ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

HOTSPOT

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 a worth one point.

```
MSAppCenter.start ( "{Your App Secret}", withServices:
 '                    [MS An allies, self,                    MS Analytics, sell]
                      [MSDistribute self, [MSPush.self,       MSCrashes sel^ MS
                                                              Distribute self]
```

Answer:

Explanation:

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

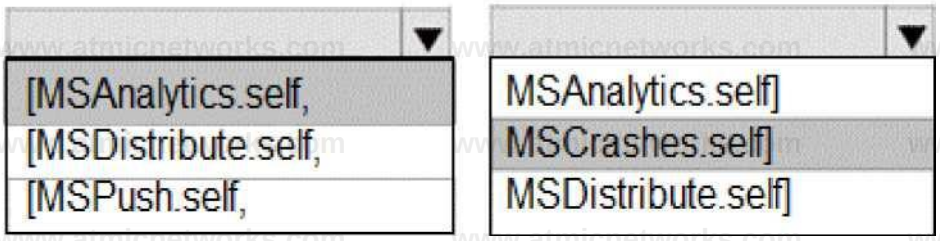

Scenario: Visual Studio App Center must be used to centralize the reporting of mobile application crashes and device types in use.

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 App Delete class in the didFinishLaunchingWithOptions method.

MSAppCenter.start("{Your App Secret}", withServices: [MSAnalytics.self, MSCrashes.self])

Reference:https://docs.microsoft.com/en-us/appcenter/sdk/getting-started/ios

Question: 2

HOTSPOT

How should you confrere the release retention policy for the investment planning depletions suite?

To answer, select the appropriate options in the answer area.

NOTE: Eachcorrect selection is worth one point.

Required secrets:

- Certificate
- Personal access token
- Shared Access Authorization token
- Username and password

Storage location:

- Azure Data Lake
- Azure Key Vault
- Azure Storage with H1 IH access
- Azure Storage with HI I PS access

Answer:

Explanation:

Required secrets:

- Certificate
- Personal access token
- Shared Access Authorizatiion token
- Username and password

Storage location:

- Azure Data Lake
- Azure Key Vault
- Azure Storage with H1 IP access
- Azure Storage with HI I PS access

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.

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

Question: 3

HOTSPOT

You need to configure a cloud service to store the secrets required by the mobile applications to call the share.

What should you include in thesolution? To answer, select the appropriate options m the answer area, NOTE: Each correct selection is worth one point.

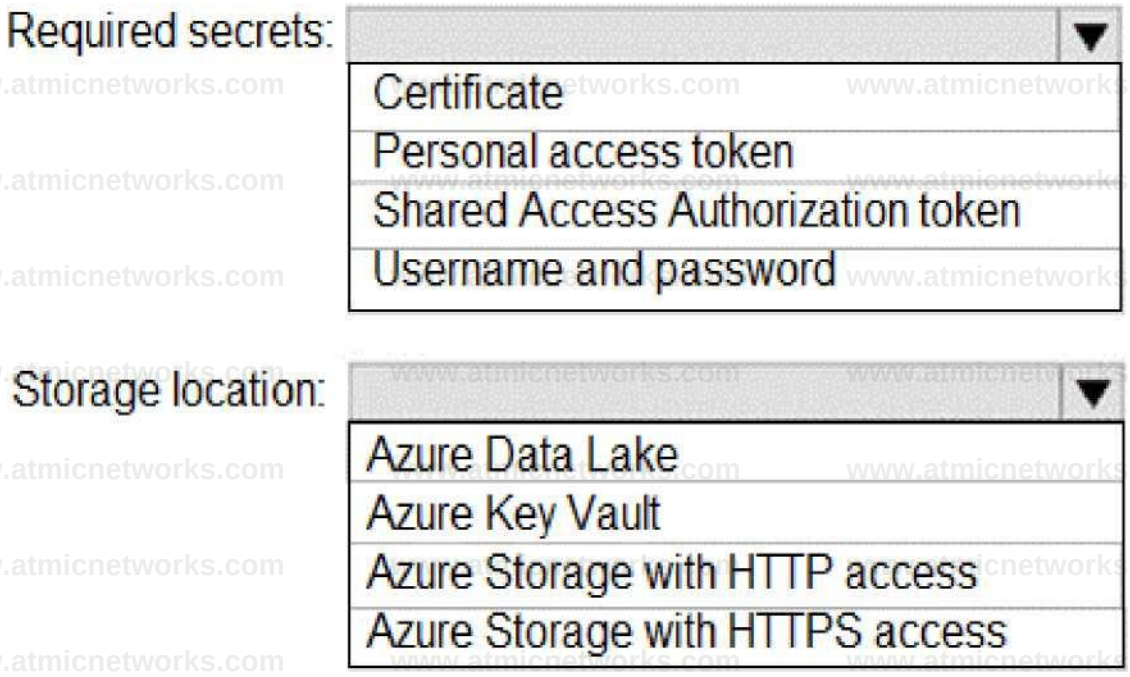

Answer:

Explanation:

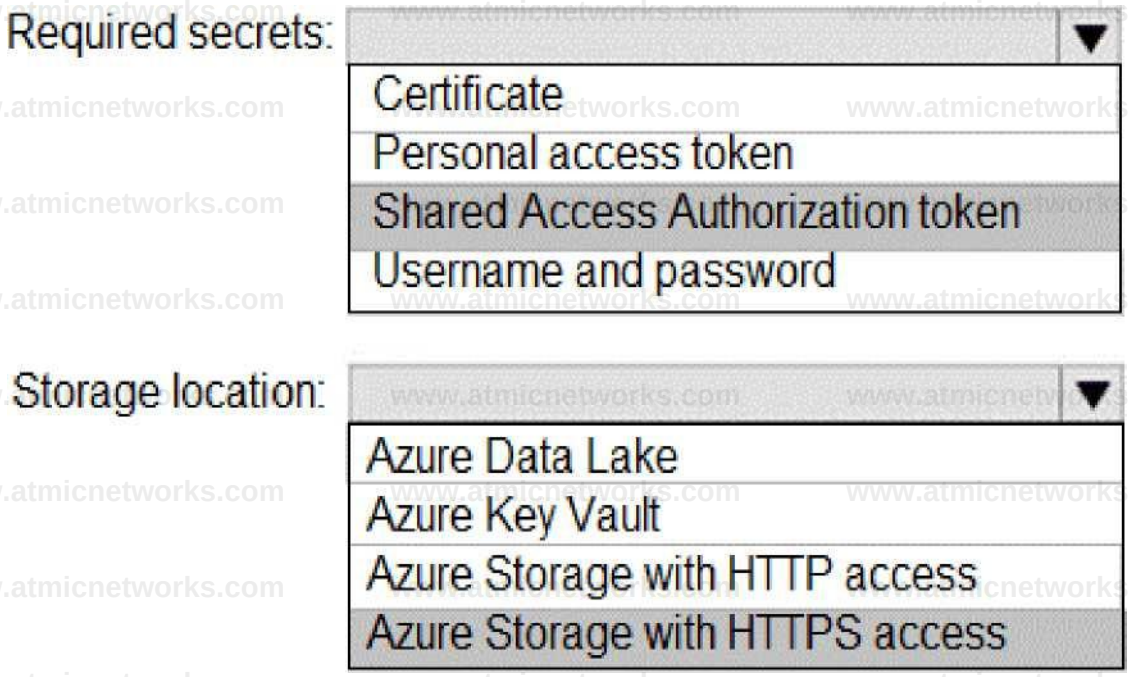

Every request made against a storage service must be authorized, unless the request is for a blob or container resource that has been made available for public or signed access. One option for authorizing a request is by using Shared Key.

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.

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

Question: 4

To resolve the current technical issue, what should you do to theRegister-AzureRmAutomationDscNodecommand?

A. Change the value of theConfigurationModeparameter.
B. Replace theRegister-AzureRmAutomationDscNodecmdlet withRegister-AzureRmAutomationScheduledRunbook
C. Add theAllowModuleOverwriteparameter.
D. Add theDefaultProfileparameter.

Answer: A

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 AutomationState Configuration fails to correct the configurations.

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

```
Re gi s te r- A z u reRmAu tom t i 0 nD s cN cd e
   -Res our ce G roupName 1Te s tResour ce G roupf
   - Aut om ti o nA co oun tName 'L it war eAu tom t i o nAo
count'
```

```
-AzureVMName ^vmanrae
-ConfigurationMode 'ApplyOnly'
```

Reference:https://docs.microsoft.com/en-us/powershell/module/azurerm.automation/register-azurermautomationdscnode?view=azurermps-6.13.0

Question: 5

What should you use to implementthe code quality restriction on the release pipeline for the investment planning applications suite?

A. a trigger
B. a pre deployment approval
C. a post-deployment approval
D. a deployment gate

Answer: D

Explanation:

Question: 6

HOTSPOT

How should youconfigure 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.

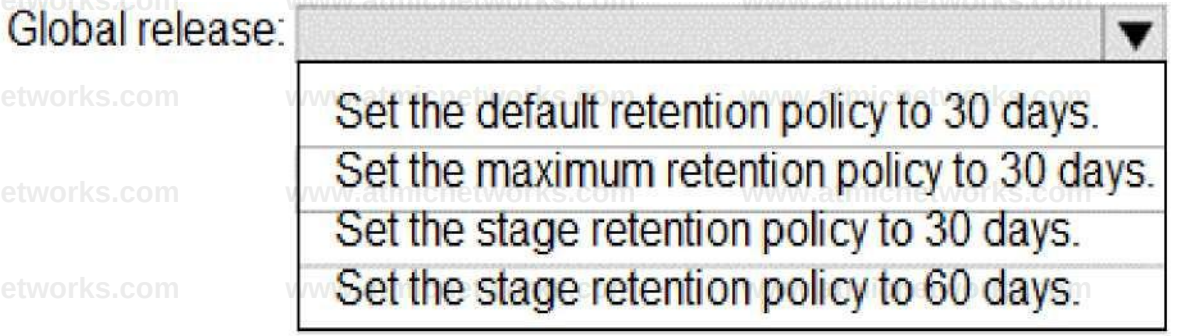

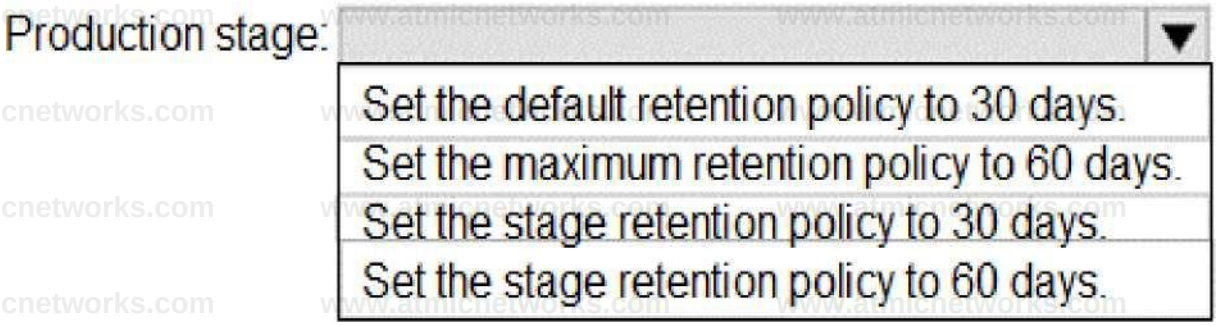

Answer:

Explanation:

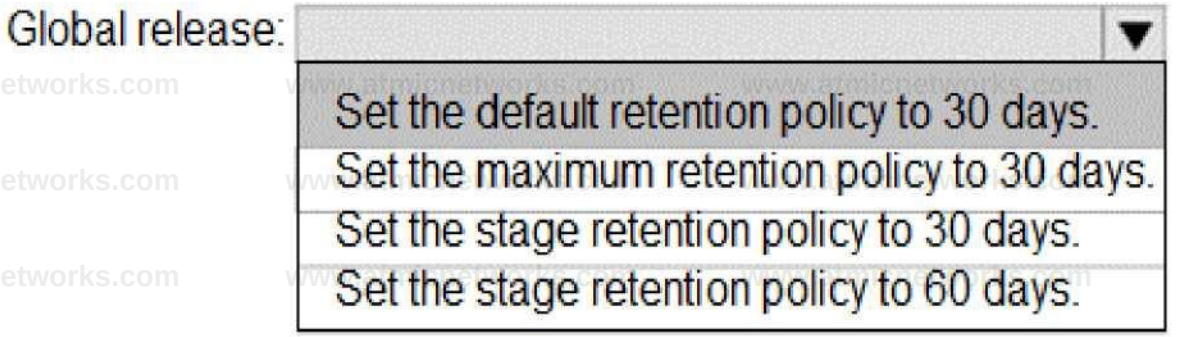

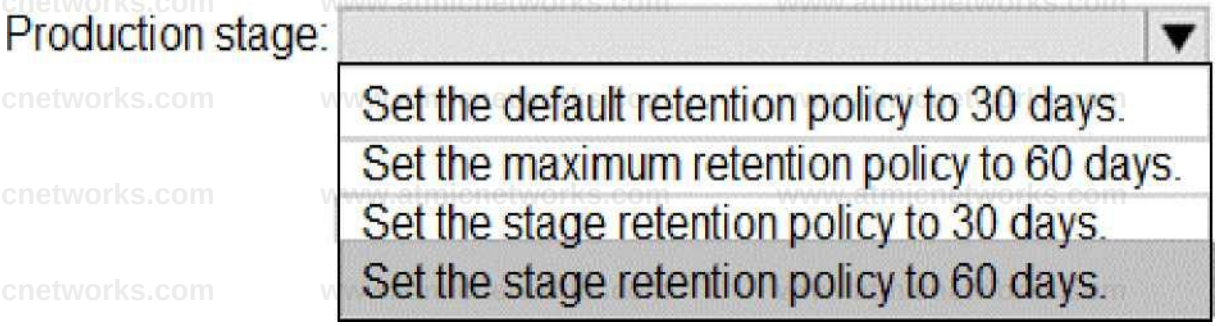

Scenario: By default, all releases must remain available for 30 days, except for production releases, which must be kept for 60 days.

