time. The solution must only apply to db1. To complete this task, sign in to the Microsoft Azure portal.

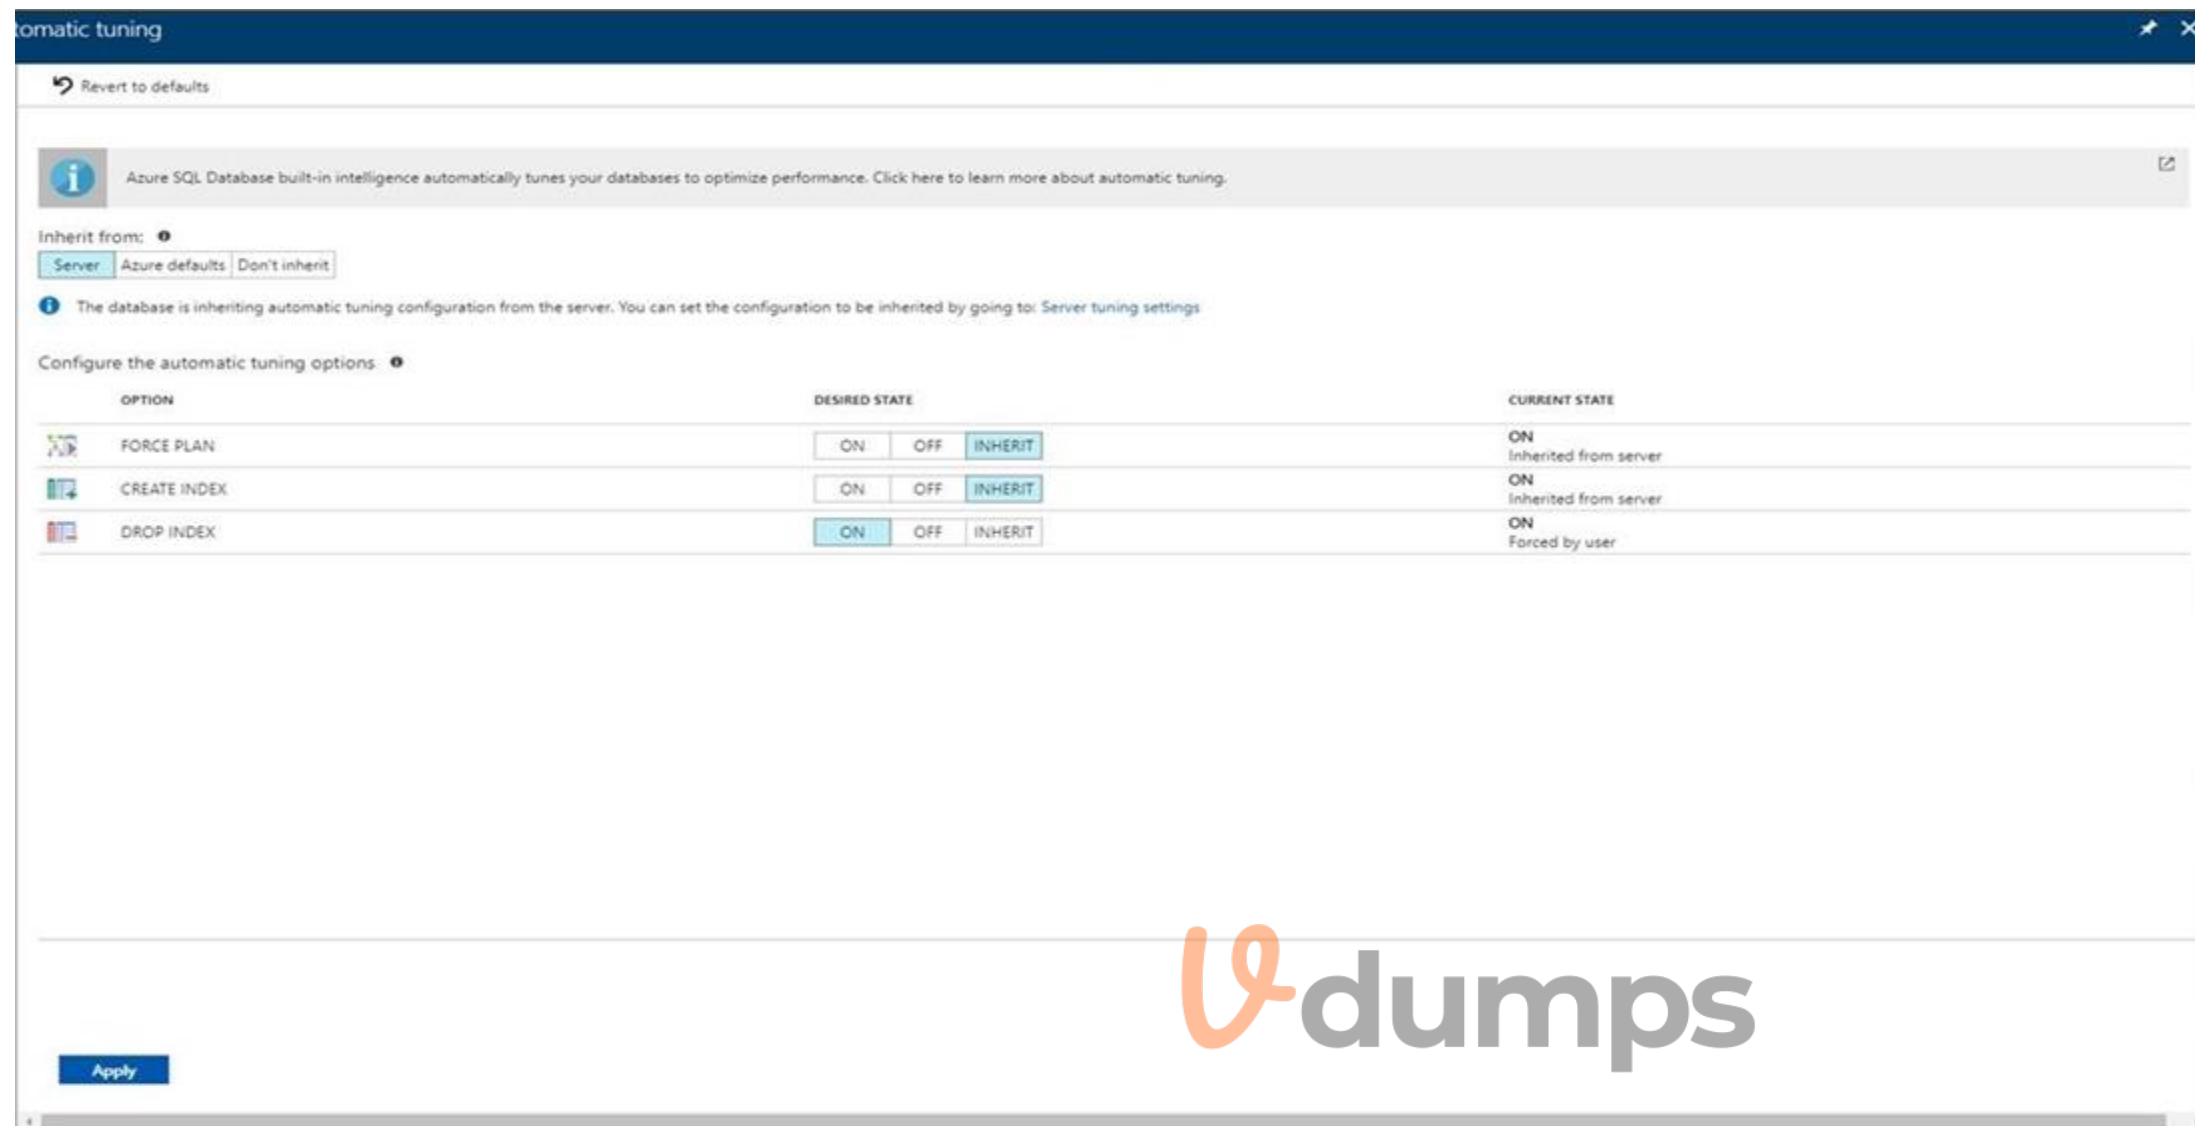
A. See solution below.

**Correct Answer: A**

**Section:**

**Explanation:**

1. To enable automatic tuning on a single database, navigate to the database in the Azure portal and select Automatic tuning.



2. Select the automatic tuning options you want to enable and select Apply.

Note: Individual automatic tuning settings can be separately configured for each database. You can manually configure an individual automatic tuning option, or specify that an option inherits its settings from the server.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-enable>

#### QUESTION 12

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 approval process that contains a condition. The condition requires that releases be approved by a team leader before they are deployed. You have a policy stating that approvals must occur within eight hours.

You discover that deployment fail if the approvals take longer than two hours.

You need to ensure that the deployments only fail if the approvals take longer than eight hours.

Solution: From Post-deployment conditions, you modify the Time between re-evaluation of gates option. Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Use a gate From Pre-deployment conditions instead.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/gates>

### QUESTION 13

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 approval process that contains a condition. The condition requires that releases be approved by a team leader before they are deployed. You have a policy stating that approvals must occur within eight hours.

You discover that deployment fail if the approvals take longer than two hours.

You need to ensure that the deployments only fail if the approvals take longer than eight hours.

Solution: From Pre-deployment conditions, you modify the Time between re-evaluation of gates option. Does this meet the goal?

A. Yes

B. No

**Correct Answer: A**

**Section:**

**Explanation:**

Gates allow automatic collection of health signals from external services, and then promote the release when all the signals are successful at the same time or stop the deployment on timeout. Typically, gates are used in connection with incident management, problem management, change management, monitoring, and external approval systems. Approvals and gates give you additional control over the start and completion of the deployment pipeline. Each stage in a release pipeline can be configured with pre-deployment and post-deployment conditions that can include waiting for users to manually approve or reject deployments, and checking with other automated systems until specific conditions are verified.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/gates>

### QUESTION 14

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 approval process that contains a condition. The condition requires that releases be approved by a team leader before they are deployed. You have a policy stating that approvals must occur within eight hours.

You discover that deployment fail if the approvals take longer than two hours.

You need to ensure that the deployments only fail if the approvals take longer than eight hours.

Solution: From Pre-deployment conditions, you modify the Timeout setting for pre-deployment approvals. Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Use a gate instead of an approval instead.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/gates>

### QUESTION 15

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 approval process that contains a condition. The condition requires that releases be approved by a team leader before they are deployed. You have a policy stating that approvals must occur within eight hours.

You discover that deployment fail if the approvals take longer than two hours.



You need to ensure that the deployments only fail if the approvals take longer than eight hours.

Solution: From Post-deployment conditions, you modify the Timeout setting for post-deployment approvals. Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Use Pre-deployments conditions instead.

Use a gate instead of an approval instead.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/gates>

#### QUESTION 16

You have a free tier of an Azure DevOps organization named Contoso. Contoso contains 10 private projects. Each project has multiple jobs with no dependencies. The build process requires access to resource files located in an onpremises file system.

You frequently run the jobs on five self-hosted agents but experience long build times and frequently queued builds. You need to minimize the number of queued builds and the time it takes to run the builds.

What should you do?

- A. Configure the pipelines to use the Microsoft-hosted agents.
- B. Register additional self-hosted agents.
- C. Purchase self-hosted parallel jobs.
- D. Purchase Microsoft-hosted parallel jobs.

**Correct Answer: B**

**Section:**

**Explanation:**

If you want Azure Pipelines to orchestrate your builds and releases, but use your own machines to run them, use self-hosted parallel jobs. For self-hosted parallel jobs, you'll start by deploying our self-hosted agents on your machines. You can register any number of these self-hosted agents in your organization.

Incorrect:

Not D: Microsoft-hosted CI/CD

If you want to run your jobs on machines that Microsoft manages, use Microsoft-hosted parallel jobs. Your jobs run on our pool of Microsoft-hosted agents. We provide a free tier of service by default in every organization.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/licensing/concurrent-jobs>

#### QUESTION 17

SIMULATION

You need to ensure that an Azure web app named az400-9940427-main supports rolling upgrades. The solution must ensure that only 10 percent of users who connect to az400-9940427-main use update versions of the app. The solution must minimize administrative effort.

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

- A. See solution below.

**Correct Answer: A**

**Section:**

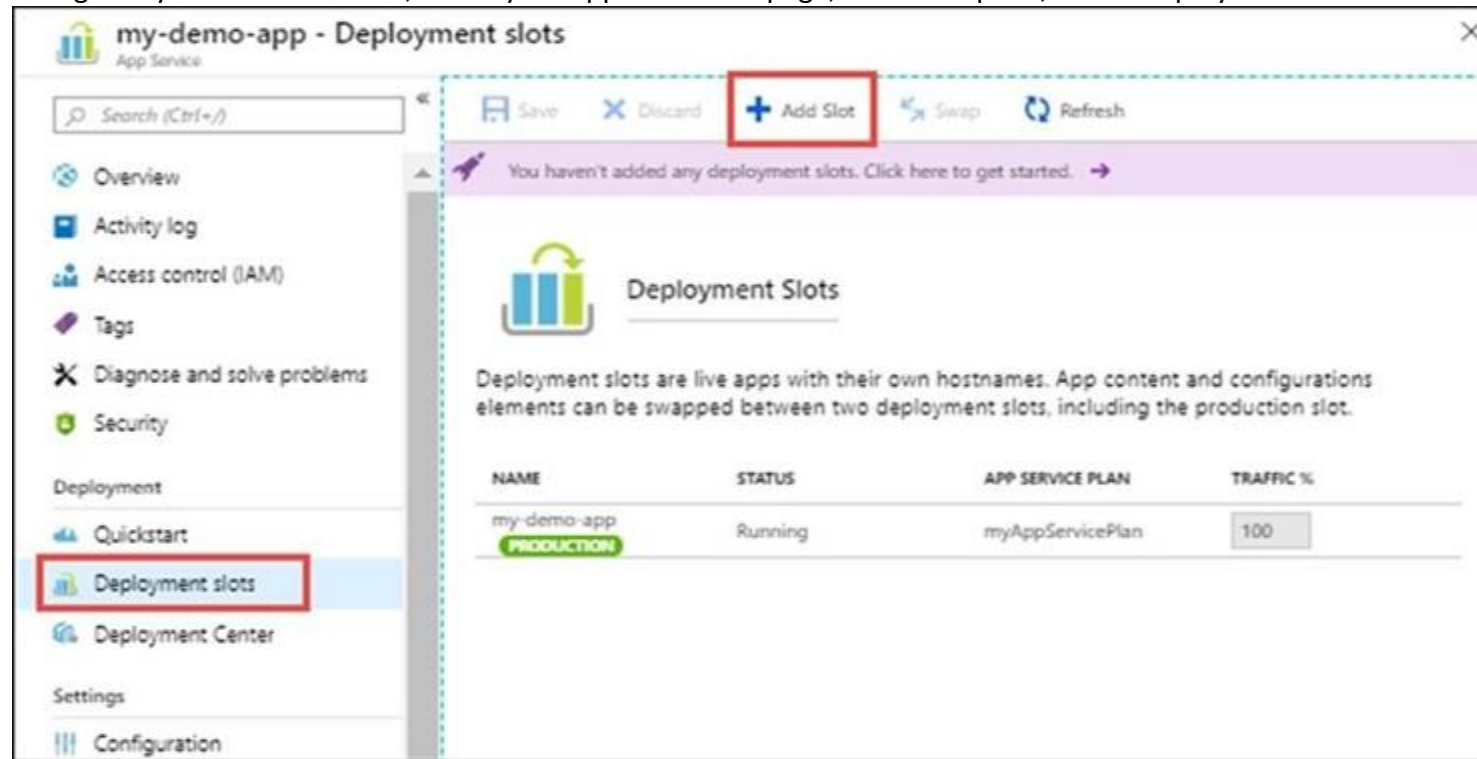
**Explanation:**

Set up staging environments in Azure App Service

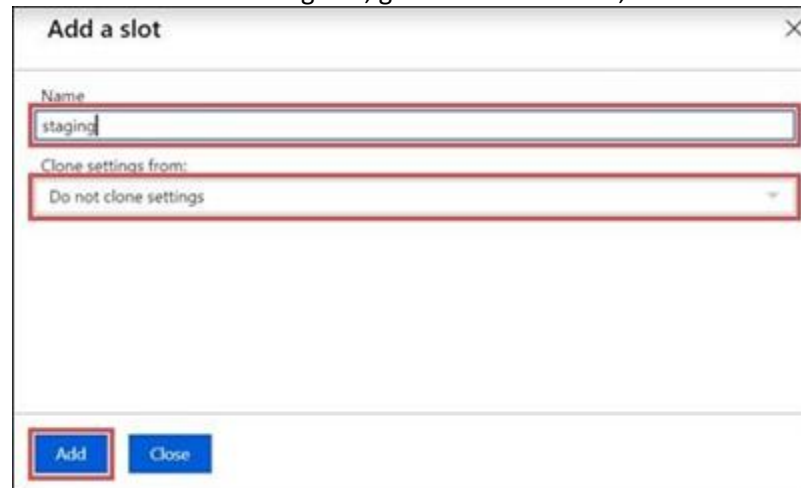
1. Open Microsoft Azure Portal



2. Log into your Azure account, select your app's resource page, in the left pane, select Deployment slots > Add Slot.



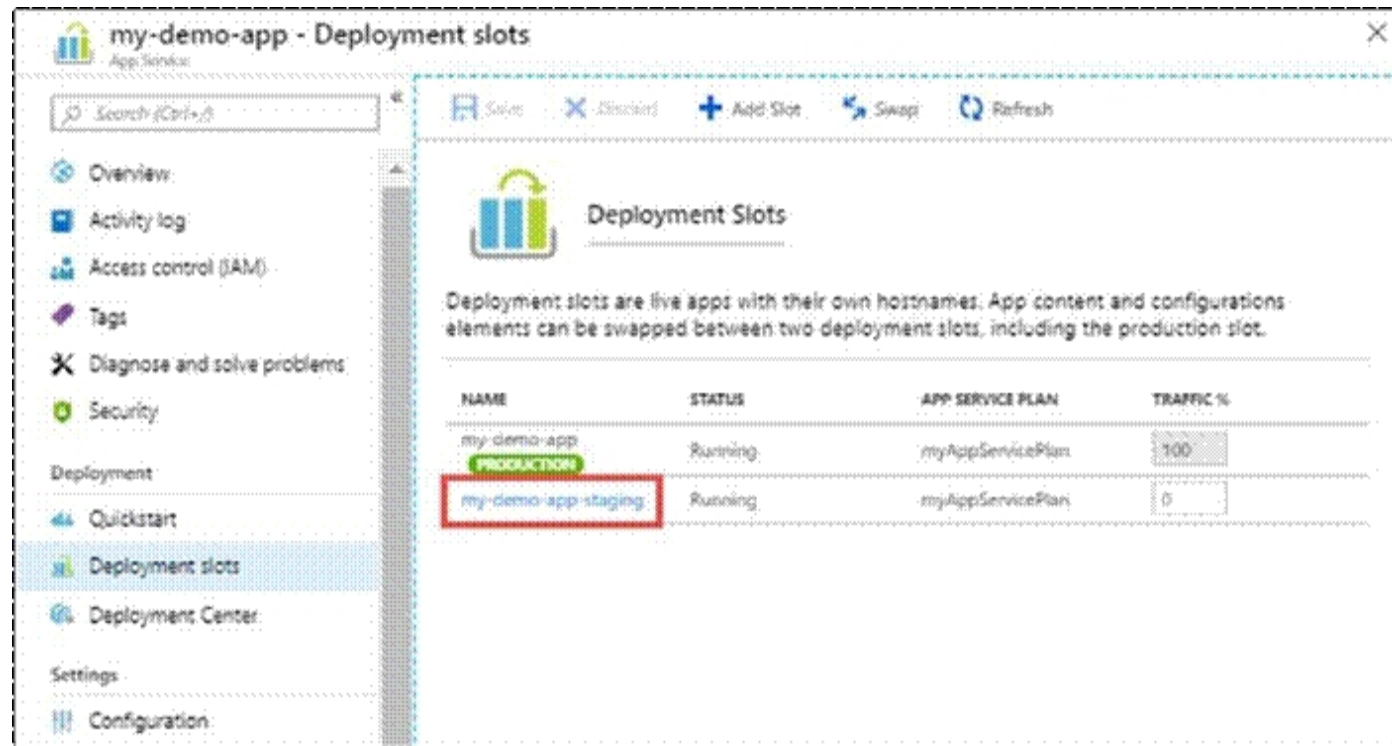
3. In the Add a slot dialog box, give the slot a name, and select whether to clone an app configuration from another deployment slot. Select Add to continue.