Box 1: Set the defaultretention 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.

Reference:https://docs.microsoft.com/en-us/azure/devops/pipelines/policies/retention

Question: 7

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.

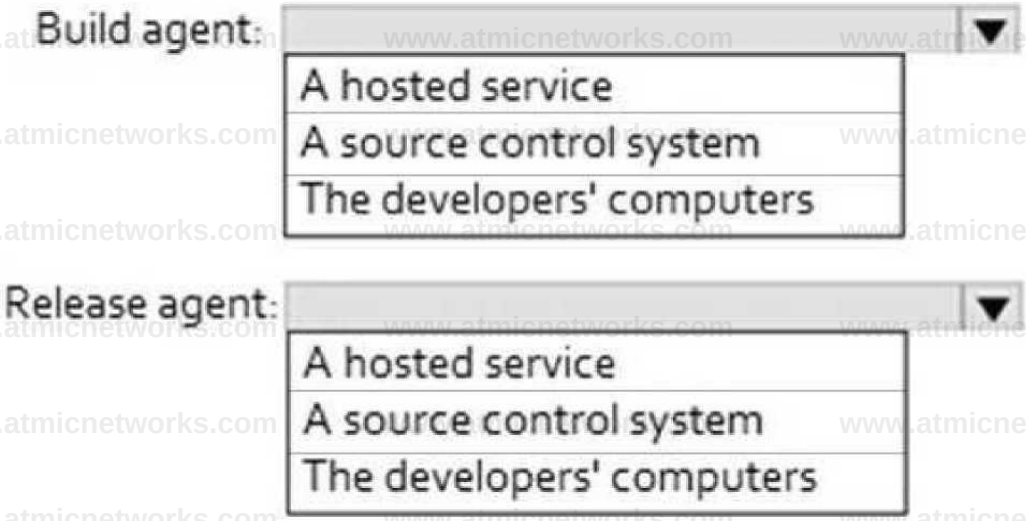

Answer:Explanation:

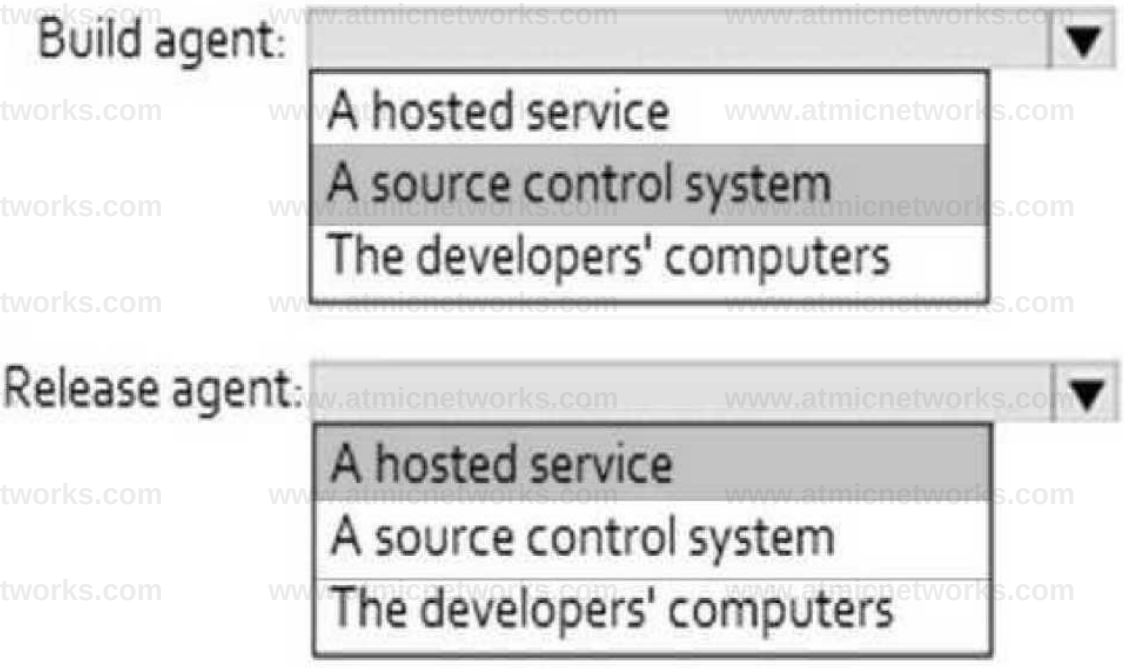

Box 1: A source control system
A source control system, also called a version control system, allows developers to collaborate on code and track changes.Source control is an essential tool for multi-developer projects.

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 agentmeets 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.

Reference:
https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/v2-osx?view=azure-devops

Question: 8
Which branching strategy should you recommend for the investment planning applications suite?

A. release isolation
B. main only
C. development isolation
D. feature isolation

Answer: D

Explanation:

Scenario: A branching strategy that supports developing newfunctionality 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.

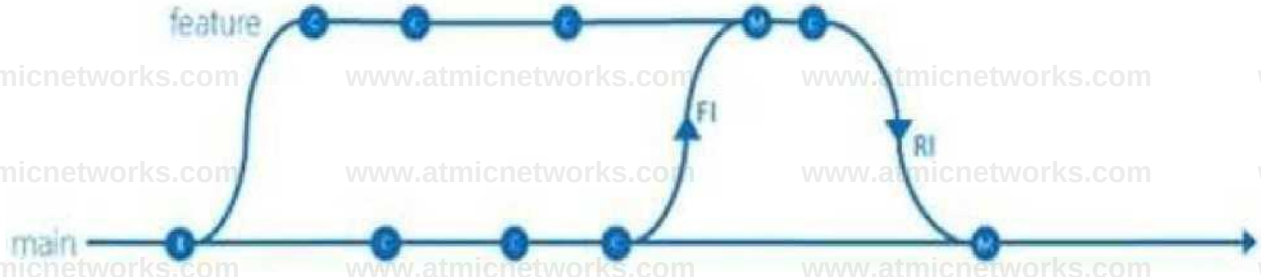

When you need to work on a particularfeature, it might be a good idea to create a feature branch.

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.

Reference:

https://docs.microsoft.com/en-us/azure/devops/repos/tfvc/branching-strategies-with- tfvc?view=azure-devops

Question: 9

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.

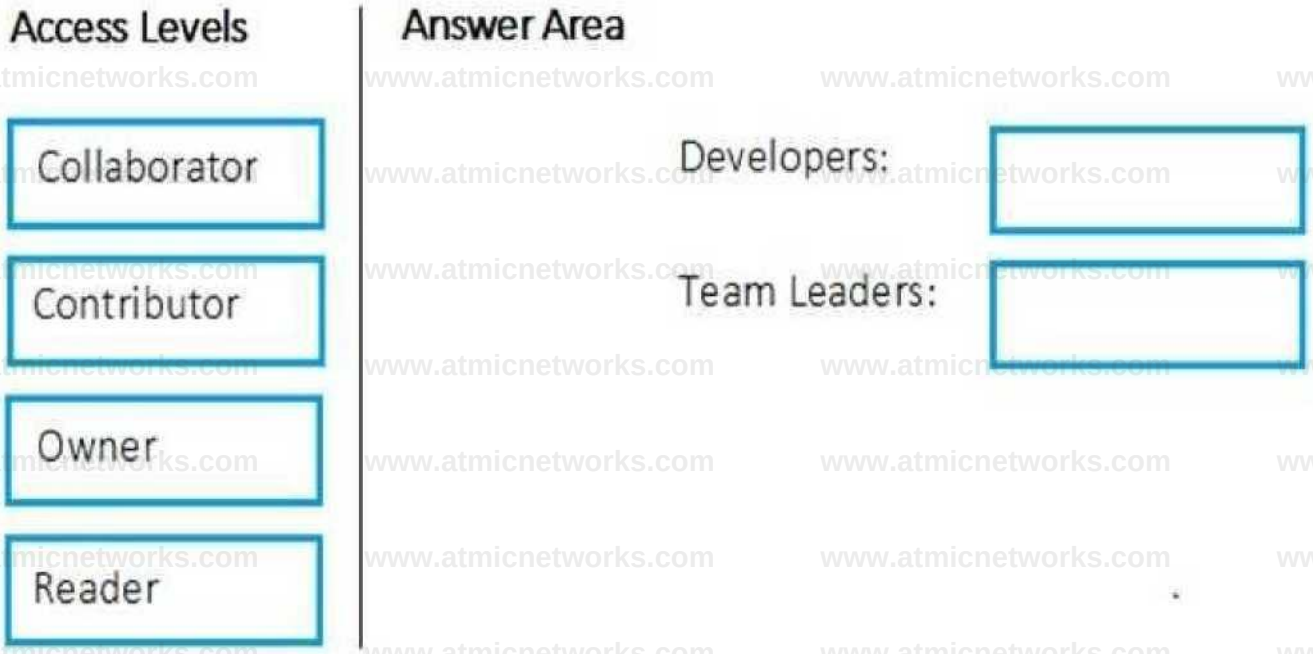

Answer:

Explanation:

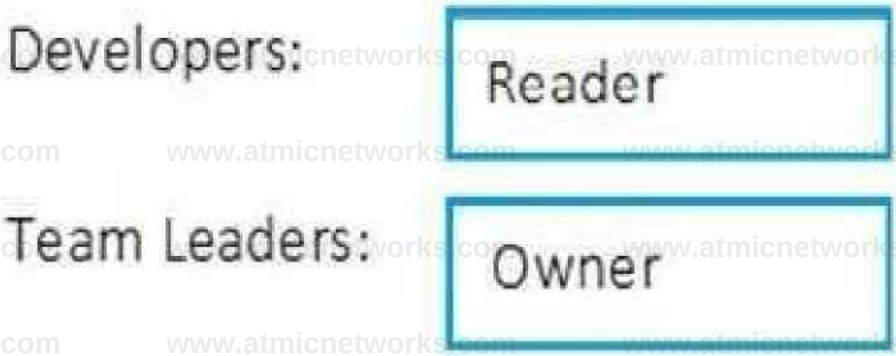

Box 1: Reader
Members of a group named Developers must be able to install packages.

Feeds have four levels of access: Owners, Contributors, Collaborators, and Readers. Ownerscan 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	7	7	7	y
Save packages from upstream sources		y	7	7
Push packages			/	y
Unlist/deprecate packages			V	V
Delete/unpublish package				7
Edit feed permissions				7
Rename and delete feed				y

Question: 10

You are using GitHub asa source code repository.

You create a client-side Git 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 git commit parameter shouldyou use?

A. --squash
B. --no-verify
C. --message ''
D. --no-post-rewrite

Answer: B

Explanation:

The commit-msg hook is invoked by git-commit and git-merge, and can be bypassed with the --noverify option.

Reference:
https://git-scm.com/docs/githooks

Question: 11

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

Answer: B

Explanation:

Batching CI runs

If you have many team members uploadingchanges 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 therepository. 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: 12

HOTSPOT

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 optionsin the answer area.

NOTE: Each correct selection is worth one point.

Microsoft. OperationalInsights/workspaces - ExpandProperties

Get-AzResource	-ResourceGroupName
Get-AzResourceGroup	-ResourceId
Get-AzResourceProvider	-ResourceType

Answer:

Explanation:

Microsoft.OperationalInsights/workspaces - ExpandProperties

Get-AzResource	-ResourceGroupName
Get-AzResourceGroup	-ResourceId
Get-AzResourceProvider	-ResourceType

Box 1: Get-AzResource

Use the followingcommand to examine the access control mode for all workspaces in the subscription:

PowerShell

```
Get-Az Resource –Resource Type Microsoft. Operational Insights/workspaces –Expand Properties | for each {s.Name + ": " + $_.Properties.features.enableLogAccessUsingOnlyResourcePermissions
```

Box 2: -Resource Type

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/logs/manage-access

Question: 13

HOTSPOT

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 onepoint.

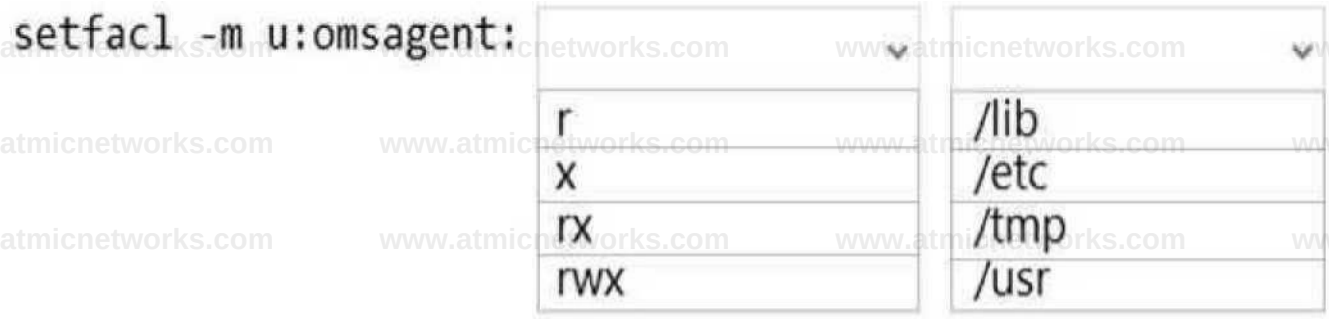

Answer:

Explanation:

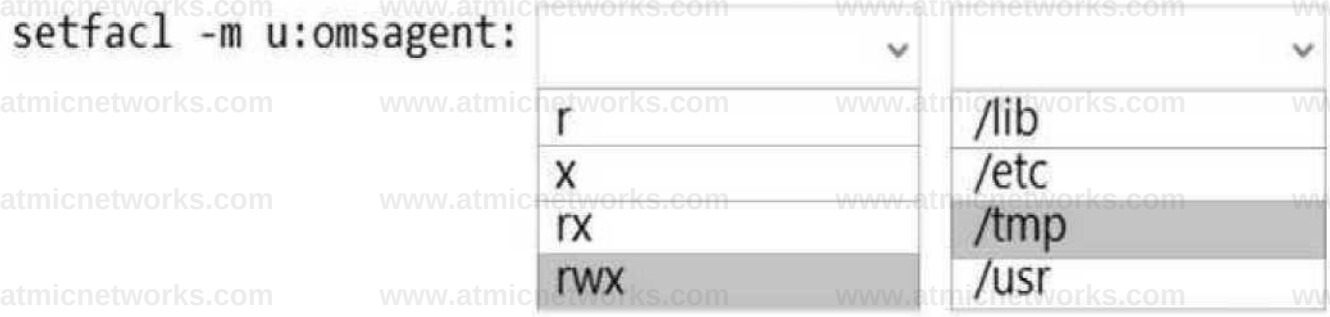

Box 1: rwx

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: 14

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 onepoint.

A. authorization credentials
B. the project name
C. the deployment group name
D. the organization URL
E. the agent pool name

Answer: C,E

Explanation:

Unattended config:

The agent can be set up from a script with no human intervention. You must pass--unattended and the answers to all questions.

To configure an agent, it must know the URL to your organization or collection and credentials of someone authorized to set up agents. All other responses are optional.

Reference:

https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/v2-windows

Topic 2, Contoso Case Study: 2

Overview

Existing Environment

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

Requirements

Contoso plans to improve itsIT development and operations processes implementing Azue DevOps principles. Contoso has an Azure subscription and creates an Azure DevelPs organization.

The Azure DevOps organization includes:

The Docker extension

A deployment pool named Pool7 that contains 10 Azure virtual machines that run Windows Server 2016.

The Azure subscription contains an Azure Automation account.

Planned Changes

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

Project name	Project details
Project 1	Projectl 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 TheTeam2 members will not have permissions to Project2.
Project 3	Projects 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	Projects 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/folderl of the repository
Project 6	Projects will provide support for build and deployment pipelines. Deployment will be allowed only if the number of current work items representing active software bugsis 0.
Project 7	Project? will contain a target deployment group named Group? that maps to Pool?. Project? will use Azure Automation State Configuration to maintain the desired state of the computers in Group?

Technical Requirements

Contosoidentities the following technical requirements:

- Implement build agents rot Project 1.
- Whenever possible, use Azure resources
- Avoid using deprecated technologies
- Implement a code flow strategy for Project2 that will:
- Enable Team 2 to submit pull requests for Project2.
- Enable Team 2 to work independently on changes to a copy of Project?
- Ensure that any intermediary changes performed by Tram2 on a copy of Project2 will be subject to the same restrictions as the ones defied in the build policy of Project2.
- Whenever possible. Implement automation and minimize administrative effort.
- Implement Protect3, Project5, Project6, and Project7 based on the planned changes.
- Implement Project4 and configure the project to push Docker images to Azure Container Reentry.

Question: 15

DRAG DROP

You need to configure Azure Automation for the computer 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.

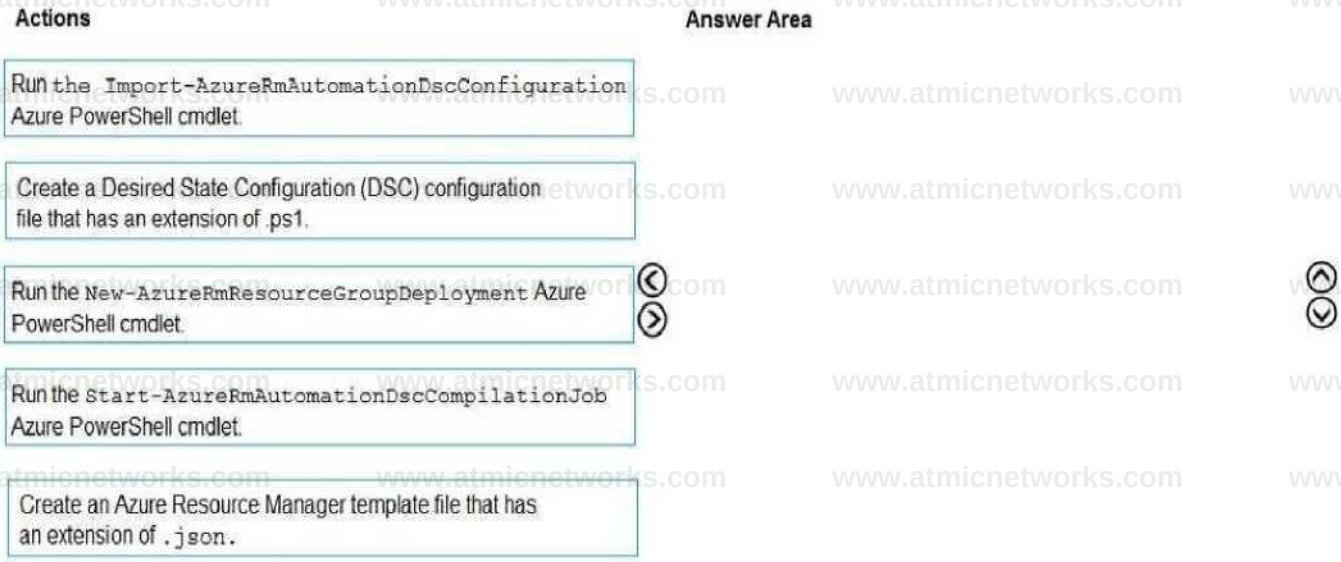

Answer:

Explanation:

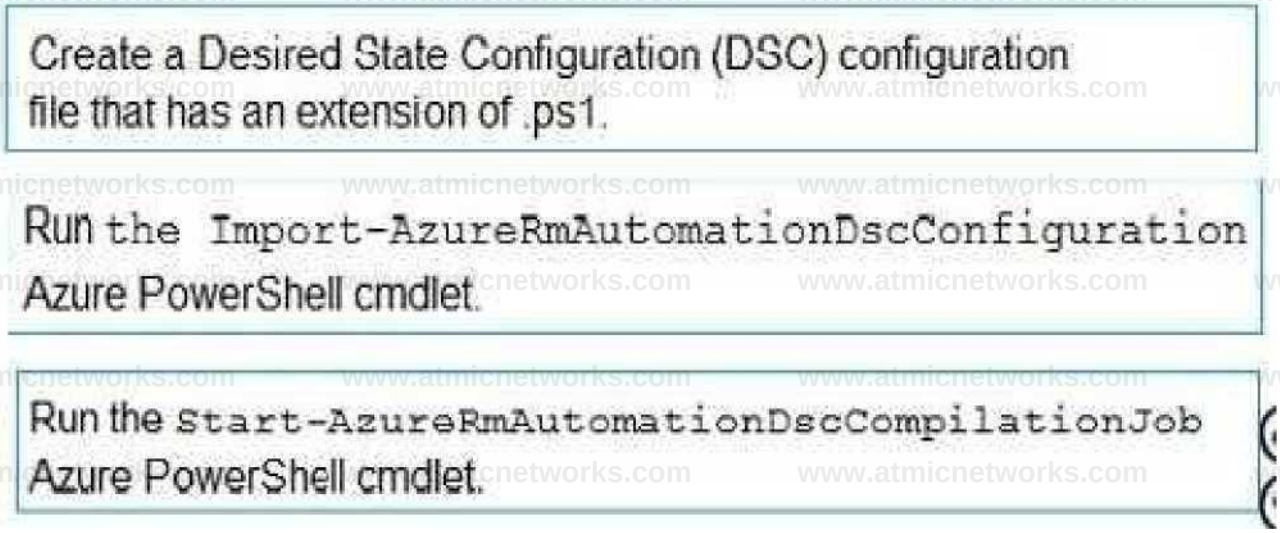

Step 1: Create a Desired State Configuration (DSC) configuration file that has an extension of .ps1.

Step 2: Run the Import-AzureRmAutomationDscConfiguration Azure Powershell cmdlet
TheImport-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.

Reference:
https://docs.microsoft.com/en-us/powershell/module/azurerm.automation/import-azurermautomationdscconfiguration

https://docs.microsoft.com/en-us/powershell/module/azurerm.automation/start-azurermautomationdsccompilationjob

Question: 16

You add the virtual machines as managed nodes in Azure Automation State Configuration.
You need to configure the computer in Group7.
What should you do?

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

Answer: A

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 contains10 Azure virtual machines that run Windows Server 2016

Project 7	Project? will contain a target deployment group named Group? that maps to Pool? Project? will use Azure Automation State Configuration to maintain the desired state of the computers in Group?.

Reference:https://docs.microsoft.com/en-us/powershell/module/azurerm.automation/register-azurermautomationdscnode

Question: 17

In Azure DevOps, you create Project3.
You need to meet the requirements of the project.
What should you do first?

A. From Azure DevOps, create a service endpoint.
B. From Sonar Qube, obtain an authentication token.
C. From Azure DevOps, modify the build definition.
D. From Sonar Qube , create a project.

Answer: A

Explanation:

The first thing to do is to declare your Sonar Qube server as a service endpoint in your VSTS/DevOps project settings.

Reference:https://docs.sonarqube.org/display/SCAN/Analyzing+with+SonarQube+Extension+for+vst s-TFS

Question: 18

HOTSPOT

How should you configure the filters for the Project5 trigger? To answer, select the appropriate option in the answer area.

NOTE: Each correct selectionis worth one point.

$^{et\ a}$ ▼ /folded

branch filter to exclude
branch filter to include path fitter to exclude
path fitter to include

Seta $_{T}$.

branch filter to exclude
branch filter to include path filter to exclude path filter to include

Answer:

Explanation:

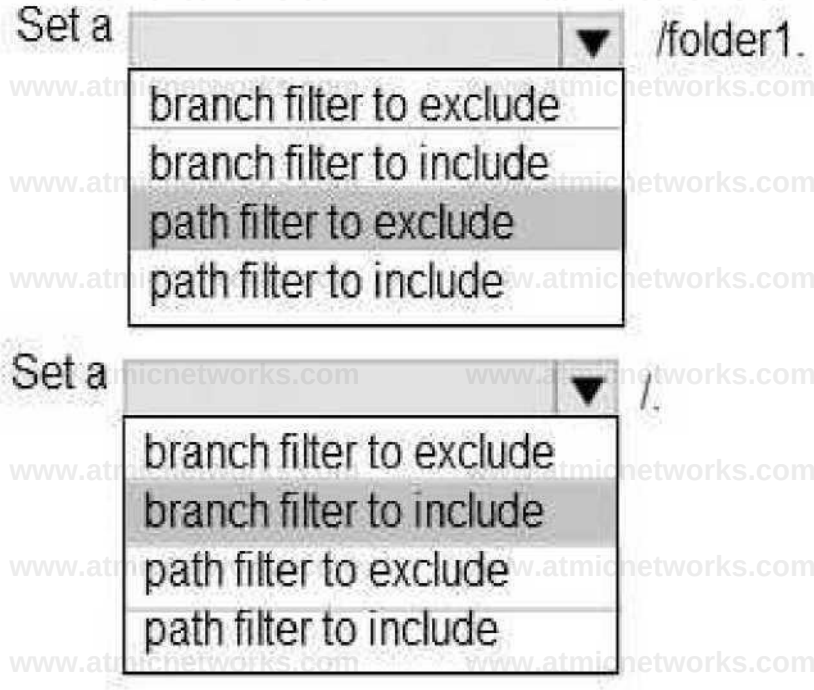

@

Scenario:

Projects will contain a Git repository in Azure Reports and a continuous integration tngger that will initiate a build in response to any change except for changes within/folderl of the repository.

Reference:https://docs.microsoft.com/en-us/azure/devops/pipelines/build/triggers

Question: 19

DRAG DROP

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 in the correct order.

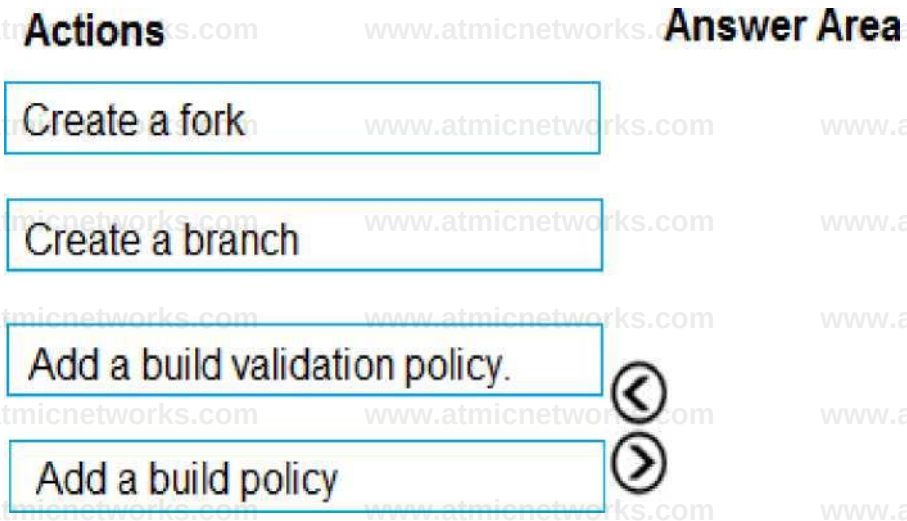

Create a repository

Add an application access policy.