4. After the slot is added, select Close to close the dialog box. The new slot is now shown on the Deployment slots page. By default, Traffic % is set to 0 for the new slot, with all customer traffic routed to the production slot

5. Select the new deployment slot to open that slot's resource page.





6. Change TRAFFIC % to 10

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots>

#### QUESTION 18

You have an Azure DevOps project named Project1 and an Azure subscription named Sub1. Sub1 contains an Azure SQL database named DB1. You need to create a release pipeline that uses the Azure SQL Database Deployment task to update DB1. Which artifact should you deploy?

- A. a BACPAC
- B. a DACPAC
- C. an LDF file
- D. an MDF file

**Correct Answer: B**

**Section:**

**Explanation:**

Use Azure SQL Database Deployment task in a build or release pipeline to deploy to Azure SQL DB using a DACPAC or run scripts using SQLCMD.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/tasks/deploy/sql-azure-dacpac-deployment>

#### QUESTION 19

Your company has a project in Azure DevOps for a new web application.

The company uses ServiceNow for change management.

You need to ensure that a change request is processed before any components can be deployed to the production environment. What are two ways to integrate ServiceNow into the Azure DevOps release pipeline? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.


- A. Define a deployment control that invokes the ServiceNow REST API.
- B. Define a pre-deployment gate before the deployment to the Prod stage.
- C. Define a deployment control that invokes the ServiceNow SOAP API.
- D. Define a post-deployment gate after the deployment to the QA stage.

**Correct Answer: B, D**

**Section:**

**Explanation:**

An example of a release pipeline that can be modeled through a release pipeline is shown below:



In this example, a release of a website is created by collecting specific versions of two builds (artifacts), each from a different build pipeline. The release is first deployed to a Dev stage and then forked to two QA stages in parallel. If the deployment succeeds in both the QA stages, the release is deployed to Prod ring 1 and then to Prod ring 2. Each production ring represents multiple instances of the same website deployed at various locations around the globe.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/release>

#### QUESTION 20

Your company develops an application named App1 that is deployed in production.

As part of an application update, a new service is being added to App1. The new service requires access to an application named App2 that is currently in development. You need to ensure that you can deploy the update to App1 before App2 becomes available. You must be able to enable the service in App1 once App2 is deployed. What should you do?

- A. Implement a feature flag.
- B. Create a fork in the build.
- C. Create a branch in the build.
- D. Implement a branch policy.



**Correct Answer: A**

**Section:**

**Explanation:**

Feature flags support a customer-first DevOps mindset, to enable (expose) and disable (hide) features in a solution, even before they are complete and ready for release.

Incorrect Answers:

C: Branch policies are an important part of the get workflow and enable you to:

Isolate work in progress from the completed work in your master branch

Guarantee changes build before they get to master

Reference:

<https://docs.microsoft.com/en-us/azure/devops/migrate/phase-features-with-feature-flags>

#### QUESTION 21

You have a private distribution group that contains provisioned and unprovisioned devices.

You need to distribute a new iOS application to the distribution group by using Microsoft Visual Studio App Center. What should you do?

- A. Request the Apple ID associated with the user of each device.
- B. Register the devices on the Apple Developer portal.
- C. Create an active subscription in App Center Test.
- D. Add the device owner to the organization in App Center.

**Correct Answer: B**

**Section:**

**Explanation:**

When releasing an iOS app signed with an ad-hoc or development provisioning profile, you must obtain tester's device IDs (UDIDs), and add them to the provisioning profile before compiling a release. When you enable the distribution group's Automatically manage devices setting, App Center automates the before mentioned operations and removes the constraint for you to perform any manual tasks. As part of automating the workflow, you must provide the user name and password for your Apple ID and your production certificate in a .p12 format.

App Center starts the automated tasks when you distribute a new release or one of your testers registers a new device. First, all devices from the target distribution group will be registered, using your Apple ID, in your developer portal and all provisioning profiles used in the app will be generated with both new and existing device ID. Afterward, the newly generated provisioning profiles are downloaded to App Center servers.

Reference:

<https://docs.microsoft.com/en-us/appcenter/distribution/groups>

**QUESTION 22**

You are developing an iOS application by using Azure DevOps.

You need to test the application manually on 10 devices without releasing the application to the public. 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 Microsoft Intune device compliance policy.
- B. Deploy a certificate from an internal certification authority (CA) to each device.
- C. Register the application in the iTunes store.
- D. Onboard the devices into Microsoft Intune.
- E. Distribute a new release of the application.
- F. Register the IDs of the devices in the Apple Developer portal.

**Correct Answer: E, F**

**Section:**

**QUESTION 23**

You have a private distribution group that contains provisioned and unprovisioned devices.

You need to distribute a new iOS application to the distribution group by using Microsoft Visual Studio App Center. What should you do?

- A. Select Register devices and sign my app.
- B. Create an active subscription in App Center Test.
- C. Create an unsigned build.
- D. Add the device owner to the collaborators group.

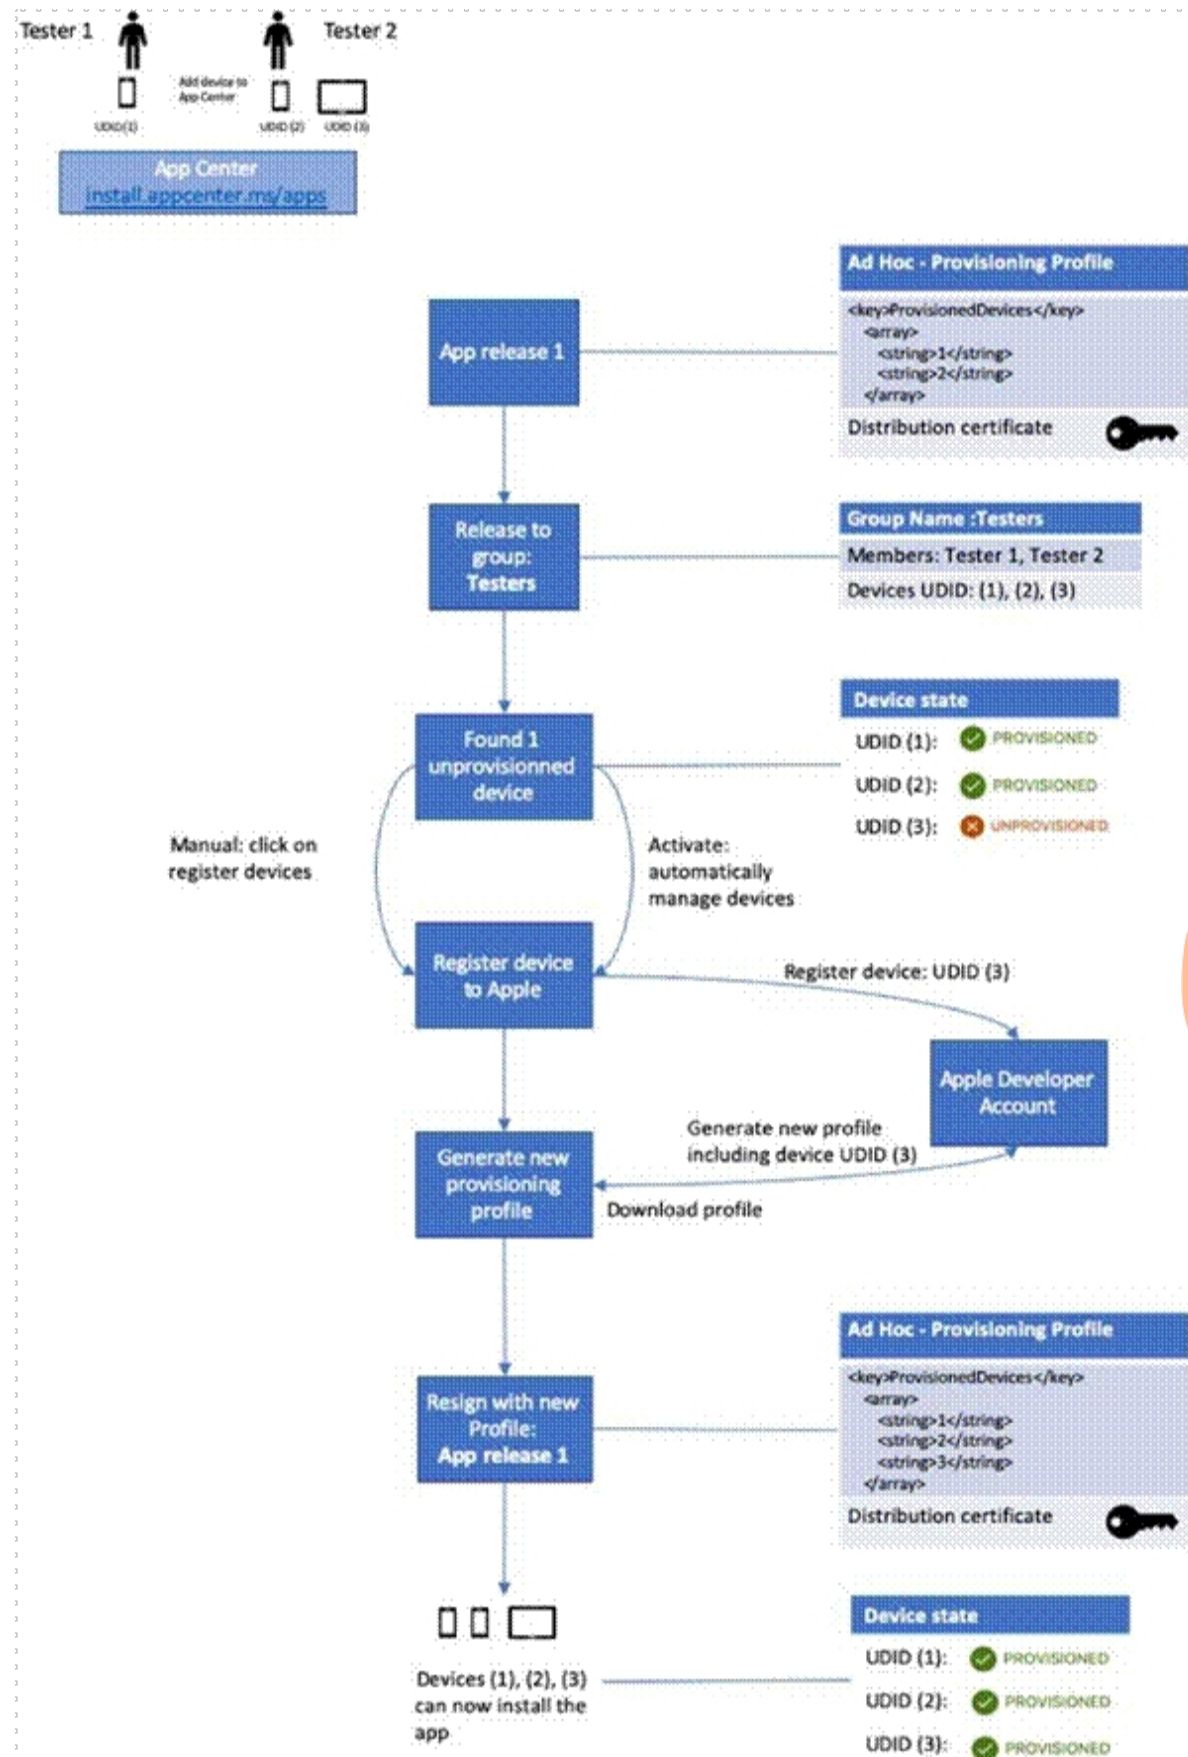
**Correct Answer: A**

**Section:**

**Explanation:**

The following diagram displays the entire app re-signing flow in App Center.





**vdumps**

Incorrect Answers:

C: The application build must be signed.

D: The device owner does not need to be added.

Reference:

<https://docs.microsoft.com/en-us/appcenter/distribution/auto-provisioning>

## QUESTION 24

### SIMULATION

You plan to deploy a website that will be hosted in two Azure regions.

You need to create an Azure Traffic Manager profile named az40011566895n1-tm in a resource group named RG1lod11566895. The solution must ensure that users will always connect to a copy of the website that is in the same country. To complete this task, sign in to the Microsoft Azure portal.

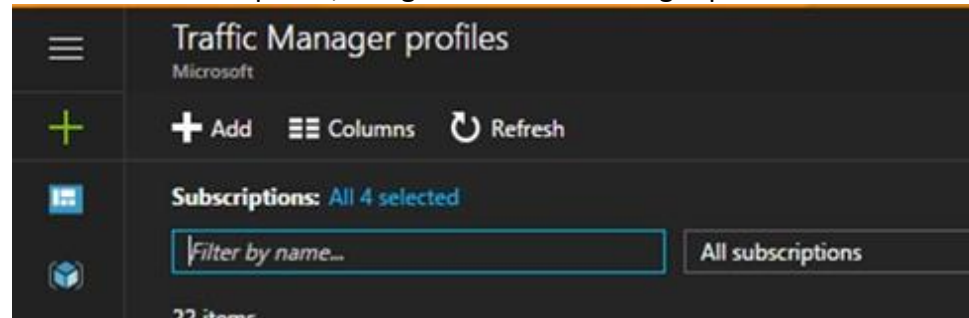
A. See solution below.

**Correct Answer: A**

**Section:**

**Explanation:**

1. Go to the Azure portal, navigate to Traffic Manager profiles and click on the Add button to create a routing profile.

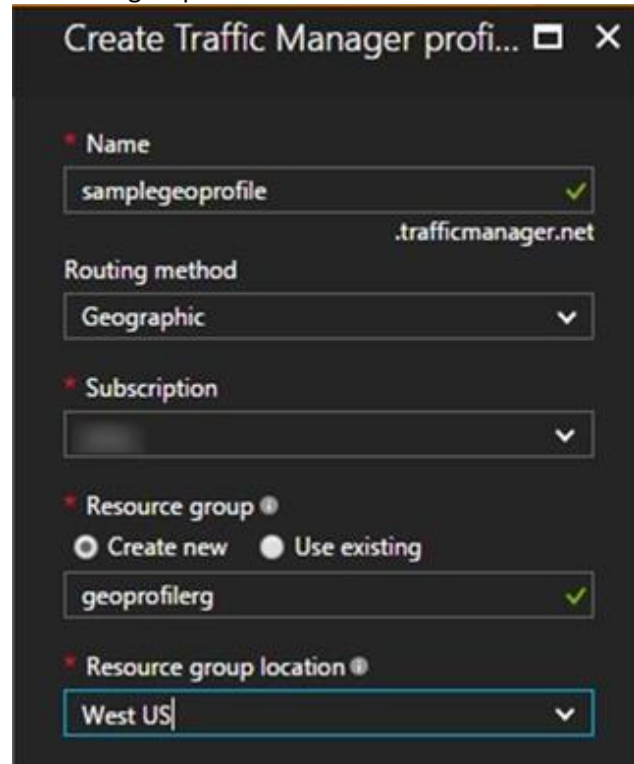


2. In the Create Traffic Manager profile, enter, or select these settings:

Name: az40011566895n1-tm

Routing method: Geographic

Resource group: RG1lod11566895



The logo for 'Vdumps' is displayed in a large, light gray font. The 'V' is stylized with a blue and orange gradient.

Note: Traffic Manager profiles can be configured to use the Geographic routing method so that users are directed to specific endpoints (Azure, External or Nested) based on which geographic location their DNS query originates from. This empowers Traffic Manager customers to enable scenarios where knowing a user's geographic region and routing them based on that is important.

Reference:

<https://azure.microsoft.com/en-us/blog/announcing-the-general-availability-of-geographic-routing-capability-in-azure-traffic-manager/>

## QUESTION 25

### HOTSPOT

Your company uses get 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 time 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.