Answer:

Explanation:

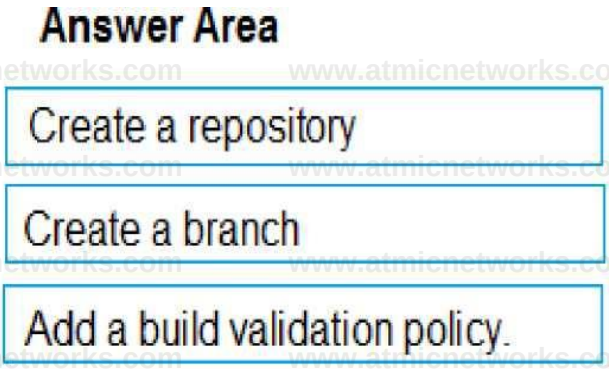

Step 1: Createa repository

A Git 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 branch

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

Step 3: Add a build validation policy

When a build validation policy is enabled, a new build is queued when a new pull request is created or when changes are pushed to an existing pull request targeting this branch. The build policy then evaluates the results of the build to determine whether the pull request can be completed.

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.

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

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

Question: 20

DRAG DROP

You need to configure Azure Automation for the computers in Pool7.

Whichthree actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them m the correct order.

Actions

Run the Rex-
Ar a *r* t tat * tow r c Hr o-upOtp 1 oywnt Azure PowerShell cmdlet

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

Run the import -
A:ur«*mAut omationDtcConfiMration Azure PowerShell cmdlet

Run the start ■>
A:ureRmAutoaation01cCoapilatlonJob Azure PowerShell cmdlet

Create a Desired Stale Configuration (DSQ configuration file that has an extension of psi.

Answer Area

1
2

Question: 21

DRAG DROP

You need to implement Project6.

Answer:

Explanation:

Actions

Run the ne*-
Azure RmRM our c Hr oupDep 1 oymant Azure PowerShell cmdlet

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

Answer Area

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

2 Run the import - ArureRmAutomat ionDscConf igurat ion Azure PowerShell cmdlet.

3 Run the start- AxureftwAutomitlonDicComp1lationJob Azure PowerShell cmdlet.

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 m the correct order.

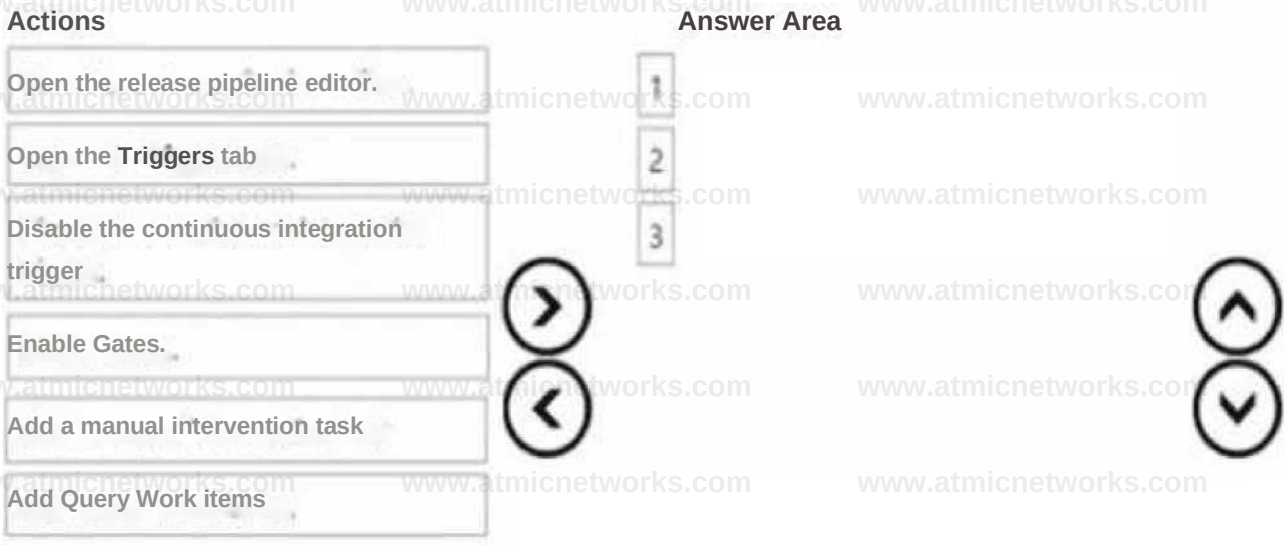

Answer:

Explanation:

Open the release pipeline editor.

Enable Gates.

Add Query Work Items.

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 thepre-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.

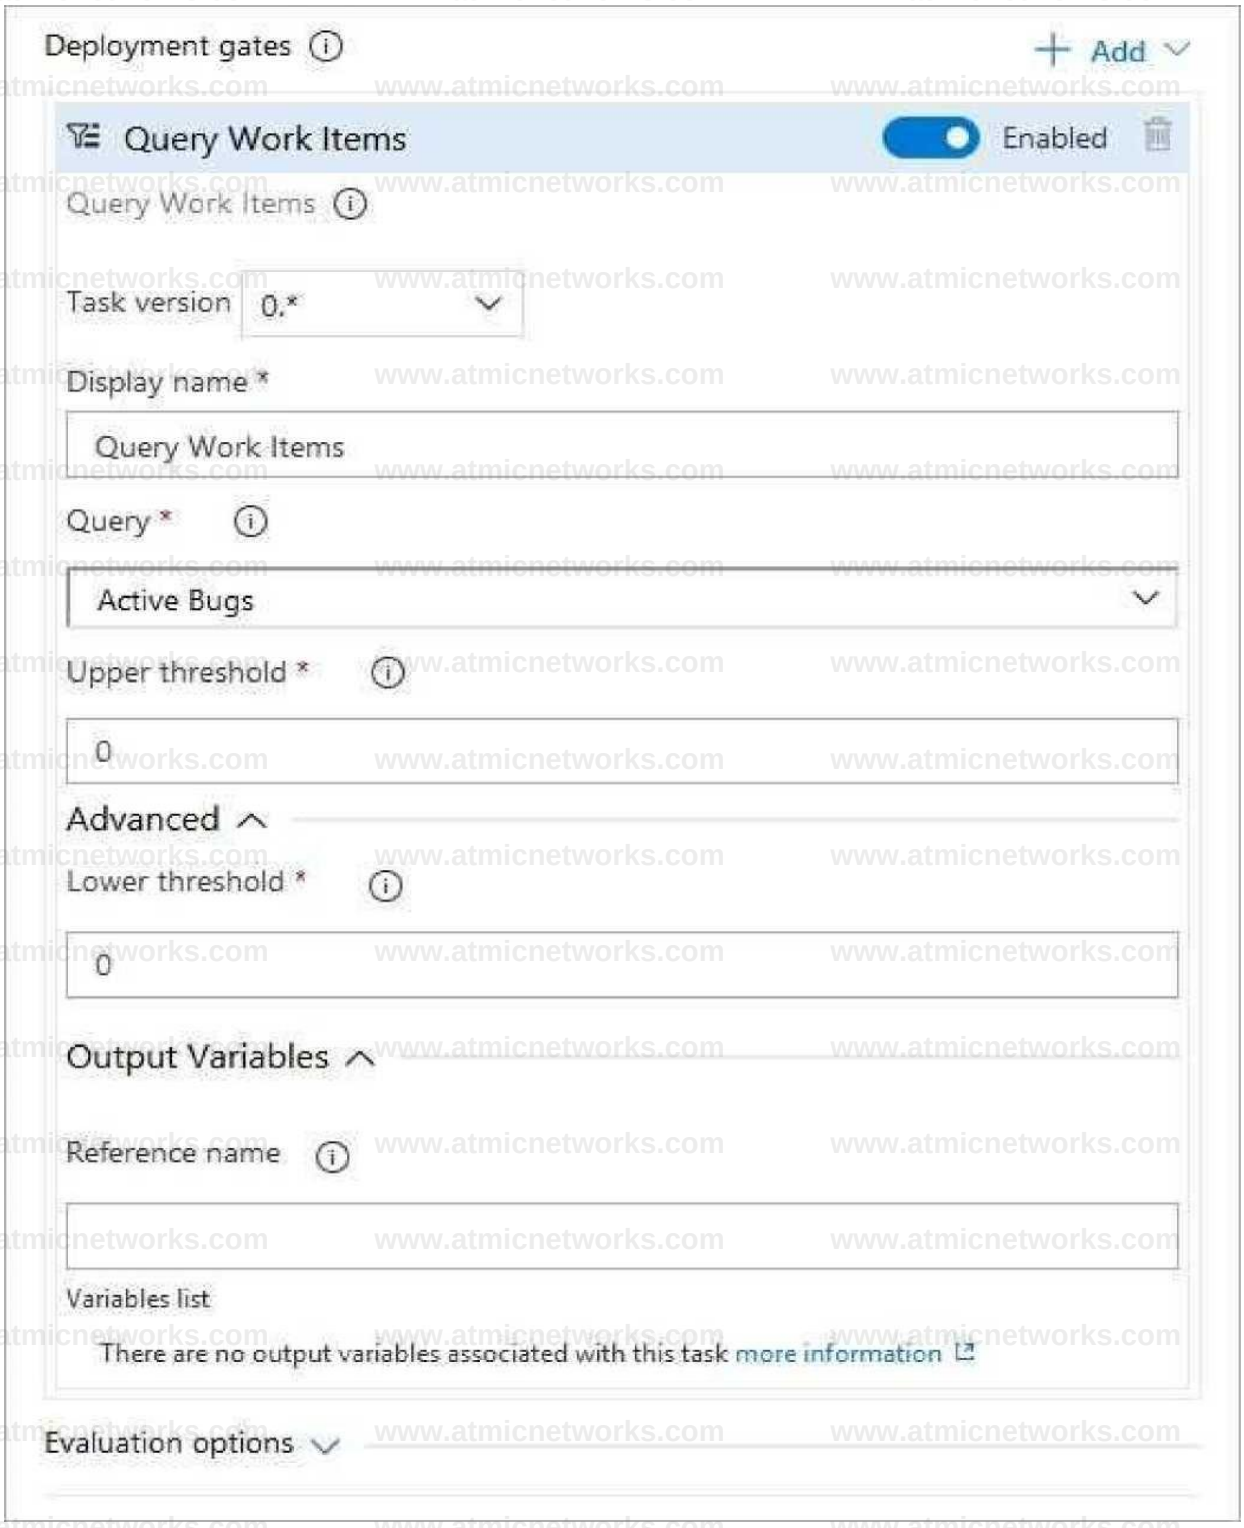

Note: A case for release gate is:
Incident and issues management. Ensure the required status for work items, incidents, and issues. For example, ensure deployment occurs only if no priority zero bugs exist, and validation that there are no active incidents takes place after deployment.

Reference:
https://docs.microsoft.com/en-us/azure/devops/pipelines/release/deploy-using-approvals?view=azure-devops#configure-gate

Question: 22

DRAG DROP

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 ofactions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create a repository	
Add a build policy for the fork.	
Create a branch.	
Add a build policy for the master branch.	
Add an application access policy.	
Create a fork.	

Answer:

Explanation:

Answer Area

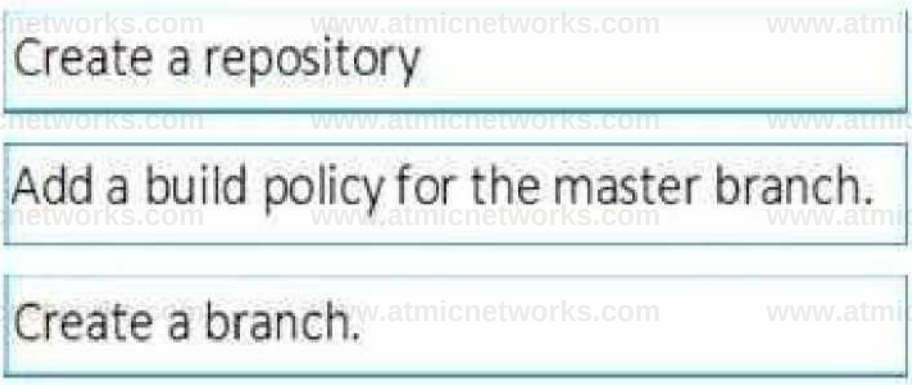

Question: 23

You need to implement Project4.

What should you do first?

A. Add the FROM instruction in the Dockerfile file.
B. Adda 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.

Answer: C

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 ofload 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: 24

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 tothe answer area and arrange them in the correct order.

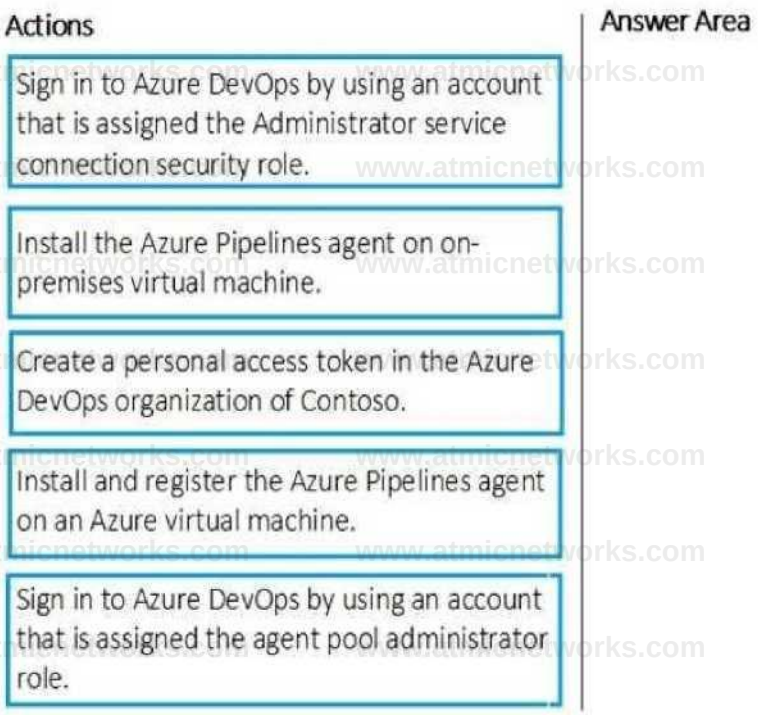

Answer:

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 DevOpsorganization of Contoso.