Hot Area:

Answer Area

Branch lifetime:	<table border="1"><tr><td></td><td>▼</td></tr><tr><td>Long-lived</td><td></td></tr><tr><td>Short-lived</td><td></td></tr></table>		▼	Long-lived		Short-lived			
	▼								
Long-lived									
Short-lived									
Branch type:	<table border="1"><tr><td></td><td>▼</td></tr><tr><td>Master</td><td></td></tr><tr><td>Feature</td><td></td></tr><tr><td>Integration</td><td></td></tr></table>		▼	Master		Feature		Integration	
	▼								
Master									
Feature									
Integration									

Answer Area:

Answer Area

Branch lifetime:	<table border="1"><tr><td></td><td>▼</td></tr><tr><td>Long-lived</td><td></td></tr><tr><td>Short-lived</td><td></td></tr></table>		▼	Long-lived		Short-lived			
	▼								
Long-lived									
Short-lived									
Branch type:	<table border="1"><tr><td></td><td>▼</td></tr><tr><td>Master</td><td></td></tr><tr><td>Feature</td><td></td></tr><tr><td>Integration</td><td></td></tr></table>		▼	Master		Feature		Integration	
	▼								
Master									
Feature									
Integration									



Section:

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 26

HOTSPOT

You manage the get repository for a large enterprise application.



You need to minimize the data size of the repository.  
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**

```
git gc
```

- aggressive
- auto
- force
- no-prune

```
git --expire now
```

- merge
- prune
- rebase
- reset

Answer Area:



**Answer Area**

```
git gc
```

- aggressive
- auto
- force
- no-prune

```
git
```

- merge
- prune
- rebase
- reset

```
--expire now
```

**Section:**

**Explanation:**

Box 1: --aggressive

Cleanup unnecessary files and optimize the local repository:

```
git gc --aggressive
```

Box 2: prune

Prune all unreachable objects from the object database:

```
git prune
```

Reference:

<https://gist.github.com/Zoramite/2039636>

**QUESTION 27**

DRAG DROP

You provision an Azure Kubernetes Service (AKS) cluster that has RBAC enabled. You have a Helm chart for a client application.

You need to configure Helm and Tiller on the cluster and install the chart.

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:**



Commands	Answer Area
helm install	
kubectl create	
helm completion	
helm init	
helm serve	

Correct Answer:

Commands	Answer Area
	kubectl create
	helm init
helm completion	helm install
helm serve	



**Section:**

**Explanation:**

Step 1: Kubectl create

You can add a service account to Tiller using the --service-account <NAME> flag while you're configuring Helm (step 2 below). As a prerequisite, you'll have to create a role binding which specifies a role and a service account name that have been set up in advance.

Example: Service account with cluster-admin role

```
$ kubectl create -f rbac-config.yaml
```

```
serviceaccount "tiller" created
```

```
clusterrolebinding "tiller" created
```

```
$ helm init --service-account tiller
```

Step 2: helm init

To deploy a basic Tiller into an AKS cluster, use the helm init command.

Step 3: helm install

To install charts with Helm, use the helm install command and specify the name of the chart to install.

References:

<https://docs.microsoft.com/en-us/azure/aks/kubernetes-helm>

[https://docs.helm.sh/using\\_helm/#tiller-namespaces-and-rbac](https://docs.helm.sh/using_helm/#tiller-namespaces-and-rbac)

### QUESTION 28

DRAG DROP

You are implementing an Azure DevOps strategy for mobile devices using App Center.

You plan to use distribution groups to control access to releases.

You need to create the distribution groups shown in the following table.

Name	Use
Group1	Application testers who are invited by email
Group2	Early release users who use unauthenticated public links
Group3	Application testers for all the apps of your company

Which type of distribution group should you use for each group? To answer, drag the appropriate group types to the correct locations. Each group 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:

**Answer Area**

Group1:

Group2:

Group3:

Private

Public

Shared



Correct Answer:

**Answer Area**

Group1:   
 Group2:   
 Group3:

**Section:**

**Explanation:**

Box1: Private

In App Center, distribution groups are private by default. Only testers invited via email can access the releases available to this group.

Box 2: Public

Distribution groups must be public to enable unauthenticated installs from public links.

Box 3: Shared

Shared distribution groups are private or public distribution groups that are shared across multiple apps in a single organization.

Reference:

<https://docs.microsoft.com/en-us/appcenter/distribution/groups>



**QUESTION 29**

**HOTSPOT**

You need to create deployment files for an Azure Kubernetes Service (AKS) cluster. The deployments must meet the provisioning storage requirements shown in the following table.

Deployment	Requirement
Deployment 1	Use files stored on an SMB-based share from the container's file system.
Deployment 2	Use files on a managed disk from the container's file system.
Deployment 3	Securely access X.509 certificates from the container's file system.

Which resource type should you use for each deployment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Hot Area:**

Deployment 1:	<input type="text"/>	▼
	azurekeyvault-flexvolume	
	blobfuse-flexvol	
	kubernetes.io/azure-disk	
	kubernetes.io/azure-file	
	volume.beta.kubernetes.io/storage-provisioner	
Deployment 2:	<input type="text"/>	▼
	azurekeyvault-flexvolume	
	blobfuse-flexvol	
	kubernetes.io/azure-disk	
	kubernetes.io/azure-file	
	volume.beta.kubernetes.io/storage-provisioner	
Deployment 3:	<input type="text"/>	▼
	azurekeyvault-flexvolume	
	blobfuse-flexvol	
	kubernetes.io/azure-disk	
	kubernetes.io/azure-file	
	volume.beta.kubernetes.io/storage-provisioner	



Answer Area:

Deployment 1:	▼
	azurekeyvault-flexvolume
	blobfuse-flexvol
	kubernetes.io/azure-disk
	kubernetes.io/azure-file
	volume.beta.kubernetes.io/storage-provisioner
Deployment 2:	▼
	azurekeyvault-flexvolume
	blobfuse-flexvol
	kubernetes.io/azure-disk
	kubernetes.io/azure-file
	volume.beta.kubernetes.io/storage-provisioner
Deployment 3:	▼
	azurekeyvault-flexvolume
	blobfuse-flexvol
	kubernetes.io/azure-disk
	kubernetes.io/azure-file
	volume.beta.kubernetes.io/storage-provisioner



**Section:**

**Explanation:**

Deployment 1: Kubernetes.io/azure-file

You can use Azure Files to connect using the Server Message Block (SMB) protocol.

Deployment 2: Kubernetes.io/azure-disk

Deployment 3: azurekeyvault-flexvolume

azurekeyvault-flexvolume: Key Vault FlexVolume: Seamlessly integrate your key management systems with Kubernetes. Secrets, keys, and certificates in a key management system become a volume accessible to pods. Once the volume is mounted, its data is available directly in the container filesystem for your application.

Incorrect Answers:

blobfuse-flexvolume: This driver allows Kubernetes to access virtual filesystem backed by the Azure Blob storage.

References:

<https://docs.microsoft.com/bs-cyrl-ba/azure/aks/azure-files-dynamic-pv>

<https://docs.microsoft.com/en-us/azure/aks/azure-disks-dynamic-pv>

**QUESTION 30**

DRAG DROP

You manage the get 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:**

### Actions

### Answer Area

- Delete and recreate the repository.
- Run the `git reflog expire` command.
- Run the `git add .gitignore` command.
- Add Config.json to the .gitignore file.
- Run the `git commit` command.



Correct Answer:

### Actions

### Answer Area

- 
- Run the `git reflog expire` command.
- 
- 
- Run the `git commit` command.



- Delete and recreate the repository.
- Add Config.json to the .gitignore file.
- Run the `git add .gitignore` command.



Section:

Explanation:

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

### QUESTION 31

DRAG DROP



You have a project in Azure DevOps.

You need to associate an automated test to a test case.

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:

Actions	Answer Area
Debug the project	
Create a test project	
Create a work item	
Check in a project to the Azure DevOps repository	
Add the automated test to a build pipeline	

Correct Answer:

Actions	Answer Area
Debug the project	Create a test project
	Check in a project to the Azure DevOps repository
Create a work item	Add the automated test to a build pipeline

Section:

Explanation:

The process to associate an automated test with a test case is:

1. Create a test project containing your automated test. What types of tests are supported?
2. Check your test project into an Azure DevOps or Team Foundation Server (TFS) repository.
3. Create a build pipeline for your project, ensuring that it contains the automated test. What are the differences if I am still using a XAML build?
4. Use Visual Studio Enterprise or Professional 2017 or a later version to associate the automated test with a test case as shown below. The test case must have been added to a test plan that uses the build you just defined.

Reference:

**QUESTION 32**

DRAG DROP

You have an Azure Kubernetes Service (AKS) cluster.

You need to deploy an application to the cluster by using Azure DevOps.

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:

**Actions**

- Create a service account in the cluster.
- Create a service principal in Azure Active Directory (Azure AD).
- Add an Azure Function App for Container task to the deployment pipeline.
- Add a Helm package and deploy a task to the deployment pipeline.
- Add a Docker Compose task to the deployment pipeline.
- Configure RBAC roles in the cluster.

**Answer Area**



**Correct Answer:**

- Actions**
- Create a service account in the cluster.
  - Add an Azure Function App for Container task to the deployment pipeline.
  - Configure RBAC roles in the cluster.

- Answer Area**
- Create a service principal in Azure Active Directory (Azure AD).
  - Add a Helm package and deploy a task to the deployment pipeline.
  - Add a Docker Compose task to the deployment pipeline.

**Section:**

**Explanation:**

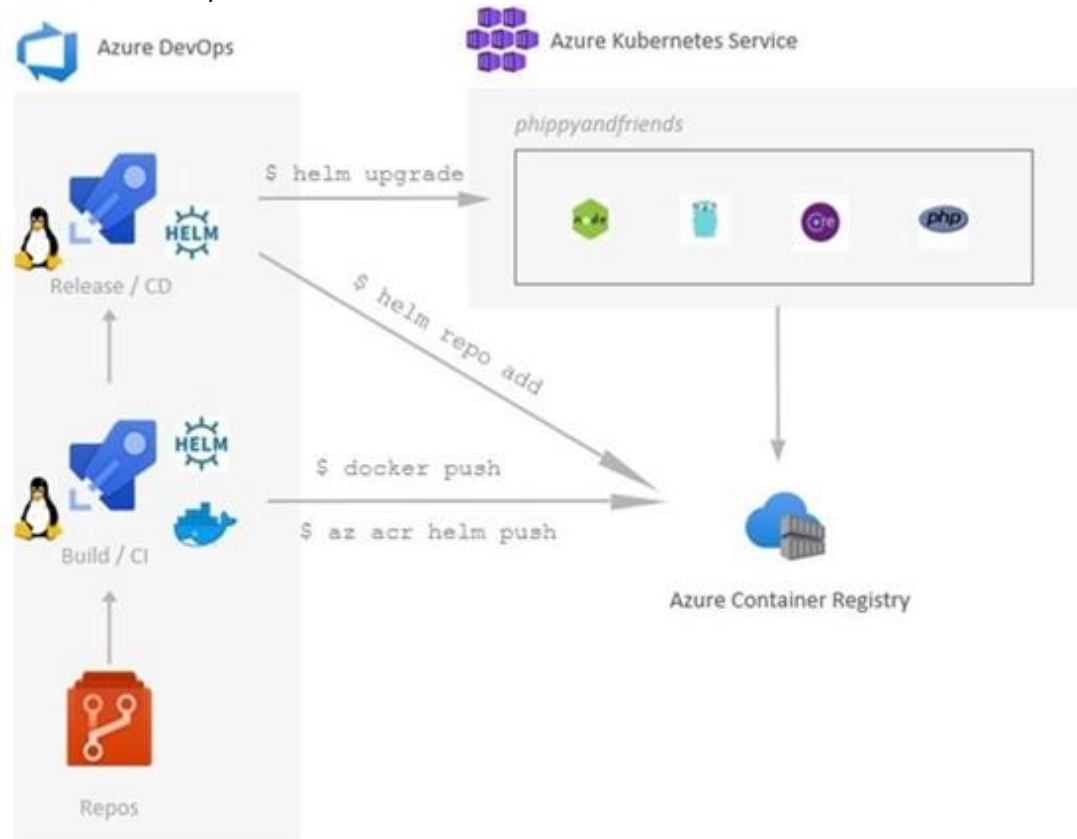
You can set up a CI/CD pipeline to deploy your apps on a Kubernetes cluster with Azure DevOps by leveraging a Linux agent, Docker, and Helm.

Step 1: Create a service principle in Azure Active Directory (Azure AD)

We need to assign 3 specific service principals with specific Azure Roles that need to interact with our ACR and our AKS. Create a specific Service Principal for our Azure DevOps pipelines to be able to push and pull images and charts of our ACR. Create a specific Service Principal for our Azure DevOps pipelines to be able to deploy our application in our AKS.

Step 2: Add a Helm package and deploy a task to the deployment pipeline

This is the DevOps workflow with containers:



Step 3: Add a Docker Compose task to the deployment pipeline.

Dockerfile file is a script leveraged by Docker, composed of various commands (instructions) and arguments listed successively to automatically perform actions on a base image in order to create a new Docker image by packaging the app.

Reference:

<https://cloudblogs.microsoft.com/opensource/2018/11/27/tutorial-azure-devops-setup-cicd-pipeline-kubernetes-docker-helm/>

### QUESTION 33

DRAG DROP

You are defining release strategies for two applications as shown in the following table.

Application name	Goal
App1	Failure of App1 has a major impact on your company. You need a small group of users, who opted in to a testing App1, to test new releases of the application.
App2	You need to minimize the time it takes to deploy new releases of App2, and you must be able to roll back as quickly as possible.

Which release strategy should you use for each application? To answer, drag the appropriate release strategies to the correct applications. Each release strategy 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:

#### Release Strategies

Blue/Green deployment

Canary deployment

Rolling deployment

Answer Area:

App1:



App2:

Correct Answer:

#### Release Strategies

Rolling deployment

Answer Area:

App1:

Canary deployment

App2:

Blue/Green deployment

Section:

**Explanation:**

**QUESTION 34**

DRAG DROP

You have an Azure DevOps organization named Contoso.

You have 10 Azure virtual machines that run Windows Server 2019. The virtual machines host an application that you build and deploy by using Azure Pipelines. Each virtual machine has the Web Server (IIS) role installed and configured.

You need to ensure that the web server configurations on the virtual machines is maintained automatically. The solution must provide centralized management of the configuration settings and minimize management overhead.

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:**

**Actions**

Create an Azure Automation account.

Install the custom Desired State Configuration (DSC) extension on the virtual machines.

Create a .zip file and upload it to Azure Blob storage.

Onboard the virtual machines to the Azure Automation account.

Complete the Desired State Configuration (DSC) configuration.

**Answer Area**



**Correct Answer:**

**Actions**

Create a .zip file and upload it to Azure Blob storage.

**Answer Area**

Create an Azure Automation account.

Install the custom Desired State Configuration (DSC) extension on the virtual machines.

Onboard the virtual machines to the Azure Automation account.

Complete the Desired State Configuration (DSC) configuration.



**Section:**

**Explanation:**

Step1: Create an Azure Automation account.

An Azure Automation account is required.

Step 2: Install the custom Desired State Configuration (DSC) extension on the virtual machines

Under the hood, and without an administrator having to remote into a VM, the Azure VM Desired State Configuration extension registers the VM with Azure Automation State Configuration.

Step 3: Onboard the virtual machines to the Azure Automation account.

Step 4: Complete the Desired State Configuration (DSC) configuration.

Create a DSC configuration.

Reference:

<https://docs.microsoft.com/en-us/azure/automation/automation-dsc-onboarding>

### QUESTION 35

#### HOTSPOT

You have a project in Azure DevOps.

You plan to create a build pipeline that will deploy resources by using Azure Resource Manager templates. The templates will reference secrets stored in Azure Key Vault.

You need to ensure that you can dynamically generate the resource ID of the key vault during template deployment.