Explanation:

Install and register the Azure Pipelines agent on an Azure virtual machine.

Scenario:

Protect 1 Project 1 will provide support for incremental builds and third party SDK components

Step 1: Sign in to Azure Develops 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 thecluster, we make it possible to test any service, regardless of type.

Reference:

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/

Topic 3, Woodgrove bank

Overview

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 yourtime 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 containexhibits 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 theNextbutton. 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 anAll Informationtab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click theQuestionbutton 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 DevOpsfor 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 Azure DevOps dashboard must display the metrics shown inthe 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: 25

HOTSPOT

You need to configure the alert for VM1.The solution must meet the technical requirements.

Which two settings should you configure? To answer, select the appropriate settings in the answer area.

NOTE:Each correct selection is worth one point.

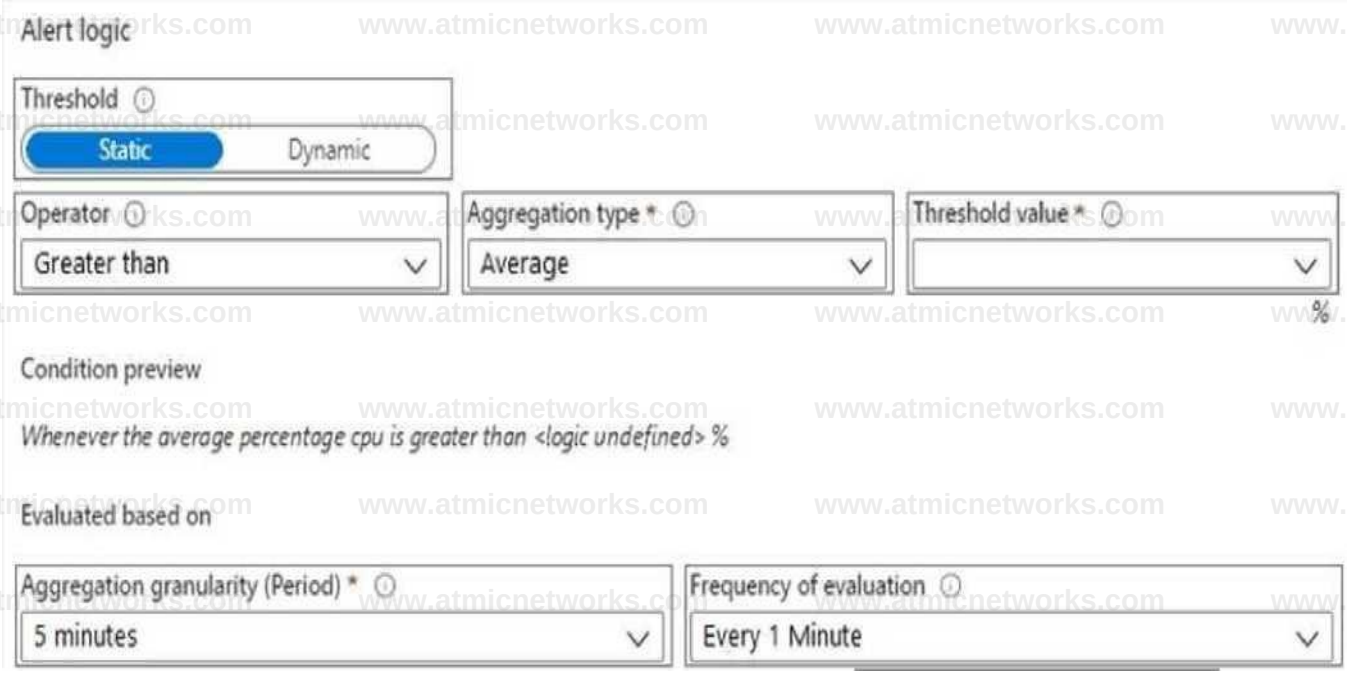

Answer:

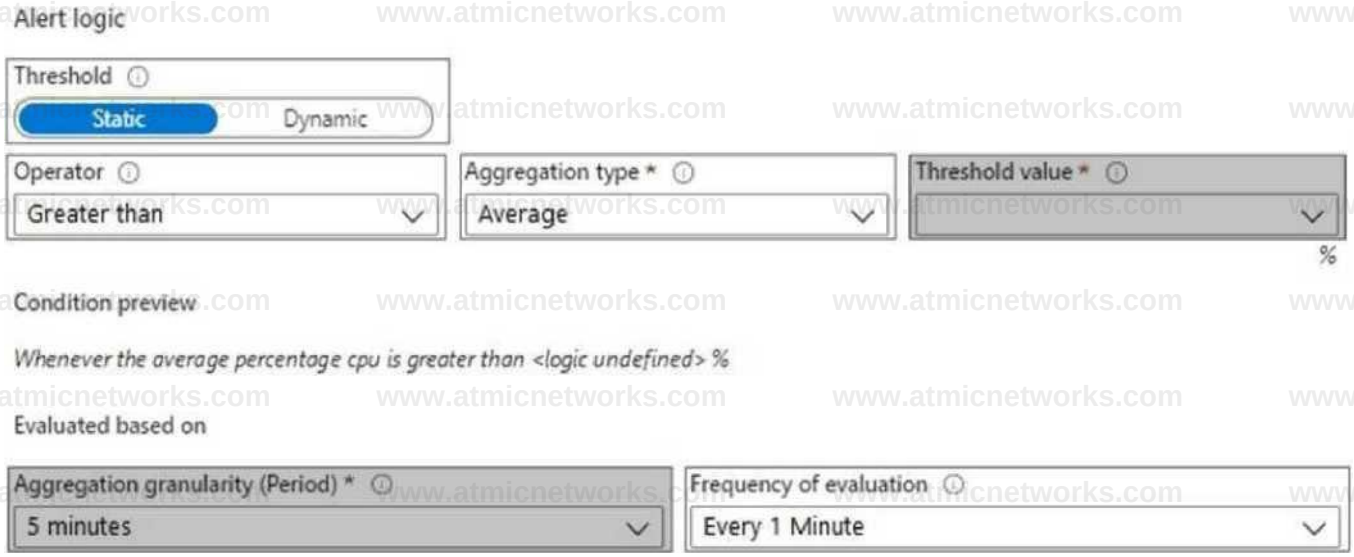

Setting 1: Threshold value Set to 80 %

Scenario: An Azure Monitor alert for VM1 must be configured to meet the following requirements: Be triggered whenaverage CPU usage exceeds 80 percent for 15 minutes.

Calculate CPU usage averages once every minute.

Setting 2: Aggregation granularity

Set to 15 minutes.

Question: 26

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)

Answer: B

Explanation:

Scenario: Access to Azure DevOps mustbe 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: 27

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)

Answer: D

Explanation:

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

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

Reference:

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

Question: 28

DRAG DROP

You need to replace the existing DevOps tools to support the planned changes.

What should you use? To answer, drag the appropriate tools to the correct targets. Each tool may be usedonce, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE:Each correct selection is worth one point.

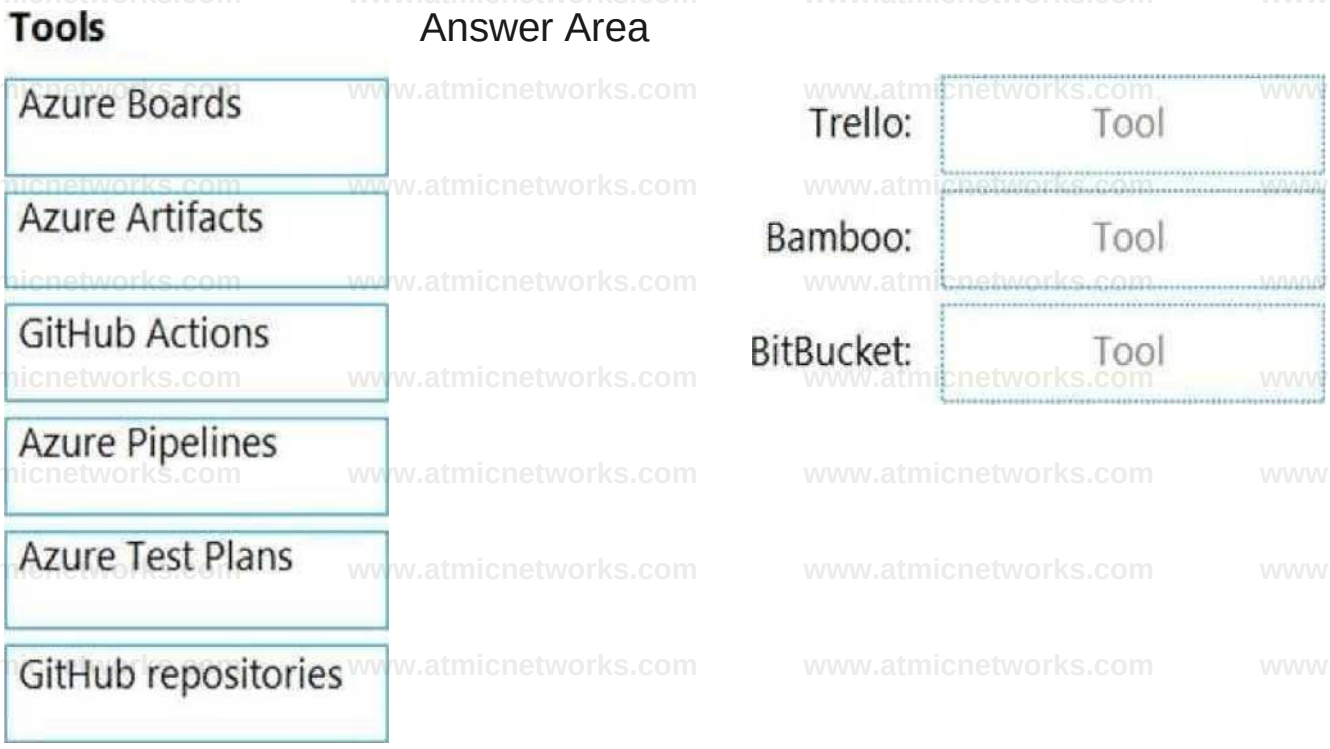

Answer:

Explanation:

Box 1: Azure Boards

Azure Boards can be used to track work with Kanban boards, backlogs, team dashboards, and custom reporting

You can create multiple Trello boards, which are spaces to store tasks (for different work contexts, or even private boards)

You can easily shareTrello boards with another person.

Box 2: Azure Pipelines

You can use Bamboo to implement CI/CD (Continuous Integration and Continuous Delivery) for a simple Azure function app using Atlassian Bamboo. Bamboo does continuous delivery of the project fromsource code to deployment. It has stages including Build, Test and Deploy.

Software teams in every industry are upgrading their continuous delivery pipeline with Bamboo. Easy build import from popular open source tools, user and group import from JIRA, seamless integration with Bitbucket, and native support for Git, Hg, and SVN means you'll be building and deploying like a champ.

Box 3: GitHub repositories

Bitbucket can be used as the Git repository, but you can use any other Git repository (Like TFS Git)for source control of the code.

Reference:

https://www.trustradius.com/compare-products/azure-devops-services-vs-trello

https://marketplace.visualstudio.com/items?itemName=ms-vsts.vss-services-bamboo

https://www.c-sharpcorner.com/article/cicd-implementation-for-an-azure-function-app-using- atlassian-bamboo-server/

Question: 29

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.

Actions Commands Cmdlets Statements

- Create an app.
- Add a secret
- Create a credential.
- Configure the ID and secret for App1
- Create a managed service identity

Answer Area

Answer:

Explanation:

Create an app

Create a managed service identity

Configure the ID and secret for App1

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: 30

You need to perform the GitHub code migration. The solution must support the planned changes for the DevOps environment.

What should you use?

A. git clone
B. GitHub Importer
C. Import repository in Azure Repos
D. git-tfs

Answer: A

Explanation:

Woodgrove Bank plans to implement the following changes to the DevOps environment:
Migrate all the source code from TFS1 to GitHub.

The Git-TFS tool is a two-way bridge between Team Foundation Version Control and Git, and can be used to perform a migration.

Reference:
https://docs.microsoft.com/en-us/devops/develop/git/migrate-from-tfvc-to-git

Question: 31

DRAG DROP

You are configuring the Azure DevOps dashboard. The solution must meet the technical requirements.

Which widget should you use for each metric? To answer, drag the appropriate widgets to the correct metrics. Each widgetmay be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE:Each correct selection is worth one point.