What should you include in the 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": "2018-05-01",
  "name" : "secrets",
  "type": [
    "Microsoft.KeyVault/vaults",
    "Microsoft.Resources/deployment",
    "Microsoft.Subscription/subscriptions"
  ],
  "properties":{
    "mode" : "Incremental",
    [
      "deployment",
      "template",
      "templateLink"
    ] :{
contentVersion" : "1.0.0.0",
  "uri" : "[uri(parameters('_artifactsLocation'),
concat('./nested/sqlserver.json',
parameters('_artifactsLocationSasToken')))]"
},
  "parameters":{
    "secret":{
      "reference":{
        "keyVault":{
          "id": "[resourceId(parameters('vaultSubscription'),
parameters('vaultResourceGroupName'),
'Microsoft.KeyVault/vaults',
parameters('vaultName'))]"
        },
        "secretName": "[parameters('secretName')]"
      }
    }
  }
}
],
```



Answer Area:

## Answer Area

```
"resources": [
  {
    "apiVersion": "2018-05-01",
    "name": "secrets",
    "type": "Microsoft.KeyVault/vaults",
    "properties": {
      "mode": "Incremental",
      "templateLink": {
        "uri": "[uri(parameters('_artifactsLocation'),
concat('./nested/sqlserver.json',
parameters('_artifactsLocationSasToken')))]"
      }
    },
    "parameters": {
      "secret": {
        "reference": {
          "keyVault": {
            "id": "[resourceId(parameters('vaultSubscription'),
parameters('vaultResourceGroupName'),
'Microsoft.KeyVault/vaults',
parameters('vaultName'))]"
          },
          "secretName": "[parameters('secretName')]"
        }
      }
    }
  }
]
```

### Section:

### Explanation:

Box 1: "Microsoft.Resources/deployments"

Reference a secret with dynamic ID. You need to reference a key vault secret that varies based on the current deployment.

Example:

```
"resources": [
  {
    "apiVersion": "2018-05-01",
    "name": "dynamicSecret",
    "type": "Microsoft.Resources/deployments",
    "properties": {
      "mode": "Incremental",
      "templateLink": {
```

Box 2: "templateLink"

In your parent template, you add the linked template and pass in a parameter that contains the dynamically generated resource ID.

References: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-manager-keyvault-parameter>





### QUESTION 36

You have a project in Azure DevOps.

You plan to deploy a self-hosted agent by using an unattended configuration script.

Which two values should you define in the configuration script? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. authorization credentials
- B. the project name
- C. the deployment group name
- D. the organization URL
- E. the agent pool name

**Correct Answer: C, E**

**Section:**

**Explanation:**

### QUESTION 37

HOTSPOT

You have an Azure virtual machine named VM1 that runs Linux.

You plan to deploy the Desired State Configuration (DSC) extension to VM1.

You need to grant the Log Analytics agent the appropriate directory permissions.

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

NOTE: Each correct selection is worth one point.



**Hot Area:**

#### Answer Area

```
setfacl -m u:omsagent:
```

r	/lib
x	/etc
rx	/tmp
rwX	/usr

**Answer Area:**

## Answer Area

```
setfacl -m u:omsagent:
```

r	/lib
x	/etc
rx	/tmp
rwX	/usr

### Section:

### Explanation:

Box 1: rwx

The Log Analytics agent for Linux runs as the omsagent user. To grant >write permission to the omsagent user, run the command `setfacl -m u:omsagent:rwx /tmp`.

Box 2: /tmp

Deploying DSC to a Linux node uses the /tmp folder.

Reference:

<https://docs.microsoft.com/en-us/azure/automation/automation-dsc-onboarding>

### QUESTION 38

#### HOTSPOT

You are using PowerShell to administer Azure Log Analytics workspaces.

You need to list the available workspaces and their properties.

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

NOTE: Each correct selection is worth one point.



### Hot Area:

## Answer Area

		Microsoft.OperationalInsights/workspaces - ExpandProperties
Get-AzResource	-ResourceGroupName	
Get-AzResourceGroup	-ResourceId	
Get-AzResourceProvider	-ResourceType	

### Answer Area:

### Answer Area

Get-AzResource	-ResourceGroupName	Microsoft.OperationalInsights/workspaces - ExpandProperties
Get-AzResourceGroup	-ResourceId	
Get-AzResourceProvider	-ResourceType	

#### Section:

#### Explanation:

Box 1: Get-AzResource

Use the following command to examine the access control mode for all workspaces in the subscription:

PowerShell

```
Get-AzResource -ResourceType Microsoft.OperationalInsights/workspaces -ExpandProperties | foreach {$_.Name + " : " + $_.Properties.features.enableLogAccessUsingOnlyResourcePermissions
```

Box 2: -ResourceType

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/logs/manage-access>

#### QUESTION 39

You have Azure Pipelines and GitHub integrated as a source code repository.

The build pipeline has continuous integration enabled.

You plan to trigger an automated build whenever code changes are committed to the repository.

You need to ensure that the system will wait until a build completes before queuing another build.

What should you implement?

- A. path filters
- B. batch changes
- C. scheduled builds
- D. branch filters

**Correct Answer: B**

#### Section:

#### Explanation:

Batching CI runs

If you have many team members uploading changes often, you may want to reduce the number of runs you start. If you set batch to true, when a pipeline is running, the system waits until the run is completed, then starts another run with all changes that have not yet been built.

Example:

```
# specific branch build with batching
```

```
trigger:
```

```
batch: true
```

```
branches:
```

```
include:
```

```
- master
```

To clarify this example, let us say that a push A to master caused the above pipeline to run. While that pipeline is running, additional pushes B and C occur into the repository. These updates do not start new independent runs



immediately. But after the first run is completed, all pushes until that point of time are batched together and a new run is started.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/repos/github>

#### QUESTION 40

You are using GitHub as a source code repository.

You create a client-side get hook on the commit-msg event. The hook requires that each commit message contain a custom work item tag.

You need to make a commit that does not have a work item tag.

Which get commit parameter should you use?

- A. --squash
- B. --no-verify
- C. --message "
- D. --no-post-rewrite

**Correct Answer: B**

**Section:**

**Explanation:**

The commit-msg hook is invoked by git-commit and git-merge, and can be bypassed with the --no-verify option.

Reference:

<https://git-scm.com/docs/githooks>

#### 02 - Define and implement a continuous delivery and release management strategy

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.

Overview

Litware, Inc. is an independent software vendor (ISV). Litware has a main office and five branch offices.

Existing Environment

Application Architecture

The company's primary application is a single monolithic retirement fund management system based on ASP.NET web forms that use logic written in VB.NET. Some new sections of the application are written in C#.

Variations of the application are created for individual customers. Currently, there are more than 80 live code branches in the application's code base.

The application was developed by using Microsoft Visual Studio. Source code is stored in Team Foundation Server (TFS) in the main office. The branch offices access the source code by using TFS proxy servers.

Architectural Issues

Litware focuses on writing new code for customers. No resources are provided to refactor or remove existing code. Changes to the code base take a long time, as dependencies are not obvious to individual developers.

Merge operations of the code often take months and involve many developers. Code merging frequently introduces bugs that are difficult to locate and resolve.

Customers report that ownership costs of the retirement fund management system increase continually. The need to merge unrelated code makes even minor code changes expensive.

Customers report that bug reporting is overly complex.

Requirements

Planned Changes

Litware plans to develop a new suite of applications for investment planning. The investment planning applications will require only minor integration with the existing retirement fund management system.

The investment planning applications suite will include one multi-tier web application and two iOS mobile applications. One mobile application will be used by employees; the other will be used by customers.

Litware plans to move to a more agile development methodology. Shared code will be extracted into a series of packages.

Litware has started an internal cloud transformation process and plans to use cloud-based services whenever suitable.

Litware wants to become proactive in detecting failures, rather than always waiting for customer bug reports.

#### Technical Requirements

The company's investment planning applications suite must meet the following requirements:

New incoming connections through the firewall must be minimized.

Members of a group named Developers must be able to install packages.

The principle of least privilege must be used for all permission assignments.

A branching strategy that supports developing new functionality in isolation must be used.

Members of a group named Team Leaders must be able to create new packages and edit the permissions of package feeds. Visual Studio App Center must be used to centralize the reporting of mobile application crashes and device types in use. By default, all releases must remain available for 30 days, except for production releases, which must be kept for 60 days. Code quality and release quality are critical. During release, deployments must not proceed between stages if any active bugs are logged against the release. The mobile applications must be able to call the share pricing service of the existing retirement fund management system. Until the system is upgraded, the service will only support basic authentication over HTTPS. The required operating system configuration for the test servers changes weekly. Azure Automation State Configuration must be used to ensure that the operating system on each test server is configured the same way when the servers are created and checked periodically.

#### Current Technical Issue

The test servers are configured correctly when first deployed, but they experience configuration drift over time. Azure Automation State Configuration fails to correct the configurations.

Azure Automation State Configuration nodes are registered by using the following command.

```
Register-AzureRmAutomationDscNode
  -ResourceGroupName 'TestResourceGroup'
  -AutomationAccountName 'LitwareAutomationAccount'
  -AzureVMName $vmname
  -ConfigurationMode 'ApplyOnly'
```

#### QUESTION 1

To resolve the current technical issue, what should you do to the Register-AzureRmAutomationDscNode command?

- A. Change the value of the ConfigurationMode parameter.
- B. Replace the Register-AzureRmAutomationDscNode cmdlet with Register-AzureRmAutomationScheduledRunbook
- C. Add the AllowModuleOverwrite parameter.
- D. Add the DefaultProfile parameter.

**Correct Answer: A**

**Section:**

**Explanation:**

Change the ConfigurationMode parameter from ApplyOnly to ApplyAndAutocorrect.

The Register-AzureRmAutomationDscNode cmdlet registers an Azure virtual machine as an APS Desired State Configuration (DSC) node in an Azure Automation account. Scenario: Current Technical Issue

The test servers are configured correctly when first deployed, but they experience configuration drift over time. Azure Automation State Configuration fails to correct the configurations. Azure Automation State Configuration nodes are registered by using the following command.

```
Register-AzureRmAutomationDscNode
  -ResourceGroupName 'TestResourceGroup'
  -AutomationAccountName 'LitwareAutomationAccount'
  -AzureVMName $vmname
  -ConfigurationMode 'ApplyOnly'
```

References: <https://docs.microsoft.com/en-us/powershell/module/azurerms/azurerms/register-azurermsautomationdscnode?view=azurermps-6.13.0>

#### QUESTION 2

Which branching strategy should you recommend for the investment planning applications suite?

- A. release isolation
- B. main only
- C. development isolation
- D. feature isolation


**Correct Answer: D**

**Section:**

**Explanation:**

Scenario: A branching strategy that supports developing new functionality in isolation must be used.

Feature isolation is a special derivation of the development isolation, allowing you to branch one or more feature branches from main, as shown, or from your dev branches.



When you need to work on a particular feature, it might be a good idea to create a feature branch.

Incorrect Answers:

A: Release isolation introduces one or more release branches from main. The strategy allows concurrent release management, multiple and parallel releases, and codebase snapshots at release time.

B: The Main Only strategy can be folder-based or with the main folder converted to a Branch, to enable additional visibility features. You commit your changes to the main branch and optionally indicate development and release milestones with labels.

C: Development isolation: When you need to maintain and protect a stable main branch, you can branch one or more dev branches from main. It enables isolation and concurrent development. Work can be isolated in development branches by feature, organization, or temporary collaboration.

References:

<https://docs.microsoft.com/en-us/azure/devops/repos/tfvc/branching-strategies-with-tfvc?view=azure-devops>

### QUESTION 3

What should you use to implement the code quality restriction on the release pipeline for the investment planning applications suite?

- A. a pre-deployment approval
- B. a deployment gate
- C. a post-deployment approval
- D. a trigger

**Correct Answer: A**

**Section:**

**Explanation:**

When a release is created from a release pipeline that defines approvals, the deployment stops at each point where approval is required until the specified approver grants approval or rejects the release (or re-assigns the approval to another user).

Scenario: Code quality and release quality are critical. During release, deployments must not proceed between stages if any active bugs are logged against the release.

References: <https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/approvals>

### QUESTION 4

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 question on 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 sections 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 on 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.

Application Architecture

The company's primary application is a single monolithic retirement fund management system based on ASP.NET web forms that use logic written in VB.NET. Some new sections of the application are written in C#.

Variations of the application are created for individual customers. Currently, there are more than 80 live code branches in the application's code base.

The application was developed by using Microsoft Visual Studio. Source code is stored in Team Foundation Server (TFS) in the main office. The branch offices access the source code by using TFS proxy servers.

Architectural Issues

Litware focuses on writing new code for customers. No resources are provided to refactor or remove existing code. Changes to the code base take a long time, as dependencies are not obvious to individual developers.

Merge operations of the code often take months and involve many developers. Code merging frequently introduces bugs that are difficult to locate and resolve.

Customers report that ownership costs of the retirement fund management system increase continually. The need to merge unrelated code makes even minor code changes expensive.

Customers report that bug reporting is overly complex.

Planned changes

Litware plans to develop a new suite of applications for investment planning. The investment planning applications will require only minor integration with the existing retirement fund management system.

The investment planning applications suite will include one multi-tier web application and two iOS mobile application. One mobile application will be used by employees; the other will be used by customers.

Litware plans to move to a more agile development methodology. Shared code will be extracted into a series of packages.

Litware has started an internal cloud transformation process and plans to use cloud-based services whenever suitable.

Litware wants to become proactive in detecting failures, rather than always waiting for customer bug reports.

Technical requirements

The company's investment planning applications suite must meet the following requirements:

New incoming connections through the firewall must be minimized.

Members of a group named Developers must be able to install packages.

The principle of least privilege must be used for all permission assignments.

A branching strategy that supports developing new functionality in isolation must be used.

Members of a group named Team Leaders must be able to create new packages and edit the permissions of package feeds. Visual Studio App Center must be used to centralize the reporting of mobile application crashes and device types in use. By default, all releases must remain available for 30 days, except for production releases, which must be kept for 60 days. Code quality and release quality are critical. During release, deployments must not proceed between stages if any active bugs are logged against the release. The mobile applications must be able to call the share pricing service of the existing retirement fund management system. Until the system is upgraded, the service will only support basic authentication over HTTPS. The required operating system configuration for the test servers changes weekly. Azure Automation State Configuration must be used to ensure that the operating system on each test server is configured the same way when the servers are created and checked periodically.

Current Technical Issue

The test servers are configured correctly when first deployed, but they experience configuration drift over time. Azure Automation State Configuration fails to correct the configurations.

Azure Automation State Configuration nodes are registered by using the following command.

```
Register-AzureRmAutomationDscNode
-ResourceGroupName 'TestResourceGroup'
-AutomationAccountName 'LitwareAutomationAccount'
-AzureVMName $vmname
-ConfigurationMode 'ApplyOnly'
```

How should you configure the release retention policy for the investment planning applications suite? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Hot Area:**

## Answer Area

Global release:  ▼

Set the default retention policy to 30 days.
Set the maximum retention policy to 30 days.
Set the stage retention policy to 30 days.
Set the stage retention policy to 60 days.

Production stage:  ▼

Set the default retention policy to 30 days.
Set the maximum retention policy to 60 days.
Set the stage retention policy to 30 days.
Set the stage retention policy to 60 days.

Answer Area:

## Answer Area

Global release:  ▼

Set the default retention policy to 30 days.
Set the maximum retention policy to 30 days.
Set the stage retention policy to 30 days.
Set the stage retention policy to 60 days.

Production stage:  ▼

Set the default retention policy to 30 days.
Set the maximum retention policy to 60 days.
Set the stage retention policy to 30 days.
Set the stage retention policy to 60 days.

**Section:**

**Explanation:**

Scenario: By default, all releases must remain available for 30 days, except for production releases, which must be kept for 60 days.

Box 1: Set the default retention policy to 30 days

The Global default retention policy sets the default retention values for all the build pipelines. Authors of build pipelines can override these values.

Box 2: Set the stage retention policy to 60 days

You may want to retain more releases that have been deployed to specific stages.

References: <https://docs.microsoft.com/en-us/azure/devops/pipelines/policies/retention>

**QUESTION 5**

**HOTSPOT**

Where should the build and release agents for the investment planning application suite run? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.





Hot Area:

### Answer Area

Build agent:

- A hosted service
- A source control system
- The developers' computers

Release agent:

- A hosted service
- A source control system
- The developers' computers

Answer Area:

### Answer Area

Build agent:

- A hosted service
- A source control system
- The developers' computers

Release agent:

- A hosted service
- A source control system
- The developers' computers

Section:

Explanation:

Box 1: A source control system

A source control system, also called a version control system, allows developers to collaborate on code and track changes. Source control is an essential tool for multi-developer projects. Box 2: A hosted service

To build and deploy Xcode apps or Xamarin.iOS projects, you'll need at least one macOS agent. If your pipelines are in Azure Pipelines and a Microsoft-hosted agent meets your needs, you can skip setting up a self-hosted macOS agent.

Scenario: The investment planning applications suite will include one multi-tier web application and two iOS mobile applications. One mobile application will be used by employees; the other will be used by customers.

References:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/v2-osx?view=azure-devops>

Exam K

### QUESTION 1

You use get for source control.

You need to commit a 3-G3 ZIP file that contains virtual machines used for testing. The solution must meet the following requirements:

The file must be versioned.

The file must be associated with the corresponding code commits.

Which two actions should you include in the solution? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Store files in Azure Storage and enable blob versions.
- B. Install the get IFS extension and associate the extension to ZIP files.
- C. Use GZip to compress the file before committing the file.
- D. Install the git-stash extension and associate the extension to ZIP files.
- E. Install the git-fat extension and associate the extension to ZIP files.

**Correct Answer: B, D**

**Section:**

### QUESTION 2

You have the services shown in the following table.

Name	Interface type
Service1	HTTP
Service2	HTTPS

You manage a project by using Azure Boards.

You need to notify the services Of build Status changes.

Which services can be notified by using a webhook?

- A. Service1 only
- B. Service2 only
- C. Service1 and Service2 only

**Correct Answer: C**

**Section:**

### QUESTION 3

You are designing a YAML template for use with Azure Pipelines. The template Will include the Outputfile parameter.

Which two methods can you use to reference the parameter? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.



A. `$(parameters['outputfile'])`

B. `${parameters.outputfile}`

C. `$(parameters.outputfile)`

D. `$(parameters[outputfile])`

E. `${parameters['outputfile']}`

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

**Correct Answer: C, D**

**Section:**

#### QUESTION 4

You have a GitHub Enterprise account.  
You need to enable push protection for secret scanning of the account repositories.  
What should you do first?

- A. Purchase Premium Plus support.
- B. Enforce multi-factor authentication (MFA).
- C. Purchase a GitHub Advanced Security license.
- D. Create an access policy for secrets.

**Correct Answer: C**

**Section:**

#### QUESTION 5

Your company is building a new solution in Java.  
The company currently uses a SonarQube server to analyze the code of .NET solutions.  
You need to analyze and monitor the code quality of the Java solution.  
Which task types should you add to the build pipeline?

- A. Chef



- B. Gradle
- C. Octopus
- D. Gulp

**Correct Answer: B**

**Section:**

**Explanation:**

SonarQube is a set of static analyzers that can be used to identify areas of improvement in your code. It allows you to analyze the technical debt in your project and keep track of it in the future. With Maven and Gradle build tasks, you can run SonarQube analysis with minimal setup in a new or existing Azure DevOps Services build task.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/java/sonarqube?view=azure-devops>

**QUESTION 6**

**HOTSPOT**

You have an Azure subscription that contains an Azure key vault named Vault1, an Azure pipeline named Pipeline1, and an Azure SQL database named DB1.

Pipeline1 is used to deploy an app that will authenticate to DB1 by using a password.

You need to store the password in Vault1. The solution must ensure that the password can be accessed by Pipeline1.

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**



**Answer Area:**

**Answer Area**

Store the password as a:

- Secret
- Certificate
- Key

Secret

Grant Pipeline1 access to Vault1 by modifying the:

- Access policies
- Access control (IAM) settings

Access policies

- Security settings

**Section:**

**Explanation:**

**QUESTION 7**

You use Calendar Versioning (CalVer) for code assets. You need to store an optional tag of beta as part of the version. Which part of the version should you use for the tag?

- A. micro
- B. minor
- C. major
- D. modifier

**Correct Answer: D**

**Section:**

**QUESTION 8**

You have a GitHub repository that contains multiple versions of an Azure Pipelines template. You plan to deploy multiple pipelines that will use a template stored in the repository. You need to ensure that you use a fixed version of the template. What should you use to reference which version of the template repository to use?

- A. the runner
- B. the branch
- C. the SHA-based hashes
- D. the serial

**Correct Answer: B**

**Section:**

**QUESTION 9**

You have a project in Azure DevOps named Project1 that references an Azure Artifacts feed named Feed1. You have a package named Package1 that has the versions shown in the following table.



Version	Description
1.0.3	Manually pushed to Feed1
1.4.0	Manually pushed to Feed1
2.0.0	Available from an upstream source
2.3.1	Saved from an upstream source

You need to perform a build of Project1. Which version of Package1 will be used?

- A. 1.0.3
- B. 1.4.0
- C. 2.0.0
- D. 2.3.1

**Correct Answer: D**

**Section:**

**QUESTION 10**

HOTSPOT

You have an Azure web app named webapp1 that uses the .NET Core runtime stack. You have an Azure Application Insights resource named AppInsights1 that collects telemetry data generated by webapp1.

You plan to deploy webapp1 by using an Azure DevOps pipeline.

You need to modify the sampling rate of the telemetry data processed by AppInsights1 without having to redeploy webapp1 after each modification.

What should you do? To answer, select the appropriate options in the answer area.

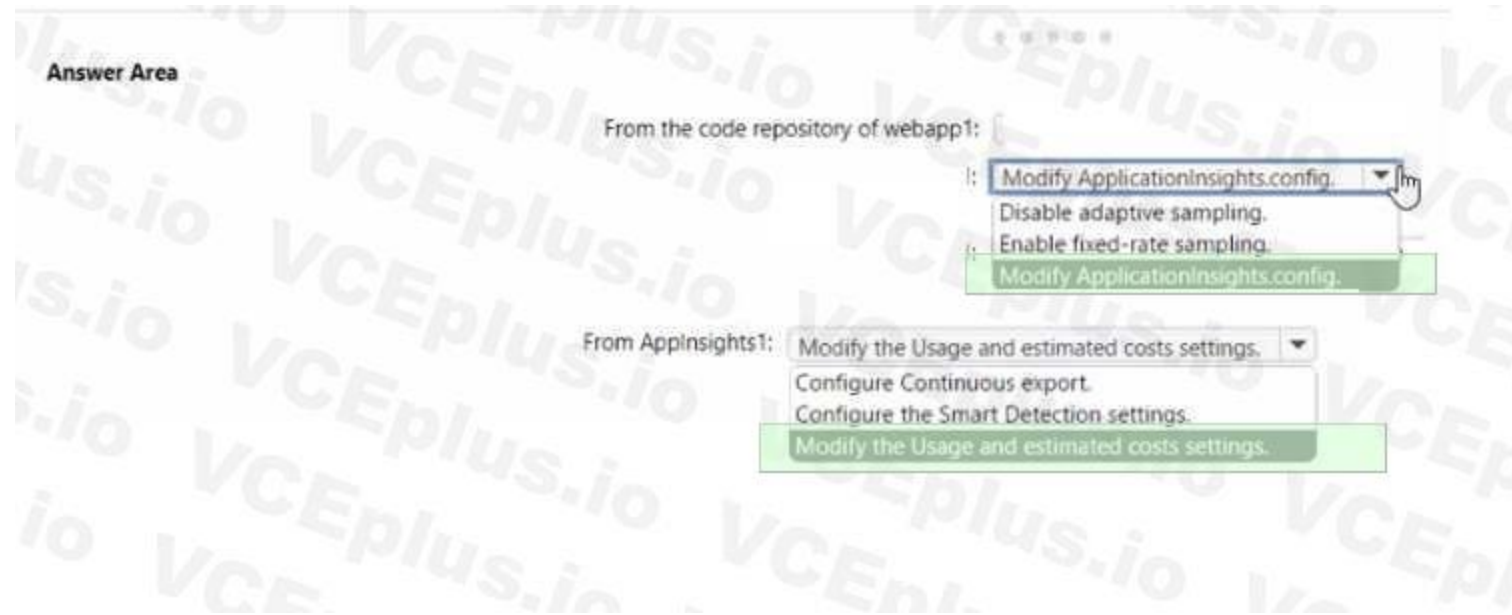
NOTE: Each correct selection is worth one point.

**Hot Area:**

The screenshot shows the 'Hot Area' for configuring Application Insights. It contains two dropdown menus:

- From the code repository of webapp1:**
  - Modify ApplicationInsights.config.
  - Disable adaptive sampling.
  - Enable fixed-rate sampling.
  - Modify ApplicationInsights.config.
- From AppInsights1:**
  - Modify the Usage and estimated costs settings.
  - Configure Continuous export.
  - Configure the Smart Detection settings.
  - Modify the Usage and estimated costs settings.

**Answer Area:**



**Section:**

**Explanation:**

**QUESTION 11**

You are creating a dashboard in Azure Boards.  
You need to visualize the time from when work starts on a work item until the work item is closed.  
Which type of widget should you use?

- A. cycle time
- B. velocity
- C. cumulative flow
- D. lead time

**Correct Answer: D**

**Section:**

**QUESTION 12**

You manage projects by using Azure Boards.  
You have a current work item name itemA that is dependant on a work item named item3.  
You need to define the dependency for item

- A. What should you do in the web portal for Azure DevOps?
- B. From Backlogs, open the context menu, select Add link and then select item3. Set Link type to Related and add the ID of itemA
- C. From itemA open the Links tab, and then select Add link. Set Link type to Successor and add the ID of itemB.
- D. From Queries, open the context menu, select Add link, and then select Existing item. Set Link type to Affected By and add the ID of itemB.
- E. From itemA, open the Links tab, and then select Add link. Set Link type to References and add the ID Of itemB.

**Correct Answer: B**

**Section:**

**QUESTION 13**

You use Azure Pipelines to build and release application code, The pipelines include validation tests that must be completed successfully before deployment proceeds from the test stage to production. You discover inconsistent test outcomes for the same source code.



You need to validate the test logic.  
What should you do?

- A. Decrease the test pass rate.
- B. Configure a parallel test runner.
- C. Enable flaky test detection.
- D. Install the Analytics extension.

**Correct Answer: B**  
**Section:**

#### QUESTION 14

You are integrating an Azure Boards project and a GitHub repository.  
You need to authenticate Azure Boards to GitHub.  
Which two authentication methods can you use? Each correct answer presents a complete solution.  
NOTE: Each correct selection is worth one point.

- A. a trusted root certificate
- B. a publisher certificate
- C. Azure Active Directory (Azure AD)
- D. GitHub user credentials
- E. a personal access token (PAT)

**Correct Answer: C, D**  
**Section:**



#### QUESTION 15

You have an Azure subscription that contains Azure DevOps build pipelines.  
You to implement pipeline caching by using the cache task  
HOW should you complete the YAML definition? TO answer, select the appropriate options in the answer are a.

A.

```
inputs:
  key: ["yarn" | "$(Agent.OS)" | yarn.lock']
  path: $(YARN_CACHE_FOLDER)
displayName: Cache Yarn packages
- script: yarn --frozen-lockfile
```

**Correct Answer: A**  
**Section:**

#### QUESTION 16

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 need to recommend an integration strategy for the build process of a Java application. The solution must meet the following requirements:  
The build must access an on-premises dependency management system.  
The build outputs must be stored as Server artifacts in Azure DevOps. The source code must be stored in a get repository in Azure DevOps.  
Solution: Configure the build pipeline to use a Microsoft-hosted agent pool running a Linux image. Include the Java Tool Installer task in the build pipeline.  
Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section:**

**Explanation:**

To run your jobs, you'll need at least one agent. A Linux agent can build and deploy different kinds of apps, including Java and Android apps.

If your pipelines are in Azure Pipelines and a Microsoft-hosted agent meets your needs, you can skip setting up a private Linux agent.

The Azure Pipelines agent pool offers several virtual machine images to choose from, each including a broad range of tools and software. We support Ubuntu, Red Hat, and CentOS.

Reference: <https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/v2-linux?view=azure-devops> <https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/hosted?view=azure-devops&tabs=yaml>

#### QUESTION 17

DRAG DROP

You are using the Dependency Tracker extension in a project in Azure DevOps.

You generate a risk graph for the project.

What should you use in the risk graph to identify the number of dependencies and the risk level of the project? To answer, drag the appropriate elements to the correct data points. Each element 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:**

The screenshot shows a drag-and-drop interface. On the left, under the heading 'Elements', there are four rectangular boxes: 'Link color', 'Link length', 'Link width', and 'Node color'. On the right, under the heading 'Answer Area', there are two dashed rectangular boxes. The top box is labeled 'Number of dependencies' and the bottom box is labeled 'Risk level'. The background of the interface has a faint, repeating watermark of the word 'QUESTION'.

**Correct Answer:**



**Section:**

**Explanation:**

Box 1: Link width

The width of the lines indicates how many dependencies exist in that area, the thicker the link the more dependencies as indicated in the legend.

Box 2: Link color

Reference: <https://docs.microsoft.com/en-us/azure/devops/boards/extensions/dependency-tracker?view=azure-devops#risk-graph>

**QUESTION 18**

You use release pipelines in Azure Pipelines to deploy an app. Secrets required by the pipeline are stored as pipeline variables. Logging of commands is enabled for the Azure Pipelines agent. You need to prevent the values of the secrets from being logged.

What should you do?

- A. Store the secrets in the environment variables instead of the pipeline variables.
- B. Pass the secrets on the command line instead of in the pipeline variables.
- C. Apply a prefix of secret to the name of the variables.
- D. Echo the values of the secrets to the command line.

**Correct Answer: A**

**Section:**

**Explanation:**

Don't set secret variables in your YAML file. Operating systems often log commands for the processes that they run, and you wouldn't want the log to include a secret that you passed in as an input. Use the script's environment or map the variable within the variables block to pass secrets to your pipeline.

Incorrect Answers:

B: Never pass secrets on the command line.

C: Adding a prefix does not make the variable a secret. The `issecret` property makes it secret but does not prevent logging of the secret. D: Never echo secrets as output.

Reference:

<https://docs.microsoft.com/en-us/azure/devops/pipelines/process/variables?view=azure-devops&tabs=yaml%2Cbatch> <https://docs.microsoft.com/en-us/azure/devops/pipelines/scripts/loggingcommands?view=azure-devops&tabs=bash>

**QUESTION 19**

You are creating a YAML-based Azure pipeline to deploy an Azure Data factory instance that has the following requirements;

- If a Data Factory instance exists already, the instance must be overwritten.
- No other resources in a resource group named Fabrikam must be affected.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
steps:
- task: AzureResourceManagerTemplateDeployment@3
  inputs:
    deploymentScope: 'Resource Group'
    azureResourceManagerConnection: 'Fabrikam Corporate(a41fb3ed-a2aa-42f0-a7ac-8fcc6ef0c5db)'
    subscriptionId: 'a41de0ed-a2aa-42f0-a7ac-8fcc6ef0c5db'
    action: 
  resourceGroupName: 'Fabrikam'
  location: 'West US'
  templateLocation: 'Linked artifact'
  deploymentMode: 
```

A.

```
steps:
- task: AzureResourceManagerTemplateDeployment@3
  inputs:
    deploymentScope: 'Resource Group'
    azureResourceManagerConnection: 'Fabrikam Corporate(a41fb3ed-a2aa-42f0-a7ac-8fcc6ef0c5db)'
    subscriptionId: 'a41de0ed-a2aa-42f0-a7ac-8fcc6ef0c5db'
    action: 'Create Or Update Resource Group'
  resourceGroupName: 'Fabrikam'
  location: 'West US'
  templateLocation: 'Linked artifact'
  deploymentMode: 'Incremental'
```



**Correct Answer: A**

**Section:**

### QUESTION 20

DRAG DROP

You have an Azure Kubernetes Service (AKS) pod that hosts an app named App1.

You need to configure the AKS container to restart automatically if the container stops responding. The solution must check the status of App1 once every three seconds.

How should you complete the deployment? To answer, drag the appropriate values to the correct targets. 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:**

Values

Answer Area

```
apiVersion: 2019-12-01
location: eastus
name: App1
properties:
  containers:
    - name: container1
      properties:
        image: mycompany/myImage:1.0.1
      ports: []
      resources:
        resources:
          requests:
            cpu: 1.0
            memoryInGB: 1.5
            
          httpGet:
            path: /
            port: 8080
            
            timeoutSeconds: 1
      osType: linux
      restartPolicy: 
tags: null
type: Microsoft.ContainerInstance/containerGroups
...
```

Correct Answer:

Values

Always

InitialDelaySeconds

livenessProbe

Never

successThreshold

Answer Area

```

apiVersion: 2019-12-01
location: eastus
name: App1
properties:
  containers:
  - name: container1
    properties:
      image: mycompany/myImage:1.0.1
      ports: []
      resources:
      resources:
        requests:
          cpu: 1.0
          memoryInGB: 1.5
      readinessProbe
        httpGet:
          path: /
          port: 8080
          value: 1
          timeoutSeconds: 1
      osType: linux
      restartPolicy: periodSeconds
tags: null
type: Microsoft.ContainerInstance/containerGroups
...

```

Section:  
Explanation:

**Values**

- Always
- initialDelaySeconds
- livenessProbe
- Never
- periodSeconds
- readinessProbe
- successThreshold

**Answer Area**

```

apiVersion: 2019-12-01
location: eastus
name: App1
properties:
  containers:
  - name: container1
    properties:
      image: mycompany/myimage:1.0.1
    ports: []
    resources:
      resources:
        requests:
          cpu: 1.0
          memoryInGB: 1.5
          readinessProbe
          httpGet:
            path: /
            port: 8080
            Value: 3
            timeoutSeconds: 1
    osType: linux
    restartPolicy: periodSeconds
  tags: null
type: Microsoft.ContainerInstance/containerGroups
...

```



**QUESTION 21**

DRAG DROP

Your company has a project in Azure DevOps named Project1.

All the developers at the company have Windows 10 devices.

You need to create a get repository for Project1. The solution must meet the following requirements:

- Support large binary files.
- Store binary files outside of the repository.
- Use a standard get workflow to maintain the metadata of the binary files by using commits to the repository.

Select and Place:

**Actions**

- Perform a custom installation of Git for Windows that includes Git Virtual File System (GVFS).
- Configure personal access token (PAT)-based authentication.
- Perform a custom installation of Git for Windows that includes Git Large File Storage (LFS).
- Configure SSH key-based authentication.
- Configure Git Large File Storage (LFS) file tracking.

**Answer Area**

- 1
- 2
- 3

Correct Answer:

Actions		Answer Area
Perform a custom installation of Git for Windows that includes Git Virtual File System (GVFS).	<input type="radio"/> <input type="radio"/>	1 Perform a custom installation of Git for Windows that includes Git Large File Storage (LFS).
Configure personal access token (PAT)-based authentication.		2 Configure SSH key-based authentication.
		3 Configure Git Large File Storage (LFS) file tracking.

**Section:**

**Explanation:**

Actions		Answer Area
Perform a custom installation of Git for Windows that includes Git Virtual File System (GVFS).	<input type="radio"/> <input type="radio"/>	1 Perform a custom installation of Git for Windows that includes Git Large File Storage (LFS).
Configure personal access token (PAT)-based authentication.		2 Configure SSH key-based authentication.
		3 Configure Git Large File Storage (LFS) file tracking.

**QUESTION 22**

You have a GitHub repository that contains the source code for an app.

You need to identify all the changes made between versions 1.4.16 and 1.6.12 of the source code.

How should you complete the get command? To answer, select the appropriate options in the answer area. `git _____ | helper-script > changes.txt`

NOTE: Each correct selection is worth one point.

**Answer Area**

A. `get diff v1.4.16 v1.6.12 | helper-script > changes.txt`

**Correct Answer: A**

**Section:**

**Explanation:**

This command will compare the changes made between versions 1.4.16 and 1.6.12 of the source code in your GitHub repository, pipe the output through the helper-script and save the result to a file called "changes.txt". Please note that, this command assumes that you have a helper-script that can handle get diff output as an input and processes it further. It is not a default get command.

**QUESTION 23**

You have an Azure virtual machine that is monitored by using Azure Monitor.

The virtual machine has the Azure Log Analytics agent installed.

You plan to deploy the Service Map solution from Azure Marketplace.

What should you deploy to the virtual machine to support the Service Map solution?

- A. the Telegraf agent
- B. the Azure Monitor agent
- C. the Dependency agent
- D. the Windows Azure diagnostics extension (WAD)

**Correct Answer: C**

**Section:**

**QUESTION 24**

DRAG DROP

You need to deploy a new project in Azure DevOps that has the following requirements:

- The lead developer must be able to create repositories, manage permissions, manage policies, and contribute to the repository.
- Developers must be able to contribute to the repository and create branches, but NOT bypass policies when pushing builds.
- Project managers must only be able to view the repository.
- The principle of least privilege must be used.

You create a new Azure DevOps project team for each role.

To which Azure DevOps groups should you add each team? To answer, drag the appropriate groups to the correct teams. Each group 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:**

**Azure DevOps groups**

- Build Administrators
- Contributors
- Project Administrators
- Project Collection Administrators
- Project Collection Valid Users

**Answer Area**

Project manager:

Lead developer:

Developer:



**Correct Answer:**

**Azure DevOps groups**

- Build Administrators
- 
- 
- 
- Project Collection Valid Users

**Answer Area**

Project manager:

Lead developer:

Developer:

**Section:**

**Explanation:**

**QUESTION 25**

You use GitHub for source control and project-related discussions.

You receive a notification when an entry is made to any team discussion.

You need to ensure that you receive email notifications only for discussions in which you commented or in which you are mentioned. Which two Notifications settings should you clear? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.



- A. Participating
- B. Automatically watch repositories
- C. Automatically watch teams
- D. Watching

**Correct Answer: B, D**

**Section:**

**QUESTION 26**

DRAG DROP

You have an Azure Repos repository named repo1.

You need to clone repo1. The solution must clone only a directory named src/web.