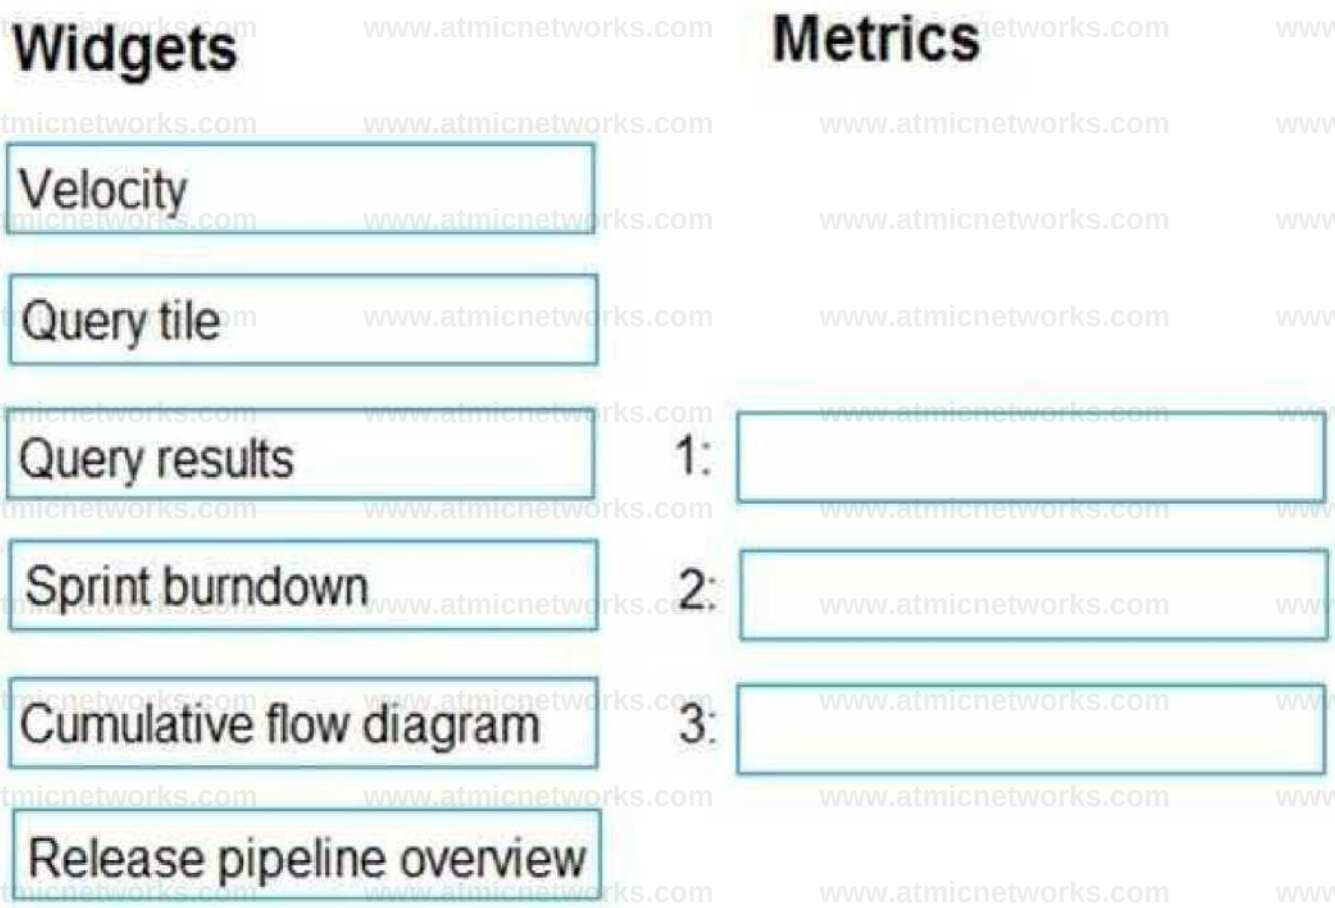

Answer:

Explanation:

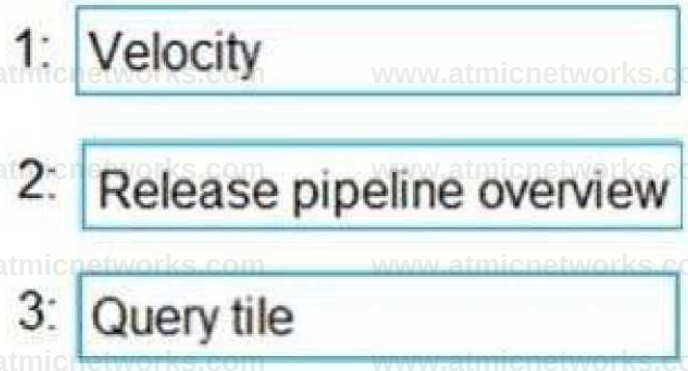

Woodgrove Bank identifies the following technical requirements:
The Azure DevOps dashboard must display the metrics shown in the following table:

Box 1: Velocity
Velocity displays your team velocity. It shows what your team delivered as compared to plan.

Box 2: Release pipeline overview
Release pipeline overview shows the status of environments in a release definition.

Box 3: Query tile
Query tile displays the total number of results from a query.

Reference:
https://docs.microsoft.com/en-us/azure/devops/report/dashboards/widget-catalog?view=azure- devops

Question: 32

You plan to deploy a new database environment. The solution must meet the technical requirements. You need to prepare the database for the deployment. How should you formatthe export?

A. NDF
B. MDF
C. BACPAC
D. DACPAC

Answer: D

Explanation:

Question: 33

You need to meet the technical requirements for monitoring App1. What should you use?

A. Splunk
B. Azure Application Insights
C. Azure Advisor
D. App Service logs

Answer: B

Explanation:

Question: 34

You need !0 the merge the POC branch into the default branch. The solution must meet the technical requirements. Which command should you run?

A. it push
B. git merge -- allow-unrelated-histories
C. git rebase
D. git merge --squash

Answer: C

Explanation:

Question: 35

You need to the merge the POC branch into the main branch. The solution must meet the technical requirements.

Which command should you run?

A. git push
B. git merge --squash
C. git merge --allow-unrelated-histories
D. git rebase

Answer: A

Explanation:

Topic 4, Mix Questions Set

Question: 36

Note: This question is part of a series les 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 m this section, you will NOT be able to return to it As a result, these questions will not appear in the review screen

You integrate a cloud-hosted Jenkins server and a new Azure Dev Ops deployment.

You need Azure Dev Ops Io send a notification to Jenkins when a developer commits changes to a branch in Azure Repos.

Solution: You create an email subscription to an Azure DevOps notification.

Does this meet the goal?

A. Yes

B. NO

Answer: B

Explanation:

You can create a service hook for Azure DevOps Services and TFS with Jenkins.

Reference:
https://docs.microsoft.com/en-us/azure/devops/service-hooks/services/jenkins

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 m.ght 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 win NOT be able to return to it. As a result, these questions win not appear in the review screen.

You integrate a cloud- hosted Jenkins server and a newAzure DevOps deployment

You need Azure DevOps to send a notification to Jenkins when a developer commits changes to a branch in Azure Repos.

Solution: You create a service hook subscription that uses the code pushed event.

Does this meet the goal?

A. Yes
B. B. NO

Answer: A

Explanation:

You can create a service hook for Azure DevOps Services and TFS with Jenkins.

Reference:
https://docs.microsoft.com/en-us/azure/devops/service-hooks/services/jenkins

Question: 38

Note: This question part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the staled goals. Some question sets might have more than one correct solution, whale 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 integrate a cloud-hosted Jenkins server and a new Azure DevOps depsoyment.

You need Azure DevOps to send a notification to Jenkins when a developer commits changes to a branch in Azure Repos.

Solution: You add a trigger to the build pipeline.

Does this meetthe goal?

A. Yes
B. NO

Answer: A

Explanation:

You can create a service hook for Azure DevOps Services and TFS with Jenkins.

Reference:
https://docs.microsoft.com/en-us/azure/devops/service-hooks/services/jenkins

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.

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 haw a pokey stating that approvals must occur within eight hours.

You discover that deployments only if the approvals take longer than two hours.

You need to ensure that the deployments only fail if the approvals take longer than hours.

Solution From Post -deployment conditions, you modify the Timeout setting for post-deployment approvals.

Does this meet the goal?

A. Yes
B. NO

Answer: B

Explanation:

Question: 40

Note: This question is part of * 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 sett 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 reviewscreen.

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 poky stating that approvals must occur within eighthour.

You discover that deployments 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

Answer: B

Explanation:

Use a gate From Pre-deployment conditions instead.

Reference:https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/gates

Question: 41

Note: This question n part of a series of questions that present the samescenario. 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 deployments fail if the approvals lake 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

Answer: B

Explanation:

Use a gate instead of an approval instead.

Reference:https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/gates

Question: 42

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 prefect in Azure DevOps for a new web application.

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

Solution: from the Triggers tab of the build pipeline, you select Enable continuous integration Does the meet the goal?

A. Yes
B. No

Answer: A

Explanation:

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

Question: 43

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 ft. As a result these questions will not appear in the review screen.

You company has a prefect inAzure 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 selected Batch changes while a build is in progress Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

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

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 thePull request triggersetting.

Does the meet the goal?

A. Yes
B. No

Answer: B

Explanation:

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

Select the Triggers tab.

Enable Continuous integration.

Reference:

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

Question: 45

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:2.1-sdk
02 COPY ./
03 RUN dotnet publish -c Release -o out
04 FROM microsoft/dotnet:2.1-sdk
05 COPY -from=0 /out /
```

```
06 WORKDIR /
07 ENTRYPOINT ["dotnet"; "appl.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

Answer: C

Explanation:

https://github.com/dotnet/dotnet-docker/blob/master/samples/dotnetapp/README.md

Question: 46

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

The company uses Azure DevOps for its 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. AzureDevOps self-hosted agents on Azure DevTest Labs virtual machines.
B. AzureDevOpsself-hosted agents on virtual machine that run on Azure Stack
C. AzureDevOps self-hosted agents on Hyper-V virtual machines
D. a Microsoft-hosted agent

Answer: B

Explanation:

Azure Stack offers virtual machines (VMs) as one type of an on-demand, scalablecomputing 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: 47

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

Answer: A

Explanation:

Use Publish Code CoverageResults 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.

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

Question: 48

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

- Minimizes image sues
- Minimizes the security surface area of the final image

What should you include m the recommendation?

A. multi-stage builds
B. single-stage builds
C. PowerShell Desired State Configuration (DSC)
D. Docker Swarm

Answer: A

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.

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

Question: 49

DRAG DROP

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

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

Youneed to conjure Azure Container Instances as a hosted environment for running me containers in AKS.

Which three actions should you perform m sequence?

To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

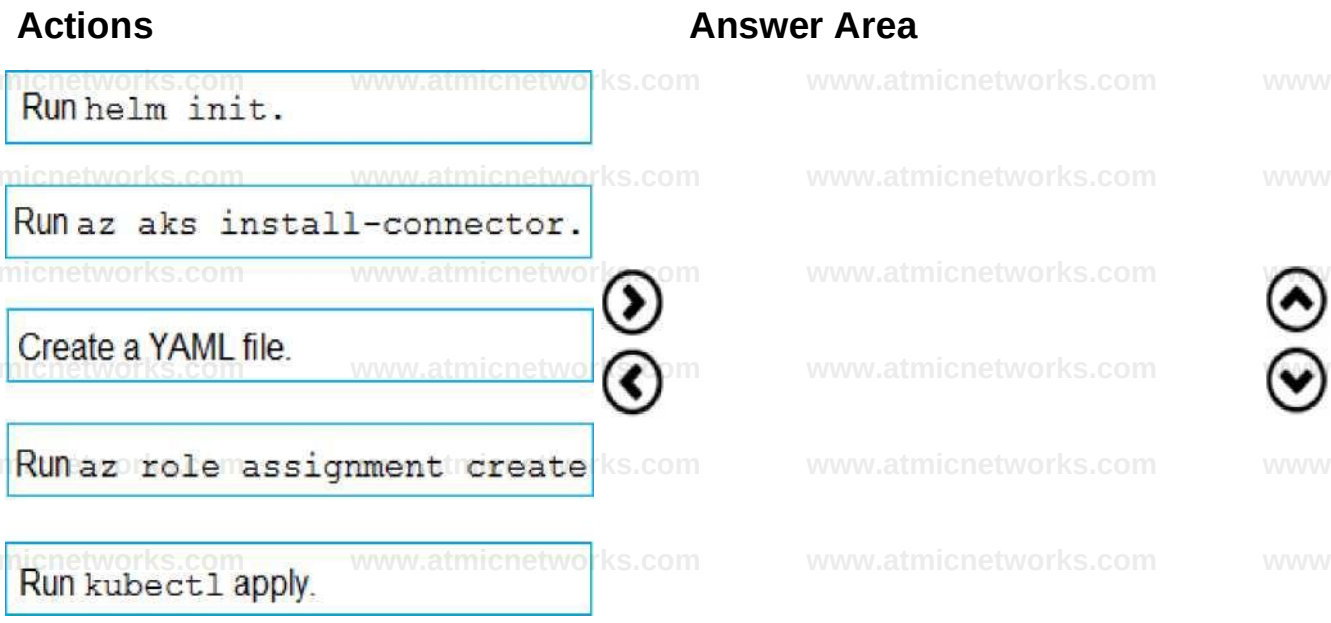

Answer:

Explanation:

Create a YAMLfile

Run kubecti apply.

Run helm init

Step 1: Create a YAML file.
If your AKS cluster is RBAC-enabled, you must create a service account and role binding for use with Tiller. To create a service account and role binding, create a file named rbac-virtual-kubelet.yaml

Step 2: Run kubectl apply.

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

Step 3: Run helm init.
Configure Helm to use the tiller service account:

helm init --service-account tiller

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

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

Question: 50

DRAG DROP

You need to use Azure Automation Sure 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 in correct. You writ receive credit for any of the orders you select.

Actions	**Answer Area**
Onboard the virtual machines to Azure Automation State Configuration.	
Check the compliance status of the node.	
Create a management group.	
Compile a configuration into a node configuration.	
Upload a configuration to Azure Automation State Configuration.	
Assign tags to the virtual machines.	

Answer:

Explanation:

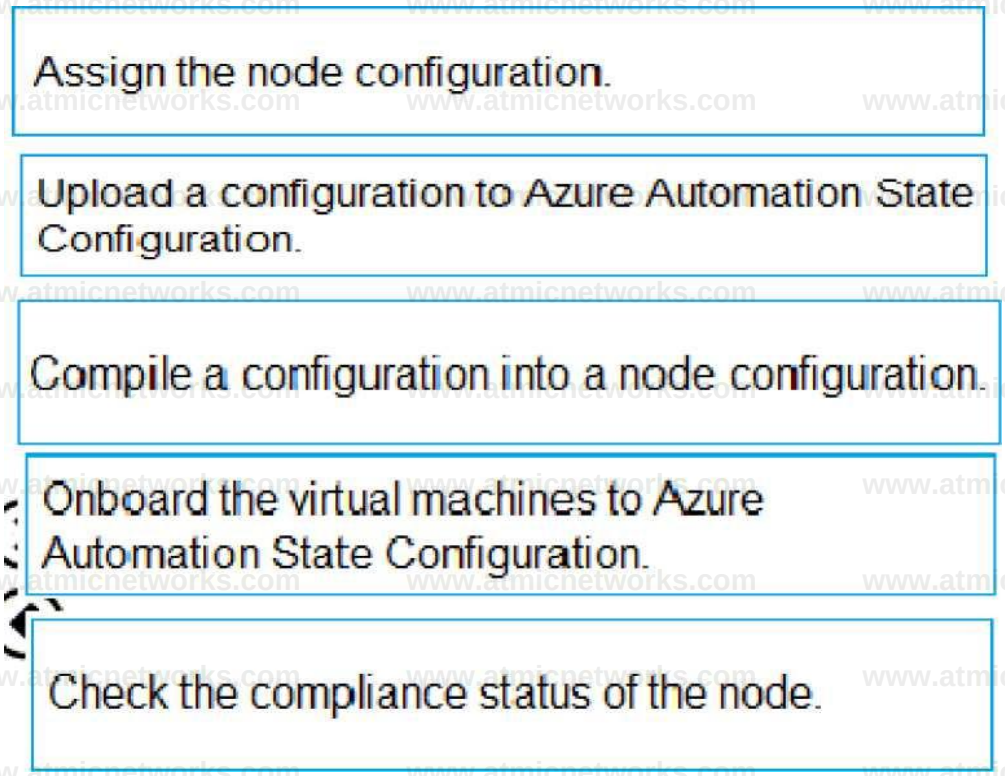

Step 1: Assign the node configuration.
You create a simple DSC configuration that ensures either the presence or absence of the Web-Server Windows Feature (IIS), depending on how you assign nodes.

Step 2: Upload a configuration to Azure Automation State Configuration.
You import the configuration into the Automation account.

Step 3: Compiling a configuration into a node configuration
Compiling a configuration in Azure Automation
Before you can apply a desired state to a node, a DSC configuration defining that state must be compiled into one or more node configurations (MOF document), and placed on the Automation DSC Pull Server.

Step 4: Onboard the virtual machines to Azure State Configuration
Onboarding an Azure VM for management with Azure Automation State Configuration

Step 5: Check the compliance status of the node.
Viewing reports for managed nodes. Each time Azure Automation State Configuration performs a consistency check on a managed node, the node sends a status report back to the pull server. You can view these reports on the page for that node.

On the blade for an individual report, you can see the following status information for the corresponding consistency check:
The report status — whether the node is "Compliant", the configuration "Failed", or the node is "Not

Compliant" (when the node is in ApplyandMonitor mode and the machine is not in the desired state).

Reference:https://docs.microsoft.com/en-us/azure/automation/automation-dsc-getting-started

Question: 51

Your company builds a multi tier web application.

>Youuse Azure DevOps and host the production application on Azure virtual machines.

Your team prepares an Azure Resource Manager template of the virtual machine that you mil use to test new features.

You need to create a staging environment in Azure that meetsthe following requirements:

- Minimizes the cost of Azure hosting
- Provisions the virtual machines automatically
- Use* the custom Azure Resource Manager template to provision the virtual machines

What should you do?

A. In Azure DevOps, configure new tasks in the release pipeline to create and delete the virtual machines in Azure DevTest Labs.

B. From Azure Cloud Shell, run Azure PowerShell commands to create and delete the new virtual machines in a staging resource group.

C. In Azure DevOps, configurenew tasks in the release pipeline to deploy to Azure Cloud Services.

D. In Azure Cloud Shell, run Azure CLI commands to create and delete the new virtual machines in a staging resource group.

Answer: A

Explanation:

You can use the Azure DevTest Labs Tasksextension that's installed in Azure DevOps to easily integrate your CI/CD build-and-release pipeline with Azure DevTest Labs. The extension installs three tasks:

Create a VM

Create a custom image from a VM

Delete a VM

The process makes it easy to, for example, quickly deploy a "golden image" for a specific test task and then delete it when the test is finished.

Reference:https://docs.microsoft.com/en-us/azure/lab-services/devtest-lab-integrate-ci-cd-vsts

Question: 52

HOTSPOT

You have a project Azure DevOps.

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

You need to ensure that you can dynamically generate the resource ID of thekey vault during template deployment.

What should you include in the template? To answer, select eh appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
'"resources": [ { "apiversion": "2018-05-01", "name" : "secrets",
"type": F                                               T
            "Microsoft    KeyVaulVvaults",    "Microsoft.Resources/deploymenf.    "Microsoft
  Subscription/subscriptions". "properties":{ "mode" : "Incremental",   ▼   :   {
  "deployment" "template" "temp late Link"

contentversion" : "1.0.0.0",
            "uri" : "[uri(parameters('_artifactsLocation'),
            concat('./nested/sqlserver.json',
            parameters ('_artifactsLocationSasToken')))]"
      }i
      "parameters": {
          "secret":{
            "reference":{
              "keyvault":{
                "id": "[resourceld(parameters('vaultsubscription')
                parameters('vaultResourceGroupName'),
                'Microsoft.KeyVault/vaults',
          i     parameters('vaultName'))]"
              }r
              "secretName": "[parameters('secretName')]" }
              }
             }
           }
         }
       ],
```

Answer:

Explanation:

```
"resources":[
{ "apiversion": "2018-05-01", "name" : "secrets", "type":
           v
           "Microsoft. KeyVautt/vaults", "Microsoft-Resources/deploymenf.
           "Microsoft Subscnption/subscnptions".
  "properties": { "mode" : "Incremental",   ▼ : {
  "deployment" "template" "template Link"

contentVerston" : "1.0.0.0",
            "uri" : "[uri(parameters('_artifactsLocation'),
            concat('./nested/sqlserver.json',
            parameters('_artifactslocationSasTokenf)))]"
  } r
  "parameters":{
      "secret":{
```