How should you complete the script? To answer, drag the appropriate values to the correct targets,

Each value may be used once, more than once, or not at all. You may need to drag the spirt bar between panes or scroll to view content. NOTE: Each correct selection is worth one point

**Select and Place:**

**Correct Answer:**

**Section:**

**Explanation:**

```

cd repos
scalar clone git@ssh.dev.azure.com:v3/organization/project/repo1
cd src/web
git sparse-checkout set repo1/src
...

```

**QUESTION 27**

DRAG DROP

You have an app named App1. You have a Log Analytics workspace named Workspace 1 that contains two tables named Events and Logs. App1 manages events in multiple locations and writes logs to Workspace1. You need to query Workspace1 for all log entries related to Asia that occurred during the last two days. In which order should you arrange the query statements? To answer, move all statements from the list of statements to the answer area and arrange them in the correct order.

Select and Place:

Statements	Answer Area
join ( Events	
) on RequestId	
where Timestamp > ago(2d)	
where continent == 'Asia'	
Logs	

Correct Answer:

Statements	Answer Area
	Logs
	where Timestamp > ago(2d)
	where continent == 'Asia'
	join ( Events
	) on RequestId

Section:

Explanation:

**QUESTION 28**

DRAG DROP

You have a project in Azure DevOps named Project1 that has a release pipeline in Azure Pipeline named ReleaseP1. you need to ensure that when a new release is generated for ReleaseP1, a new release note document is created. The release notes must contain new features and bug fixes. Which three actions should you perform in sequence? To answer, move the appropriate actions from the Most 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:

Articles

- Add a dashboard widget that retrieves the feature and bug fix information.
- Create a PowerShell task in ReleaseP1 that writes the retrieved data to a markdown file.
- Create a service principal.
- Create a personal access token (PAT).
- Create a query that retrieves the feature and bug fix information.

Answer Area

Correct Answer:

Articles

- Add a dashboard widget that retrieves the feature and bug fix information.
- Create a service principal.
- Create a personal access token (PAT).
- 

Answer Area

- Add a dashboard widget that retrieves the feature and bug fix information.
- Create a PowerShell task in ReleaseP1 that writes the retrieved data to a markdown file.
- Create a query that retrieves the feature and bug fix information.

Section:

Explanation:

#### QUESTION 29

You have an Azure subscription that contains multiple Azure pipelines.

You need to deploy a monitoring solution for the pipelines. The solution must meet the following requirements:

- Parse logs from multiple sources
- Identify the root cause of issues.

What advanced feature of a monitoring tool should you include in the solution?

- A. synthetic monitoring
- B. Alert Management
- C. analytics
- D. directed monitoring

Correct Answer: C

**Section:**

**Explanation:**

An analytics feature in a monitoring solution would allow you to parse logs from multiple sources and analyze them to identify the root cause of issues in your Azure pipelines. This feature would typically provide tools for searching, filtering, and visualizing log data, as well as for identifying patterns and anomalies. With analytics, you can also create custom dashboards and alerts to monitor your pipelines and quickly identify and troubleshoot any issues.

**QUESTION 30**

**DRAG DROP**

You have a web app named App1 that is hosted on multiple servers. App1 uses Application Insights in Azure Monitor. You need to compare the dairy CPU usage from the last week for all servers. How should you complete the query? To answer, drag the appropriate values to the correct targets.

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:**

**Values**

bin(timestamp,1d)

bin(timestamp,1h)

project timechart

render chart

render timechart

**Answer Area**

...

performanceCounters

| where counter == "% Processor Time"

| where timestamp >= ago{7d}

| summarize avg(value) by cloud\_RoleInstance,

|

...

**Correct Answer:**

**Values**

bin(timestamp,1h)

project timechart

render chart

**Answer Area**

...

performanceCounters

| where counter == "% Processor Time"

| where timestamp >= ago{7d}

| summarize avg(value) by cloud\_RoleInstance, bin(timestamp,1d)

| render timechart

...

**Section:**

**Explanation:**

**QUESTION 31**

DRAG DROP

You have a GitHub organization that contains three users named User 1, User2, and User3. You have a project that contains a repository named repo1. You need to configure permissions for repo1. The solution must meet the following requirements:

- Ensure that User 1 can actively push to repo1.
- Ensure that User2 can manage issues and pull requests for repo1.
- Ensure that User3 can manage repo1.
- Prevent User3 from accessing sensitive data in repo1.

Which role should you assign to each user?

**Select and Place:**

The screenshot shows a 'Select and Place' interface. On the left, under the heading 'Roles', there are five buttons: Admin, Maintain, Read, Triage, and Write. On the right, under the heading 'Answer Area', there are three boxes labeled 'User1: Role', 'User2: Role', and 'User3: Role'. A large watermark 'Vdumps' is visible in the center of the image.

**Correct Answer:**

The screenshot shows the 'Correct Answer' for the question. On the left, under the heading 'Roles', there are three buttons: Maintain, Triage, and an empty box. On the right, under the heading 'Answer Area', there are three boxes labeled 'User1: Read', 'User2: Write', and 'User3: Admin'. A large watermark 'Vdumps' is visible in the center of the image.

**Section:**

**Explanation:**

**QUESTION 32**

Your team uses Azure Pipelines to deploy applications.

You need to ensure that when a failure occurs during the build or release process, all the team members are notified by using Microsoft Teams. The solution must minimize development effort. What should you do?

- A. Install the Azure Boards app for Teams and configure a subscription to receive notifications in a channel.
- B. Use Azure Automation to connect to the Azure DevOps REST API and notify the team members.
- C. Use an Azure function to connect to the Azure DevOps REST API and notify the team members.
- D. Install the Azure Pipelines app for Teams and configure a subscription to receive notifications in a channel.

**Correct Answer: D**

**Section:**

**QUESTION 33**

You are automating the build process for a Java-based application by using Azure DevOps. You need to add code coverage testing and publish the outcomes to the pipeline. What should you use?

- A. Cobertura
- B. JUnit
- C. Coverage.py
- D. Bullseye Coverage

**Correct Answer: A**

**Section:**

**QUESTION 34**

You have project in Azure DevOps. You create the following template named Template1.yml.

```
steps:  
- script: npm install  
- script: yarn install  
- script: npm run compile
```

You create the following pipeline named File1.yml.

```
parameters:  
usersteps:  
- task: MyTask@1  
- script: echo Done
```

You need to ensure that Template1.yml runs before File1.yml. How should you update File1.yml?

- A. `parameters: usersteps: extends: template: template1.yml - task: MyTask@1 - script: echo Done`
- B. `extends: template: template1.yml parameters: usersteps: - task: MyTask@1 - script: echo Done`
- C. `parameters: usersteps: - template: template1.yml - task: MyTask@1 - script: echo Done`
- D. `template: template1.yml parameters: usersteps: - task: MyTask@1 - script: echo Done`

- A. Option A
- B. Option B
- C. Option C



D. Option D

**Correct Answer: B**

**Section:**

**QUESTION 35**

You have an Azure subscription that contains two resource groups named ContosoRG and ContosoDev, an Azure data factory named Contoso Data Factory, and a release pipeline in Azure Pipelines named Pipeline1. You plan to deploy Contoso Data Factory to ContosoRG by using Pipeline1.

You add the Azure Resource Manager (ARM) template deployment task shown the following exhibit.



A.



**Correct Answer: A**

**Section:**

**QUESTION 36**

You have an Azure subscription that contains an Azure pipeline named Pipeline1 and a GitHub repository named Repo1, Repo1 contains Bicep modules. Pipeline1 deploys Azure resources by using the Bicep modules. You need to ensure that all releases comply With Azure Policy before they are deployed to production. What should you do?

- A. Configure a deployment gate for Pipeline1 include the Azure DevOps Security and compliance assessment task.
- B. Create an Azure DevOps build runs on the creation of a pull request assesses the code for compliance.
- C. To Pipeline1, add a step that runs a What If deployment before the deployment step.
- D. Configure a deployment gate for Pipeline1 that uses Azure Automation to run a What If deployment

**Correct Answer: A**

**Section:**

**QUESTION 37**

You need to make a custom package available to all the developers. The package must be managed centrally, and the latest version must be available for consumption in Visual Studio automatically. Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Add the package URL to the Environment settings in Visual Studio.
- B. Create a get repository in Azure Repos.
- C. Add the package URL to the NuGet Package Manager settings in Visual Studio.
- D. Upload a package to a get repository.
- E. Create a new feed in Azure Artifacts.
- F. Publish the package to a feed.

**Correct Answer: A, B, E**

**Section:**

**QUESTION 38**

You have an Azure subscription that contains 50 virtual machines

You plan to manage the configuration of the virtual machines by using Azure Automation State Configuration. You need to create the Desired State Configuration (DSO configuration files.

How should structure the code blocks?

- A. Node>Configuration>Resource
- B. Configuration>Node> Resource
- C. Configuration>ReSource>Node
- D. Resource>Configuration>Node

**Correct Answer: B**

**Section:**

**Explanation:**

In Azure Automation State Configuration, the Desired State Configuration (DSC) configuration files are used to define the desired state of resources on a system. The structure of the code blocks in a DSC configuration file should be organized in a logical and meaningful way. One way to structure the code blocks is as follows: Configuration: This block defines the overall configuration, including any parameters that are used in the configuration. Node: This block defines the target node(s) for the configuration, typically specified by the hostname or IP address of the target system. Resource: This block defines the resources that are managed by the configuration, including the resource type, module, and properties. "A configuration script consists of the following parts: The Configuration block. This is the outermost script block. You define it by using the Configuration keyword and providing a name. In this case, the name of the configuration is MyDscConfiguration. One or more Node blocks. These define the nodes (computers or VMs) that you are configuring. In the above configuration, there is one Node block that targets a computer named TEST-PC1. The Node block can accept multiple computer names. One or more resource blocks. This is where the configuration sets the properties for the resources that it is configuring. In this case, there are two resource blocks, each of which call the WindowsFeature resource." <https://docs.microsoft.com/en-us/powershell/dsc/configurations/configurations?view=dsc-1.1#configuration-syntax>

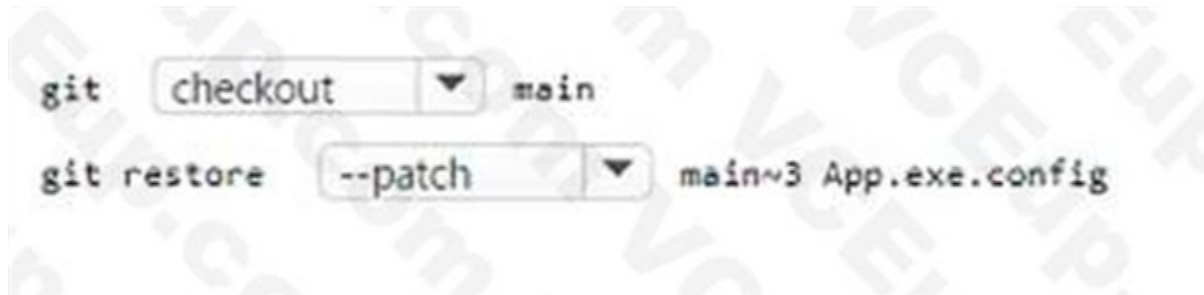


**QUESTION 39**

You use get for source control. You have an app named Appt.

In the main branch. you need to restore the third most recent revision of a file named App.exe.config. How should you complete command?

A.



**Correct Answer: A**

**Section:**

**QUESTION 40**

DRAG DROP

You have an Azure Repos repository that contains large PSD files. You need to configure get LFS to manage all the files. How should you complete the script? To answer, drag the appropriate access levels to the correct groups.

Each access 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:**



Values

```

git fetch
git lfs config add "*.psd"
git lfs migrate import --include="*.psd" --everything
git lfs track "*.psd"
git lfs update
git push

```

Answer Area

```

...
git lfs install
Value
git add .gitattributes
git commit -m "track *.psd files using Git LFS"
Value
Value
...

```

Correct Answer:

Values

```

git fetch
git lfs config add "*.psd"
git lfs track "*.psd"
git lfs migrate import --include="*.psd" --everything
git lfs update
git push

```

Answer Area

```

...
git lfs install
git lfs track "*.psd"
git add .gitattributes
git commit -m "track *.psd files using Git LFS"
git lfs migrate import --include="*.psd" --everything
git lfs update
...

```

Section:

Explanation:

#### QUESTION 41

You have an Azure subscription that contains multiple Azure pipelines.

You need to deploy a monitoring solution for the pi\*lines. The solution must meet the following requirements:

Parse logs from multiple sources. Identify the root cause of issues.

What advanced feature of a monitoring tool should include in the solution?

- A. directed monitoring
- B. synthetic monitoring
- C. analytics
- D. Alert Management

Correct Answer: B

Section:

#### QUESTION 42

Your team uses Azure Pipelines to deploy applications.

You need to ensure that when a failure occurs during the build or release process, all the team members are notified by using Microsoft Teams. The solution must minimize development effort. What should you do?

- A. Use Azure Automation to connect to the Azure DevOps REST API and notify the team members.
- B. Install the Azure Pipelines app for Teams and configure a subscription to receive notifications in a channel.
- C. Install the Azure Boards app for Teams and configure a subscription to receive notifications in a channel.
- D. Use an Azure function to connect to the Azure DevOps REST API and notify the team members.

Correct Answer: C



**Section:**

**QUESTION 43**

DRAG DROP

You have GitHub repository named repo1 that stores the code of named App1.

You need deploy workflow for repo1 by using GitHub Actions. The solution must meet the following requirements:

Scan on pushes to the main branch.

Scan on pull requests to the main branch.

Scan on pull requests to any branch that has a prefix of releases/.

Scan all the files in subdirectories of the src directory.

Exclude scanning of markdown files

**Select and Place:**

Values	Answer Area
<input type="text" value="- '**/*,.md'"/>	<pre>--- on:   push:     branches: [main]   pull_request:     branches:       - main   paths:   paths-ignore: ---</pre>
<input type="text" value="- '*,.md'"/>	
<input type="text" value="- 'release*'"/>	
<input type="text" value="- 'releases/**'"/>	
<input type="text" value="- 'src/**'"/>	
<input type="text" value="- 'src/**'"/>	



**Correct Answer:**

Values	Answer Area
<input type="text"/>	<pre>--- on:   push:     branches: [main]   pull_request:     branches:       - main   paths:     - 'release*'   paths-ignore:     - '**/*,.md' ---</pre>
<input type="text"/>	
<input type="text" value="- 'releases/**'"/>	
<input type="text" value="- 'src/**'"/>	
<input type="text" value="- 'src/**'"/>	

**Section:**

**Explanation:**

The screenshot shows a configuration interface with two columns: 'Values' and 'Answer Area'. The 'Values' column contains six text boxes with the following content: `.**/*.ed`, `*.ed`, `*release*`, `*releases/*`, `*src/*`, and `*src/**`. The 'Answer Area' column shows a configuration snippet with several fields: `on:`, `path:`, `branches: [main]`, `pull_request:`, `branches:`, `main`, `*release*`, `paths:`, `*.ed`, `paths-ignore:`, and `.**/*.ed`. The `*release*` and `.**/*.ed` values are highlighted with blue boxes, indicating they are the correct answers for the corresponding fields.

**QUESTION 44**

**DRAG DROP**

You have app named App1. You have a Log Analytics workspace named Workspace1 that contains two tables named Events and Logs. App1 manage events in multiple locations and writes logs to Workspace1. You need to query Workspace1 for all log entries related to Asia that occurred during the last two days. In which order should you arrange the query statements?

**Select and Place:**

The screenshot shows a query editor interface with two columns: 'Statements' and 'Answer Area'. The 'Statements' column contains five text boxes with the following content: `| join ( Events`, `| where continent == 'Asia'`, `) on RequestId`, `| where timestamp > ago(2d)`, and `Logs`. The `Logs` box is highlighted in blue. The 'Answer Area' column contains three numbered boxes (1, 2, 3) for ordering the statements.

**Correct Answer:**

The screenshot shows the correct answer for the query editor. The 'Statements' column contains three text boxes: `| join ( Events`, `| where continent == 'Asia'`, and an empty box. The 'Answer Area' column contains three numbered boxes (1, 2, 3) for ordering the statements. The correct order is: 1. `) on RequestId`, 2. `| where timestamp > ago(2d)`, and 3. `Logs`.

Section:

Explanation:

#### QUESTION 45

DRAG DROP

You have web app named App1 that uses Application Insights in Azure Monitor to Store log data. App1 has users in multiple locations. You need to query App1 requests from London and Paris that return error. The solution must meet the following requirements:

Return the timestamp, url, resultCode, and duration fields.

Only requests made the last hour.

How should you complete the query?

Select and Place:

**Values**

- extend
- project
- select
- timestamp >= ago(1hr)
- timestamp -gt ago(1hr)

**Answer Area**

```
...
requests
| where [ ]
| where resultCode == "404" and (client_City == "London" or client_City == "Paris")
| [ ] timestamp, url, resultCode, duration
...
```

Correct Answer:

**Values**

- extend
- project
- [ ]
- timestamp >= ago(1hr)
- [ ]

**Answer Area**

```
...
requests
| where timestamp -gt ago(1hr)
| where resultCode == "404" and (client_City == "London" or client_City == "Paris")
| select [ ] timestamp, url, resultCode, duration
...
```

Section:

Explanation:

#### QUESTION 46

DRAG DROP

You have an Azure subscription that uses Azure Monitor and contains a Log Analytics workspace.

You have an encryption key.

You need to configure Azure Monitor to use the key to encrypt log data

Select and Place:

**Actions**

- Grant the system-assigned managed identity Certificate permissions for the key vault.
- Create an Azure key vault and store the key.
- Configure the key vault properties for the cluster.
- Grant the system-assigned managed identity Key permissions for the key vault.
- Create an Azure Monitor Logs dedicated cluster that has a system-assigned managed identity.
- Link the Log Analytics workspace to the cluster.

**Answer Area**

- 1
- 2
- 3
- 4
- 5

**Correct Answer:**

**Actions**

- Grant the system-assigned managed identity Certificate permissions for the key vault.
- 
- 
- 
- 
- 

**Answer Area**

- 1
- 2
- 3
- 4
- 5

- Create an Azure key vault and store the key.
- Configure the key vault properties for the cluster.
- Grant the system-assigned managed identity Key permissions for the key vault.
- Create an Azure Monitor Logs dedicated cluster that has a system-assigned managed identity.
- Link the Log Analytics workspace to the cluster.

**Section:**

**Explanation:**

**QUESTION 47**

**HOTSPOT**

You use get for source control.

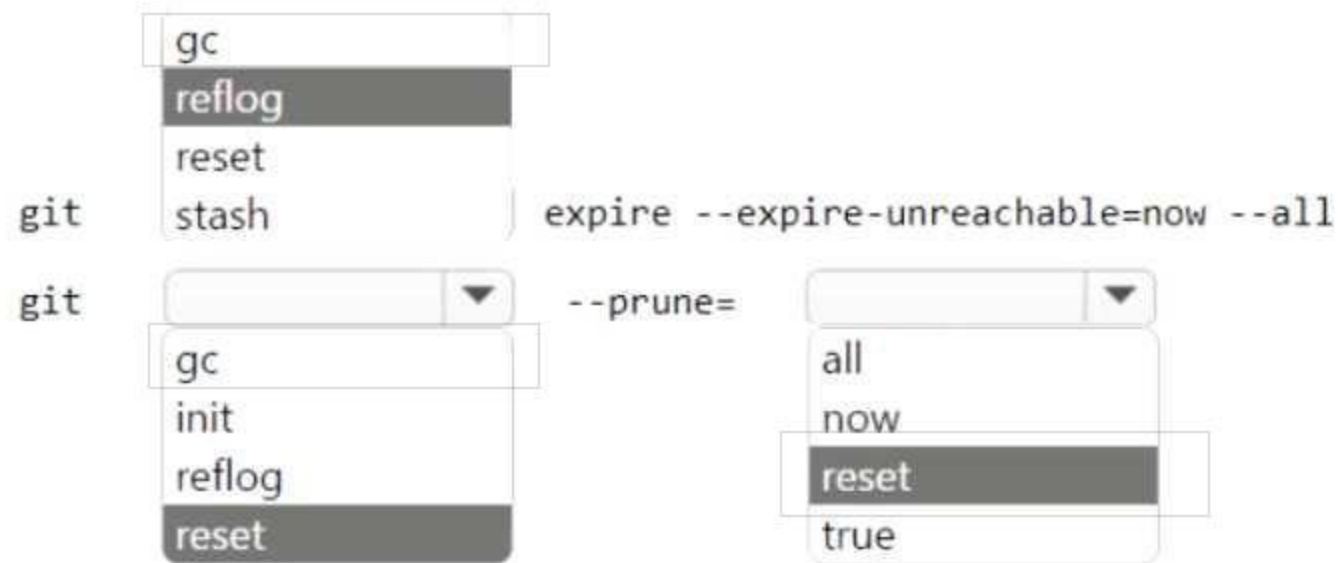
You need to optimize the performance of a repository. The solution must meet the following requirements:

Permanently remove all items referenced only in the reflog.

Remove history that is NOT in any current branch.

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

**Hot Area:**



Answer Area:



Section:

Explanation:

**QUESTION 48**

You have an Azure subscription that contains the resources shown in the following table.

Name	Type
Feed1	Azure Artifacts feed
Project1	Project in Azure DevOps

Project1 produces 9pm packages that are published to Feed 1. Feed1 is consumed by multiple projects. You need to ensure that only tested packages are available for consumption. The solution must minimize development effort.

What should you do?

- A. Create a feed view named @default. After the 9pm packages test successfully, configure a release pipeline that promotes a package to the @default view.

- B. Create a feed view named release and set @release as the default view. After the 9pm packages test successfully, configure a release pipeline that promotes a package to the @release View.
- C. Create a feed view named @release and set @release as the default view. After the 9pm packages test successfully, configure a release pipeline that tags the packages as release.
- D. Create a feed view named @default. After the 9pm packages test successfully. configure a release pipeline that tags the packages as release.

**Correct Answer: C**

**Section:**

**QUESTION 49**

Your company has an Azure DevOps project that produces Node Package Manager (npm) packages.

Multiple projects consume the packages.

You need to minimize the amount of disk space used by older packages in Azure Artifacts.

What should you modify?

- A. the retention settings of the project's pipeline
- B. the retention settings of the project's release
- C. the retention settings of the project's tests
- D. the retention settings of the company pipeline

**Correct Answer: B**

**Section:**

**Explanation:**

To minimize the amount of disk space used by older packages in Azure Artifacts, you should modify the retention settings of the project's release. This can be done by navigating to the project's release settings and adjusting the retention policy. For more information, please refer to the Microsoftdocumentation.



**QUESTION 50**

You have an Azure subscription that contains the resources shown in the following table.

Name	Type
DepPipeline1	Azure DevOps deployment pipeline
ADFPipeline1	Azure Data Factory pipeline
Vault1	Azure Key Vault

DepPipeline1 and ADFPipeline1 use a single credential that is stored in Vault1.

You need to configure ADFPipeline1 to retrieve the credential from Vault1.

Which type of activity should you use?

- A. Web
- B. Copy
- C. Lookup
- D. Get Metadata

**Correct Answer: A**

**Section:**

**QUESTION 51**

You use GitHub for source control of .NET applications.

You need to deploy a documentation solution that meets the following requirements:

Documents will be written in Markdown as developers make code changes

Changes to the documents will trigger the recompilation of a static website.

Users will access the documents from the static websites

Documents will be stored in a GitHub repository

Which two tools can you use to compile the website? Each correct answer presents a complete solution.

- A. Jekyll
- B. Medium
- C. caret
- D. WordPress
- E. DocFX

**Correct Answer: A, D**

**Section:**

**Explanation:**

#### QUESTION 52

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:

- Ensure that the credentials for the identity are managed automatically.
- Support granting privileges to the identity.

Which type of identity should you use?

- A. a service principal
- B. a user-assigned managed identity
- C. a system-assigned managed identity
- D. a user account

**Correct Answer: B**

**Section:**

**Explanation:**

#### QUESTION 53

HOTSPOT

You have a virtual machine that runs Windows Server 2019 and is managed by using Desired State Configuration (DSC). You have the following DSC configuration.





```

configuration WebConfiguration
{
    File WebsiteContent {
        Ensure          = 'Present'
        SourcePath      = 'c:\test\index.htm'
        DestinationPath = 'c:\inetpub\wwwroot'
        DependsOn      = '[WindowsFeature]Web-Server'
    }

    WindowsFeature Web-Server
    {
        Ensure = 'Present'
        Name   = 'Web-Server'
    }
}

```

You have the following Local Configuration Manager (LCM) configuration.



```

LocalConfigurationManager
{
    ConfigurationMode = "ApplyAndMonitor"
    RefreshFrequencyMins = 30
    ConfigurationModeFrequencyMins = 60
    RefreshMode = 'Push'
}

```

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:**

Statements	Yes	No
The index.htm file will be copied to the C:\Test folder before the Web-Server Windows feature is installed.	<input type="radio"/>	<input type="radio"/>
If manual changes are made to the configuration of the virtual machine, the configuration will reapply automatically.	<input type="radio"/>	<input type="radio"/>
If the Web-Server Windows feature is uninstalled from the virtual machine, the discrepancy will be reported in a log entry within 60 minutes.	<input type="radio"/>	<input type="radio"/>

**Answer Area:**

Statements	Yes	No
The index.htm file will be copied to the C:\Test folder before the Web-Server Windows feature is installed.	<input type="radio"/>	<input checked="" type="radio"/>
If manual changes are made to the configuration of the virtual machine, the configuration will reapply automatically.	<input type="radio"/>	<input checked="" type="radio"/>
If the Web-Server Windows feature is uninstalled from the virtual machine, the discrepancy will be reported in a log entry within 60 minutes.	<input checked="" type="radio"/>	<input type="radio"/>

**Section:**

**Explanation:**

No "DependsOn" are defined"

No "Mode is Apply and Monitor"

Yes "ConfigurationModeFreqMin : 60 "

-----ConfigurationModeFreqMin

:How often, in minutes, the current configuration is checked and

applied. This property is ignored if the ConfigurationMode property is set to ApplyOnly.

<https://docs.microsoft.com/en-us/powershell/dsc/managing-nodes/metaconfig?view=dsc-1.1>

<https://docs.microsoft.com/en-us/powershell/dsc/configurations/resource-depends-on?view=dsc1.1>

**QUESTION 54**

**DRAG DROP**

You have an Azure subscription that uses Azure Automation State Configuration to manage the configuration of virtual machines. You need to identify which nodes are noncompliant.

How should you complete the query? To answer, drag the appropriate values to the correct targets.

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:**



Values	Answer Area
Category	AzureDiagnostics
DscReportStatus	where <input type="text"/> == "DscNodeStatus"
Message	where <input type="text"/> contains ""
OperationName	where <input type="text"/> != "Compliant"
Resource	
ResultType	

Correct Answer:

Values	Answer Area
<input type="text"/>	AzureDiagnostics
DscReportStatus	where <input type="text" value="Category"/> == "DscNodeStatus"
Message	where <input type="text" value="OperationName"/> contains ""
<input type="text"/>	where <input type="text" value="ResultType"/> != "Compliant"
Resource	
<input type="text"/>	

Section:

Explanation:

**QUESTION 55**

**HOTSPOT**

You have a project in Azure DevOps that contains a release pipeline. The pipeline contains two stages named QA and Prod. QA deploys code to an Azure web app named webapp1. Prod deploys code to an Azure web app named webapp2. You need to ensure that code deployments to webapp2 are blocked if Azure Application Insights generates Failed requests alerts following the deployment of new code to webapp1. What should you do for each stage? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

QA:

- Add a task to configure alert rules in Application Insights.
- Configure a gate in the pre-deployment conditions.
- Configure an auto-redeploy trigger in the post-deployment conditions
- Configure a post-deployment approval in the post-deployment conditions

Prod:

- Add a task to configure an alert rule in Application Insights.
- Configure a gate in the pre-deployment conditions.
- Configure a trigger in the pre-deployment conditions.
- Configure the Deployment queue settings in the pre-deployment conditions.

Answer Area:

QA:

- Add a task to configure alert rules in Application Insights.
- Configure a gate in the pre-deployment conditions.
- Configure an auto-redeploy trigger in the post-deployment conditions
- Configure a post-deployment approval in the post-deployment conditions

Prod:

- Add a task to configure an alert rule in Application Insights.
- Configure a gate in the pre-deployment conditions.
- Configure a trigger in the pre-deployment conditions.
- Configure the Deployment queue settings in the pre-deployment conditions.

Section:

Explanation:

**QUESTION 56**

DRAG DROP

You use Exabeam Fusion SIEM and the Azure cloud platform.

You need to integrate Exabeam and Azure. The solution must use OAuth authentication.

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:

**Actions**

- Upload a certificate.
- Create a client secret.
- Register an Exabeam application in Microsoft Azure Active Directory (Azure AD), part of Microsoft Entra.
- Configure the Exabeam Azure cloud connector.
- Configure API permissions.

**Answer Area**

1

2

3

Correct Answer:

**Actions**

- Upload a certificate.
- 
- 
- 
- Configure API permissions.

**Answer Area**

1 Create a client secret.

2 Register an Exabeam application in Microsoft Azure Active Directory (Azure AD), part of Microsoft Entra.

3 Configure the Exabeam Azure cloud connector.

Section:

Explanation:

**QUESTION 57**

DRAG DROP

You have an Azure Repos repository named repo1.

You delete a branch named features/feature11.

You need to recover the deleted branch.

Which three commands should you 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:

**Commands**

```
git restore <SHA1>
git stash
git log
git checkout <SHA1>
git branch features/feature11
```

**Answer Area**

>

<

**Correct Answer:**

**Commands**

```
git restore <SHA1>
git stash
```

**Answer Area**

```
git log
git checkout <SHA1>
git branch features/feature11
```

>

<

**Section:**

**Explanation:**

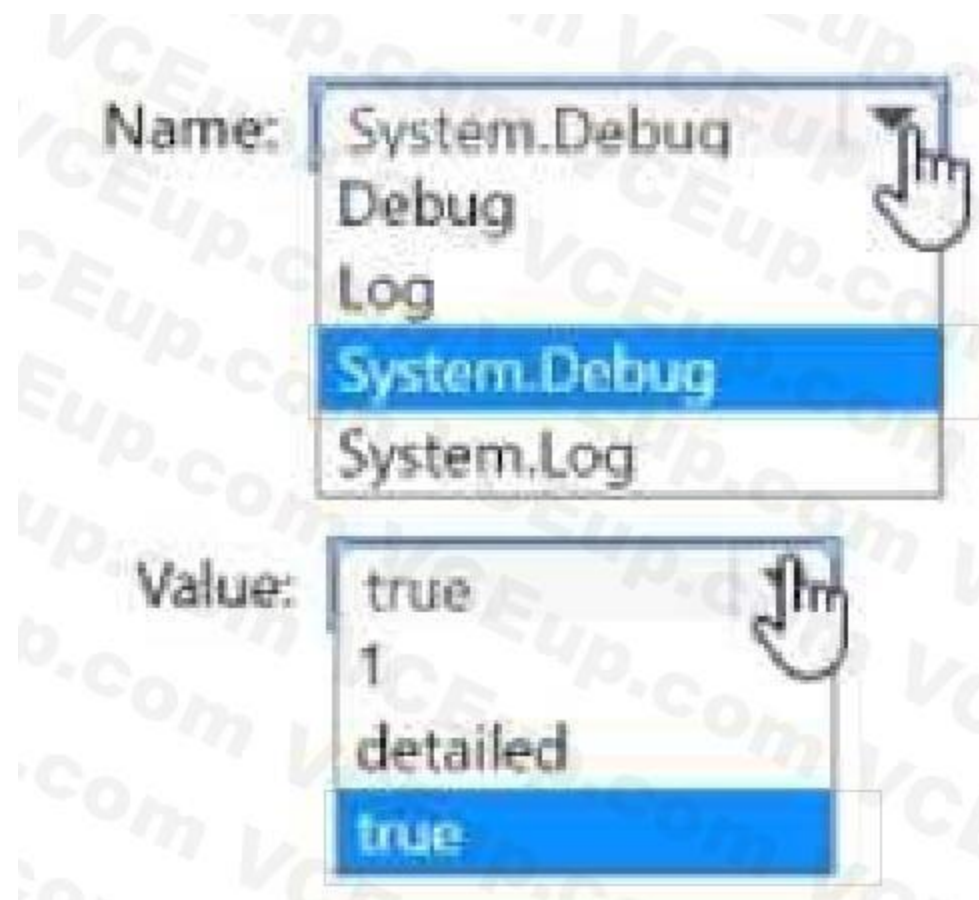
**QUESTION 58**

**HOTSPOT**

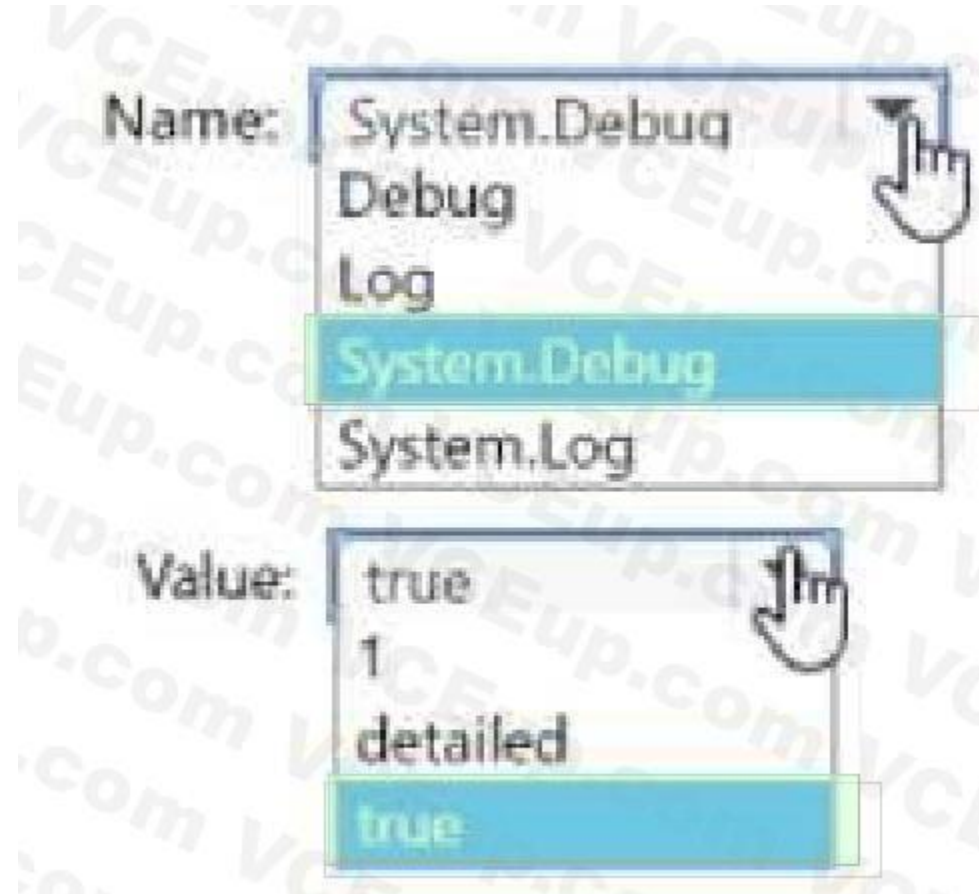
You have a project in Azure DevOps that contains a Continuous Integration/Continuous Deployment (CI/CD) pipeline. You need to enable detailed logging by defining a pipeline variable. How should you configure the variable? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Hot Area:**





Answer Area:



Section:

Explanation:



**QUESTION 59**

DRAG DROP

You are developing a full Microsoft .NET Framework solution that includes unit tests.

You need to configure SonarQube to perform a code quality validation of the C# code as part of the build pipelines. Which four tasks should you perform in sequence? To answer, move the appropriate tasks from the list of tasks to the answer area and arrange them in the correct order.

Select and Place:

Actions Commands Cmdlets Statements	Answer Area
Run Code Analysis	
Visual Studio Test	
Publish Build Artifacts	
Visual Studio Build	
Prepare Analysis Configuration	

Correct Answer:

Actions Commands Cmdlets Statements	Answer Area
	Prepare Analysis Configuration
	Visual Studio Build
Publish Build Artifacts	Visual Studio Test
	Run Code Analysis



**Section:**

**Explanation:**

Step 1: Prepare Analysis Configuration

Prepare Analysis Configuration task, to configure all the required settings before executing the build. This task is mandatory.

In case of .NET solutions or Java projects, it helps to integrate seamlessly with MSBuild, Maven and Gradle tasks.

Step 2: Visual Studio Build

Reorder the tasks to respect the following order:  
Prepare Analysis Configuration task before any MSBuild or Visual Studio Build task.

Step 3: Visual Studio Test

Reorder the tasks to respect the following order:

Run Code Analysis task after the Visual Studio Test task.



#### Step 4: Run Code Analysis

Run Code Analysis task, to actually execute the analysis of the source code.

This task is not required for Maven or Gradle projects, because scanner will be run as part of the Maven/Gradle build.

Note:



#### QUESTION 60

You have an Azure Automation account that contains a runbook. The runbook is used to configure the application infrastructure of an Azure subscription. You have a project in Azure DevOps named Project1. Project1 contains a repository that stores code for the runbook. You need to ensure that every committed change to the code will update automatically and publish the runbook to Azure Automation. What should you configure?

- A. the Connections settings for the Automation account
- B. the Service hooks settings for Project1
- C. the Source control settings for the Automation account
- D. the Service connections settings for Project1

**Correct Answer: C**

**Section:**

#### QUESTION 61

You manage source code control and versioning by using GitHub.

You need to ensure that a PowerShell script is executed automatically before rebase operations are performed. What should you use?

- A. a package
- B. GitHub Copilot
- C. a webhook
- D. a gist

**Correct Answer: C**

**Section:**

#### QUESTION 62

You manage a project by using Azure Boards. You manage the project code by using GitHub. You have three work items that have IDs of 456, 457, and 458. You need to create a pull request that will be linked to all the work items. The solution must set the state of work item 456 to done. What should you add to the commit message?

- A. Done #456, #457, #A53
- B. Fixes #456, #457, #458
- C. Fixes #AB456, #AB457, #AB4S8
- D. #AB456, #AB457, #AB458 Completed #AB456

**Correct Answer: C**

**Section:**

**QUESTION 63**

HOTSPOT

You have an Azure subscription.

You need to create a storage account by using a Bicep file.

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

NOTE: Each correct selection is worth one point.

**Hot Area:**

```
param storageAccount string
var storageAccountNameToUse = '${storageAccount}${uniqueString(resourceGroup().id)}'
resource invoiceStorage 'Microsoft.Storage/storageAccounts@2022-03-01' = [
  name: storageAccountNameToUse
  location: 'eastus'
  sku: {
    name: 'Standard_GRS'
  }
  kind: StorageV2
  properties: {
    publicAccess: 'blob'
    allowBlobPublicAccess: true
  }
  kind: StorageV2
  param:
  properties:
  type:
  var:
```

**Answer Area:**

Answer Area

```
param storageAccount string
var storageAccountNameToUse = '${storageAccount}${uniqueString(resourceGroup().id)}'

resource invoiceStorage 'Microsoft.Storage/storageAccounts@2022-05-01' = {
  name: storageAccountNameToUse
  location: 'eastus'
  sku: {
    name: 'Standard_GRS'
  }
  kind: 'StorageV2'
  properties: {
    blobCaching: 'None'
    defaultServiceVersion: '2022-05-01'
    fileEncryptionKey: ''
    fileEncryptionKeySource: 'Microsoft.KeyVault/Vaults@2021-07-01'
    fileEncryptionScope: ''
    fileEncryptionScopeKeyVaultName: ''
    fileEncryptionScopeKeyName: ''
    httpOnly: true
  }
}
```



Section:

Explanation:

**QUESTION 64**

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 need to recommend an integration strategy for the build process of a Java application. The solution must meet the following requirements:

- The builds must access an on-premises dependency management system.
- The build outputs must be stored as Server artifacts in Azure DevOps.
- The source code must be stored in a get repository in Azure DevOps.

Solution: Configure the build pipeline to use a Microsoft-hosted agent pool running the Windows Server 2022 with Visual Studio 2022 image, include the Java Tool installer task in the build pipeline.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

Section:

Explanation:

**QUESTION 65**

DRAG DROP

You have an Azure Repos Git repository named repo1.

You need to ensure that you can authenticate to repo1 by using SSH.

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:

**Actions**

- Add the SSH private key.
- Save the SSH key to the root of repo1.
- Sign in to Azure DevOps.
- Create SSH keys by using ssh-keygen.
- Add the SSH public key.
- Clone repo1.

**Answer Area**



**Correct Answer:**

**Actions**

- Add the SSH private key.
- Save the SSH key to the root of repo1.

**Answer Area**

- Sign in to Azure DevOps.
- Create SSH keys by using ssh-keygen.
- Add the SSH public key.
- Clone repo1.



**Section:**

**Explanation:**

**QUESTION 66**

DRAG DROP

You are implementing a new project in Azure DevOps.

You need to assess the performance of the protect. The solution must identify the following metrics:

- How long it takes to complete a work item
- \* The percentage of defects found in production

Which DevOps KPI should you review for each metric? To answer drag the appropriate KPIs to the correct metric. Each KPI may be used once, more than once, or not at all. You may need to drag the spirit bar between panes or scroll to view content.

NOTE Each correct selection is worth one point.

**Select and Place:**

- KPIs**
- Application failure rates
  - Bug report rates
  - Burndown trend
  - Cycle time
  - Defect escape rate
  - Deployment speed
  - Lead time
  - Mean time to recover



**Answer Area**

How long it takes to complete a work item:

The percentage of defects found in production:

**Correct Answer:**

- KPIs**
- Application failure rates
  - 
  - Burndown trend
  - 
  - Defect escape rate
  - Deployment speed
  - Lead time
  - Mean time to recover



**Answer Area**

How long it takes to complete a work item:

The percentage of defects found in production:



**Section:**

**Explanation:**

**QUESTION 67**

DRAG DROP

You have an Azure subscription that contains 50 virtual machines.

You manage the configuration of the virtual machines by using Azure Automation State Configuration.

You need to ensure that Windows Defender is installed on each virtual machine and the Windows Defender service is running.

How should you complete the Desired State Configuration (DSC) code? To answer, drag the appropriate values to the correct targets. 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:**

Values

- 'Automatic'
- 'Enabled'
- 'IncludeAllSubFeature'
- 'Manual'
- 'Present'
- 'Running'
- 'Started'

Answer Area

```

Configuration EnsureSecure {
  Node "localhost" {
    WindowsFeature WindowsDefenderInstalled {
      Ensure = 
      Name = 'Windows-Defender'
    }
    Service WindowsDefenderAvailable {
      Name = 'WinDefend'
      StartupType = 'Automatic'
      State = 
    }
  }
}

```

Correct Answer:

Values

- 'Automatic'
- 'Enabled'
- 'IncludeAllSubFeature'
- 'Manual'
- 
- 
- 'Started'

Answer Area

```

Configuration EnsureSecure {
  Node "localhost" {
    WindowsFeature WindowsDefenderInstalled {
      Ensure = 'Present'
      Name = 'Windows-Defender'
    }
    Service WindowsDefenderAvailable {
      Name = 'WinDefend'
      StartupType = 'Automatic'
      State = 'Running'
    }
  }
}

```

Section:

Explanation:

QUESTION 68

DRAG DROP

You have an Azure Repos repository named Repo1 that is used for source control.

You need to configure code scanning for Repo1.

Which three tasks should the pipeline perform in sequence? To answer, move the appropriate tasks from the list of tasks to the answer area and arrange them in the correct order

**Select and Place:**

Tasks

- Microsoft Purview Scan
- Advanced Security AutoBuild
- Advanced Security Initialize CodeQL
- Advanced Security Perform CodeQL Analysis
- Advanced Security Dependency Scanning

Answer Area

**Correct Answer:**

Tasks

- Microsoft Purview Scan
- 
- 
- 
- Advanced Security Dependency Scanning

Answer Area

- Advanced Security Initialize CodeQL
- Advanced Security AutoBuild
- Advanced Security Perform CodeQL Analysis



**Section:**

**Explanation:**

**QUESTION 69**

DRAG DROP

You have an Azure subscription.

You have the standards shown in the following table.

Name	Contents
Standard1	Security requirements
Standard2	Software prerequisites

You plan to use Azure Pipelines to build and release web apps.

You need to recommend a solution to build the pipelines. The solution must meet the following requirements:

- \* Ensure that all new pipelines meet the security requirements defined in Standards
- \* Ensure that the first stage of all new pipelines contains the software prerequisites defined in Standard2.
- \* Minimize administrative effort.

What should you use to implement each standard? To answer, drag the appropriate options to the correct standards. 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:**

Options

- A pipeline permission
- A self-hosted agent
- A stage
- An extends template
- An includes template

Answer Area

Standard1:

Standard2:

Correct Answer:

Options

- A pipeline permission
- A self-hosted agent
- A stage
- 
- 

Answer Area

Standard1:

Standard2:

Section:

Explanation:

QUESTION 70

HOTSPOT

You have an Azure Pipelines pipeline named Pipeline1 that has the following YAML definition.





```
pool:
  name: Azure Pipelines
  demands:
  - msbuild
  - visualstudio

steps:
- task: VSBuild@1
  displayName: 'Build solution **\*.sln'
  inputs:
    solution: '$(Parameters.Solution)'
    platform: '$(BuildPlatform)'
    configuration: '$(BuildConfiguration)'

- task: CopyFiles@2
  displayName: 'Copy Files to: $(Build.ArtifactStagingDirectory)'
  inputs:
    SourceFolder: '$(System.DefaultWorkingDirectory)'
    Contents: '**\bin\$(BuildConfiguration)\**'
    TargetFolder: '$(Build.ArtifactStagingDirectory)'
  condition: succeededOrFailed()

- task: PublishBuildArtifacts@1
  displayName: 'Publish Artifact: drop'
  inputs:
    PathToPublish: '$(Build.ArtifactStagingDirectory)'
  condition: succeededOrFailed()
```



For each of the following statements, select Yes if True. Otherwise select No.

NOTE: Each correct selection is worth one point.

**Hot Area:**

Answer Area

Statements	Yes	No
Pipeline1 consists of three stages.	<input type="radio"/>	<input type="radio"/>
\$(BuildPlatform) is a predefined variable.	<input type="radio"/>	<input type="radio"/>
\$(System.DefaultWorkingDirectory) is a predefined variable.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Pipeline1 consists of three stages.	<input checked="" type="radio"/>	<input type="radio"/>
\$(BuildPlatform) is a predefined variable.	<input type="radio"/>	<input checked="" type="radio"/>
\$(System.DefaultWorkingDirectory) is a predefined variable.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

QUESTION 71

DRAG DROP

You have a GitHub repository named repo1.

You migrate repo1 to an Azure Repos repository named repo2.

After the migration, changes are made to repo1.

You need to sync the changes to repo2.

How should you complete the script? To answer, drag the appropriate values to the correct targets. 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:

Values

- repo1
- repo1 clone URL
- repo1 repo URL
- repo2
- repo2 clone URL
- repo2 repo URL

Answer Area

```
...
git clone --bare [ ]
cd [ ]
git remote add --mirror=fetch upstream [ ]
git fetch upstream --tags
git push origin --all
...
```

Correct Answer:

Values

Answer Area

```
...
git clone --bare repo1 repo URL
cd repo1
git remote add --mirror=fetch upstream repo1 clone URL
git fetch upstream --tags
git push origin --all
...
```

**Section:**

**Explanation:**

#### QUESTION 72

You manage a project by using Azure Boards, and you manage the project code by using Azure Repos. You have a bug work item that has an ID of 123. You need to set the work item state to Resolved. What should you add to the commit message?

- A. Fixes #123
- B. Resolves #AB-123
- C. #123 completes
- D. #123 Resolved

**Correct Answer: A**

**Section:**

#### QUESTION 73

You manage projects by using Azure Boards. You manage project code by using GitHub. You have a work item that has an ID of 456. You need to link work item 456 to a new pull request. What are two ways to achieve this goal? Each correct answer presents a complete solution. NOTE: Each correct solution is worth one point.

- A. To the description of the pull request, add #AB456.
- B. To work item 456, add a comment that includes the URL of the pull request.
- C. In the Development section for work item 456, select Add link, and then enter the URL of the pull request
- D. From work item 456, open the Links tab, select Add link, select Existing item and then enter the URL of the commit.

**Correct Answer: A, B**

**Section:**

#### QUESTION 74

You need to use an Azure Pipelines pipeline to build and test an app and test the database of the app. The solution must meet the following requirements.

- \* The test stages must be run in parallel.
- \* The Publish\_Test\_Results stage must always be run.
- \* The test stages must be run after successful completion of the build stage.
- \* The Publish\_Test\_Results stage must be run after completion of all the test stages

Solution: You include the following elements in the YAML definition of the pipeline.



```
...
stages:
- stage: Build_App
jobs:
- stage: Test_App
dependsOn: [Build_App]
jobs:
- stage: Test_Database
dependsOn: [Build_App]
jobs:
- stage: Publish_Test_Results
jobs:
...
```

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

#### QUESTION 75

##### SIMULATION

###### Task 1

Navigate to <https://dev.azure.com>, select Start Free, and specify the following credentials:

\* UserUsefl-42147509@ExamUsers.com

\* Password: eWrSalD2!

Use the default setting to sign up for Azure DevOps and create an Azure DevOps organization. Once the organization is created, creates private project named Project1.

You need to add an external user that has an email address of Usfrr2-42147509@ExamUsers.com as a stakeholder of the User1 -42147509 Azure DevOps organization.

The user must be added to the most restrictive Azure DevOps group.

To complete this task, sign in to the Azure DevOps portal as Userl-42147509ExamUsers.com.

- A. See the solution below in explanation

**Correct Answer: A**

**Section:**

##### **Explanation:**

Step 1: Sign Up for Azure DevOps

Navigate to Azure DevOps.

Click onStart Free.

Enter the credentials:

Email: UserUsefl-42147509@ExamUsers.com

Password: eWrSalD2!

Follow the prompts to complete the sign-up process using the default settings.

Step 2: Create an Azure DevOps Organization

Once signed in, you will be prompted to create a new organization.

Enter a name for your organization and select your region.

Click onContinueto create the organization.

Step 3: Create a Private Project

In your new organization, click onNew Project.



Name the project Project1.  
Set the visibility to Private.  
Click on Create.  
Step 4: Add an External User as a Stakeholder  
Go to the Organization Settings.  
Under General, select Users.  
Click on Add users.  
Enter the email address: Usfrr2-42147S09@ExamUsers.com.  
Set the access level to Stakeholder.  
Add the user to the most restrictive group, which is typically the Readersgroup.  
Click on Add to complete the process.  
Step 5: Verify the User Addition  
Ensure that the external user has been added successfully by checking the Userslist.  
Confirm that the user has the Stakeholder access level and is part of the Readersgroup.  
By following these steps, you should be able to complete the task successfully. If you encounter any issues, feel free to ask for further assistance!

#### QUESTION 76

##### SIMULATION

##### Task 2

For the RG1 lod42147S09 resource group, you need to create an action group named DevOpsAG that can be triggered in any region by using Alerts.  
The action group must email two users named admin1 @contoso.com and admin2@contoso.com and notify members of the Owner role for the resource.  
Use only the common alert schema for notifying the resource owners.

A. See the solution below in explanation

**Correct Answer: A**

**Section:**

**Explanation:**

Step 1: Create an Action Group  
Navigate to Azure Portal:  
Go to Azure Portal and sign in with your credentials.  
Access Azure Monitor:  
In the left-hand menu, select Monitor.  
Create Action Group:  
Under Alerts, select Action groups.  
Click on + Create.  
Configure Basic Settings:  
Subscription: Select your subscription.  
Resource Group: Select RG1 lod42147S09.  
Action Group Name: Enter DevOpsAG.  
Display Name: Enter a display name for the action group.  
Step 2: Define Actions  
Add Email Notifications:  
Click on Add action.  
Action Type: Select Email/SMS message/Push/Voice.  
Action Name: Enter a name for the action (e.g., EmailAdmins).  
Email: Enter admin1@contoso.com and admin2@contoso.com.  
Click OK.  
Notify Resource Owners:  
Click on Add action again.



Action Type: SelectEmail/SMS message/Push/Voice.

Action Name: Enter a name for the action (e.g.,NotifyOwners).

Email: SelectNotify all owners.

ClickOK.

Step 3: Enable Common Alert Schema

Common Alert Schema:

In theAdvancedtab, enable theCommon alert schemaoption1.

Step 4: Review and Create

Review:

Review all the settings you have configured.

Create:

Click onReview + createand thenCreate.

By following these steps, you will have successfully created an action group namedDevOpsAGthat emails the specified users and notifies resource owners using the common alert schema

## QUESTION 77

### SIMULATION

#### Task 3

You need to create a new team dashboard named Dashboard1 for the default project team of Project1. The dashboard must display the members of the team

A. See the solution below in explanation

**Correct Answer: A**

**Section:**

**Explanation:**

Step 1: Create a New Team Dashboard

Navigate to Azure DevOps:

Go toAzure DevOpsand sign in with your credentials.

Select Your Project:

ChooseProject1from your list of projects.

Access Dashboards:

In the left-hand menu, selectDashboards.

Create a New Dashboard:

Click onNew Dashboard.

Enter the nameDashboard1.

Ensure the dashboard type is set toTeam Dashboard.

ClickCreate.

Step 2: Add the Team Members Widget

Open the Widget Catalog:

After creating the dashboard, the widget catalog will open automatically. If it doesn't, click onAdd Widget.

Search for Team Members Widget:

In the widget catalog, search forTeam Members.

Add the Widget:

Click on theTeam Memberswidget and then clickAddto place it on your dashboard.

Configure the Widget:

Once added, you can resize and move the widget to your preferred location on the dashboard.

Step 3: Save and Share the Dashboard

Save the Dashboard:

Click onSaveto save your changes.

Share the Dashboard:

You can share the dashboard URL with your team members or set permissions to control who can view or edit the dashboard.



By following these steps, you will have a new team dashboard named Dashboard1 that displays the members of the default project team for Project1