```
"reference":{
   "keyVault":{
      "id": "[resourceId(parameters('vaultSubscription')
      parameters('vaultResourceGroupName'),
       'Microsoft.KeyVault/vaults',
,      parameters('vaultName' ))]"
} /
 "secretName": "[parameters('secretName')]"
  }
```

Question: 53

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 isworth one point.

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	

Answer:

Explanation:

Answer Area

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

Box 1: A key Vault advanced access policy

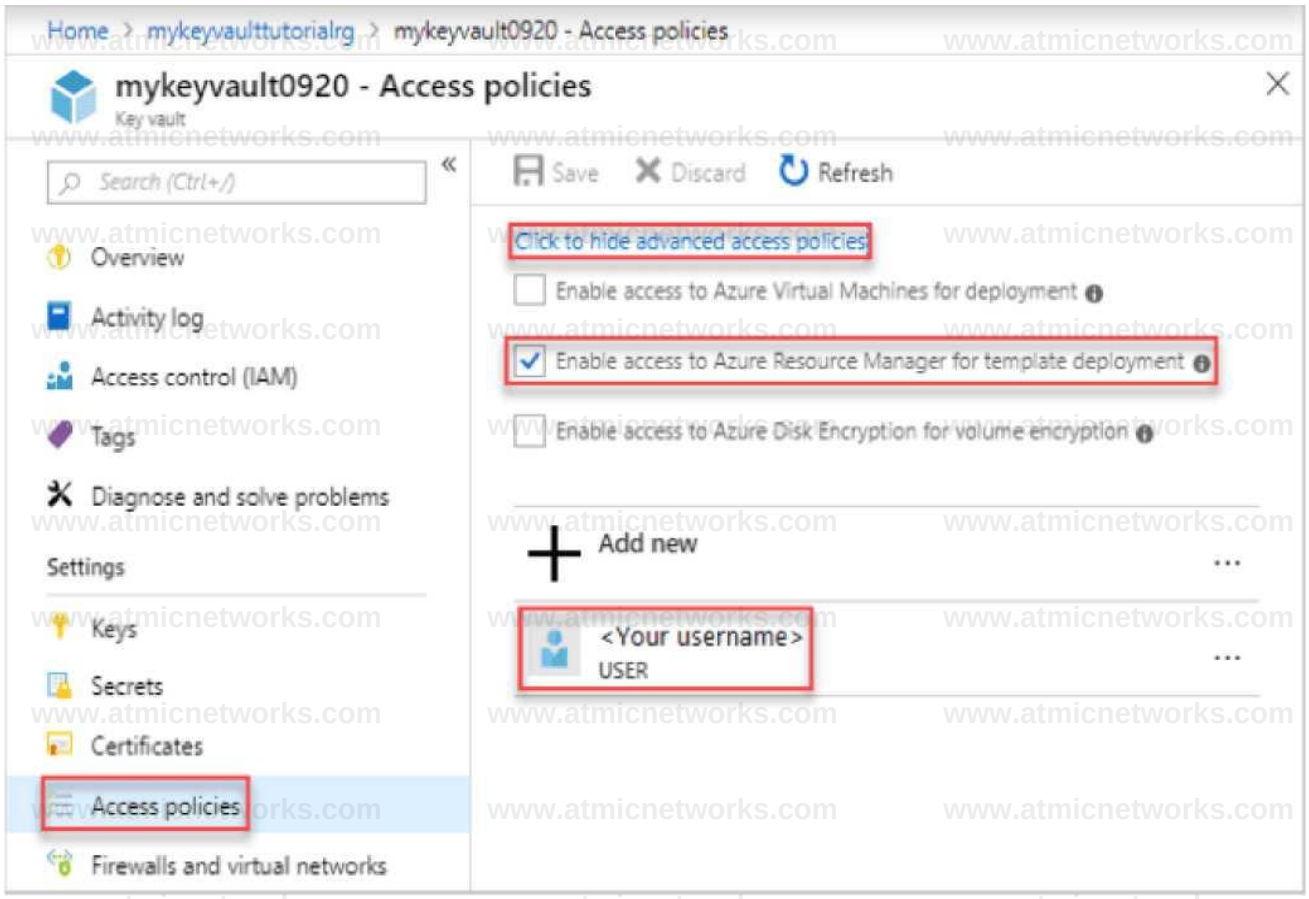

Box 2: RBAC

Management plane access control uses RBAC.
The management plane consists of operationsthat affect the key vault itself, such as:

Creating or deleting a key vault.
Getting a list of vaults in a subscription.
Retrieving Key Vault properties (such as SKU and tags).

Setting Key Vault access policies that control user and application access tokeys and secrets.

Reference:https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-manager- tutorial-use-key-vault

Question: 54

DRAG DROP

You need to recommend a solution for deploying charts by using Helm and Title to Azure Kubemets Service (AKS) in an RBAC-enabled cluster.

Which threecommands should you recommend be run m sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

Commands

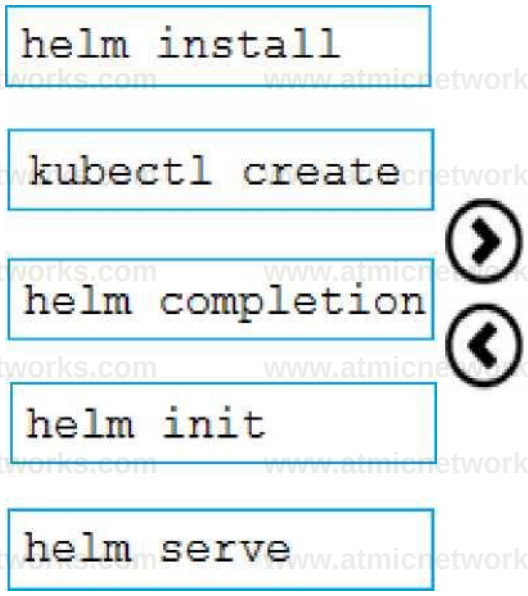

Answer Area

Answer:

Explanation:

Answer Area

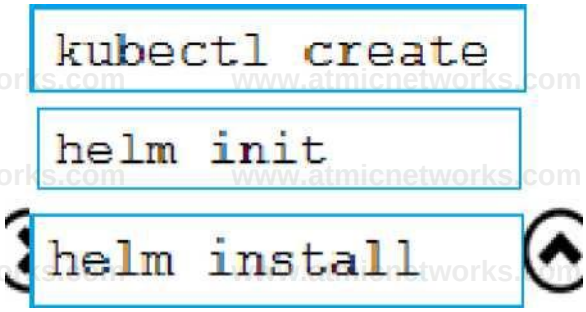

Step 1: Kubectl create

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.

Reference:

https://docs.microsoft.com/en-us/azure/aks/kubernetes-helm

https://docs.helm.sh/using_helm/#tiller-namespaces-and-rbac

Question: 55

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

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

Answer: B

Explanation:

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: 56

During a code review, you discover many quality issues. Many modulescontain unused variables and empty catch Modes. You need to recommend a solution to improve the quality o' the code. What should you recommend?

A. In a Gradle build task, select Run Checkstyle.
B. In an Xcode build task, select Use xcpretty from Advanced
C. In a Grunt build task, select Enabled from Control Options.
D. In a Maven build task, select Run PMD.

Answer: D

Explanation:

PMD is a source code analyzer. It finds common programming flaws like unused variables, empty catch blocks,unnecessary object creation, and so forth.

There is an Apache Maven PMD Plugin which allows you to automatically run the PMD code analysis tool on your project's source code and generate a site report with its results.
Reference:https://pmd.github.io/

Question: 57

DRAG DROP
You need to increase the security of your team's development process.
Which type of security tool should you recommend for each stage of the development process? To answer, drag the appropriate security toots to the correct stages. Each security tool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to viewcontent
NOTE: Each correct selection is worth one point.

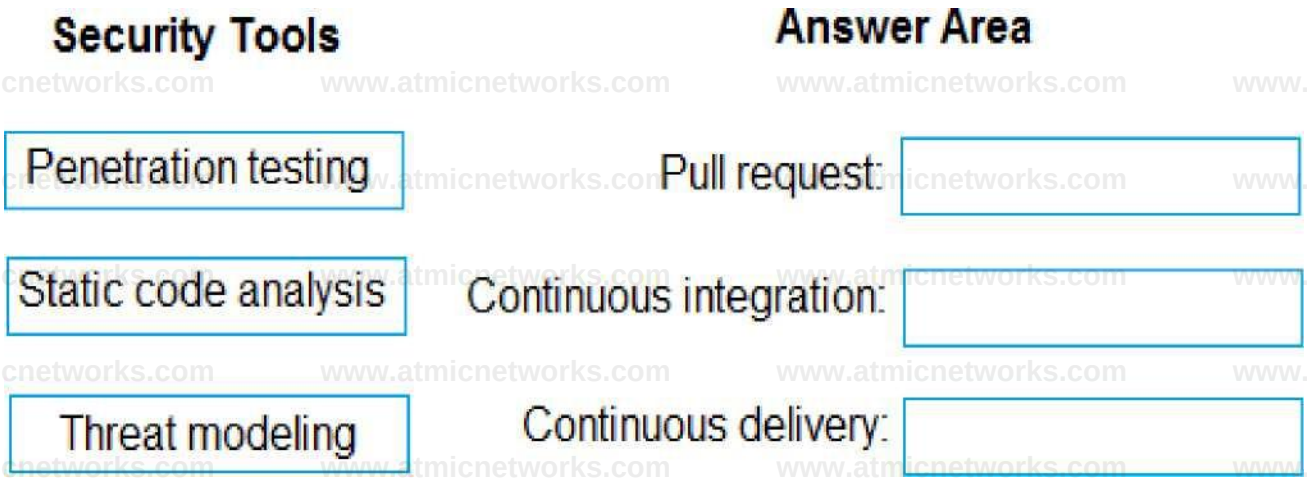

Answer:

Explanation:

https://docs.microsoft.com/en-us/azure/devops/migrate/security-validation-cicd- pipeline?view=azure-devops&viewFallbackFrom=vsts
So:
PR: Static Code Analysis
CI: Static Code Analysis

CD: PenTest

Question: 58

Your company plans to use an agile approach to software development
You need to recommendan application to provide communication between members of the development team who work in locations around the world. The application must meet the following requirements:

- Provide the ability to isolate the members of efferent project teams into separate communication channels and to keep a history of the chats within those channels.
- Be available on Windows 10, Mac OS, iOS, and Android operating systems.

- Provide the ability to add external contractors and suppliers to projects.
- Integrate directlywith Azure DevOps.

What should you recommend?

A. Octopus
B. Bamboo
C. Microsoft Project
D. Slack

Answer: D

Explanation:

Slack is a popular team collaboration service that helps teams be more productive by keeping all communications in one place and easilysearchable from virtually anywhere. All your messages, your files, and everything from Twitter, Dropbox, Google Docs, Azure DevOps, and more all together. Slack also has fully native apps for iOS and Android to give you the full functionality of Slack wherever you go.
Integrated with Azure DevOps
This integration keeps your team informed of activity happening in its Azure DevOps projects. With this integration, code check-ins, pull requests, work item updates, and build events show up directly in your team's Slack channel.

Note: Microsoft Teams would also be a correct answer, but it is not an option here.

Reference:
https://marketplace.visualstudio.com/items?itemName=ms-vsts.vss-services-slack

Question: 59

Your company uses Azure DevOps for the build pipelines and deployment pipelines of Java based projects. You need to recommend a strategy for managing technical debt.Which two actions should you include in the recommendation? Each correct answer presents part of the solution

NOTE: Each correct selection is worth one point.

A. Integrate Azure DevOps and SonarQube.
B. Integrates Azure DevelOPs and Azure DevTest Labs.
C. Configure post-deployment approvals in the deployment pipeline.
D. Configure pre-deployment approvals in the deployment pipeline.

Answer: A,C

Explanation:

Question: 60

DRAG DROP

You need to recommend project metrics for dashboards in Azure DevOps.

Which chart widgetsshould you recommend for each metric? To answer, drag the appropriate chart widgets to the correct metrics. Each chart widget 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.

Chart Widgets	Answer Area
Burndown	The elapsed time from the creation of work items to their completion:
Cycle Time	
Lead Time	The elapsed time to complete work items once they are active:
Velocity	The remaining work:

Answer:

Explanation:

Answer Area

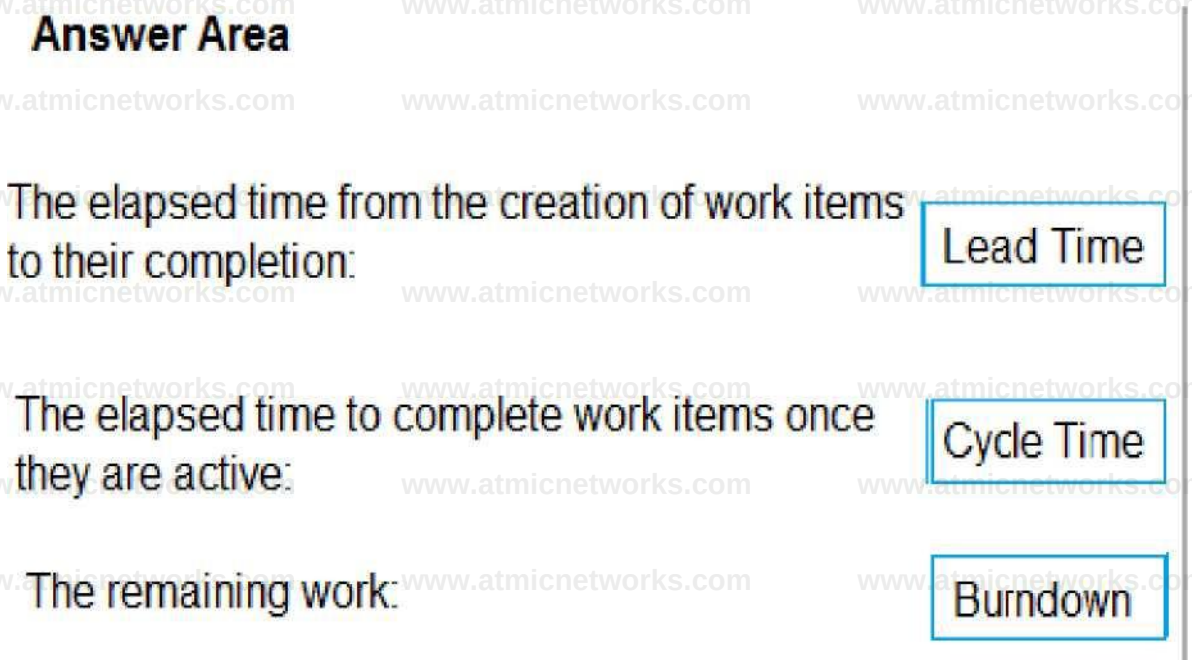

Box 1: Lead time
Lead time measures the total time elapsed from the creation of work items to their completion.

Box 2: Cycle time
Cycle time measures the time it takes for your team to complete work items once they begin actively working onthem.

Box 3: Burndown
Burndown charts focus on remaining work within a specific time period.

Incorrect Answers:
A: guide for determining how well the team estimates and meets their planned commitments

Reference:
https://docs.microsoft.com/en-us/azure/devops/report/dashboards/velocity-guidance?view=vsts

https://docs.microsoft.com/en-us/azure/devops/report/dashboards/cycle-time-and-lead- time?view=vsts

https://docs.microsoft.com/en-us/azure/devops/report/dashboards/configure-burndown-burnup- widgets?view=vsts

Question: 61

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 answerarea.

NOTE:Each correct selection is worth one point.

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:

Explanation:

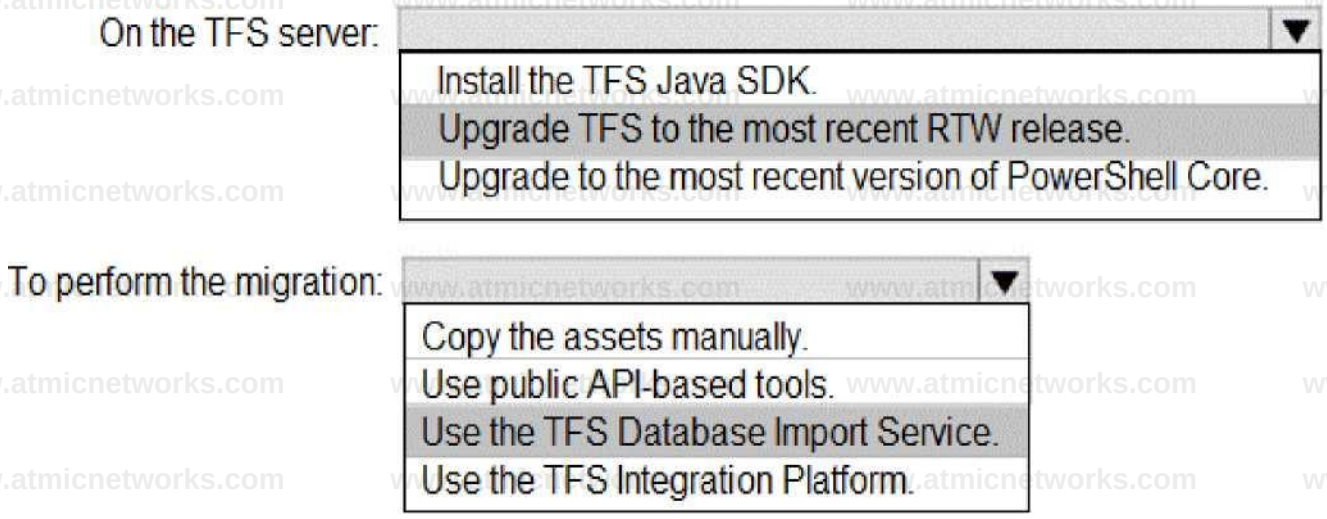

Box 1: Upgrade TFS to the most recent RTM release.
One of the major prerequisites for migrating your Team Foundation Server database is to get your database schema version as close as possible to what is currently deployed in Azure Devops Services.

Box 2: Use the TFS Database Import Service
In Phase 3 ofyour 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.

Reference: Team Foundation Server to Azure Devops Services Migration Guide

Question: 62

Your company deploys applications m Docket containers.

You want to detect known exploits in the Docket 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 deployment pipeline and a scheduled task against a running production container.

B. a task executed in the continuous integration pipeline and a scheduled task that analyzes the production container.

C. a task executed in the continuous integration pipeline and a scheduled task that analyzes the image registry

D. manual tasks performed during the planning phase and the deployment phase

Answer: C

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 usedwith Docker or Azure Container registry.

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

Question: 63

You are developinga multi-tier application. The application will use Azure App Service web apps as the front end and an Azure SQL database as the back end. The application will use Azure functions to write some data to Azure Storage.

You need to send the Azure DevOps teaman email message when the front end fails to return a status code of 200.

Which feature should you use?

A. Service Map in Azure Log Analytics

B. Profiler in Azure Application Insights

C. availability tests in Azure Application Insights

D. ApplicationMap in Azure Application Insights

Answer: D

Explanation:

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

Question: 64

DRAG DROP

Your company plans to deploy an application to the following endpoints:

- Ten virtual machines hosted in Azure.
- Ten virtual machines hosted in an on-premises data center environment

All the virtual machines have the- Azure Pipelines agent.

You need to implement a release strategy for deploying the application to the endpoints.

What should you recommend using to deploy the application to the endpoints? To answer, drag the appropriate components to the correct endpoint.

Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or soon to view content

NOTE: Each correct selection n worth one point.

Components

- A deployment group
- A management group
- A resource group
- Application roles

Answer Area

Ten virtual machines hosted in Azure: [ ]

Ten virtual machines hosted in an on-premises data center environment: [ ]

Answer:

Explanation:

Box 1: A deployment group

When authoring an Azure Pipelines or TFS Release pipeline, you can specify the deployment targets for a job using a deployment group.

If the target machines are Azure VMs,you can quickly and easily prepare them by installing the Azure Pipelines Agent Azure VM extension on each of the VMs, or by using the Azure Resource Group Deployment task in your release pipeline to create a deployment group dynamically.

Box 2: A deployment group

Reference:https://docs.microsoft.com/en-us/azure/devops/pipelines/release/deployment-groups

Question: 65

DRAG DROP

You need to configure access to Azure DevOps Agent pools to meet the forwarding requirements:

- Use a project agent pool when authoring build 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 0e%Oos organization and the project? To answer, drag the appropriate role membership 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 content

NOTE: Each correct selection is worth one point.

Roles

Administrator

Reader

Service Account

User

Answer Area

Organization: [ ]

Project: [ ]

Answer:

Explanation:

project level role: User
Organization level role: Reader

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

Question: 66

HOTSPOT

You are configuring a release pipeline in Azure DevOps as shown in the exhibit.

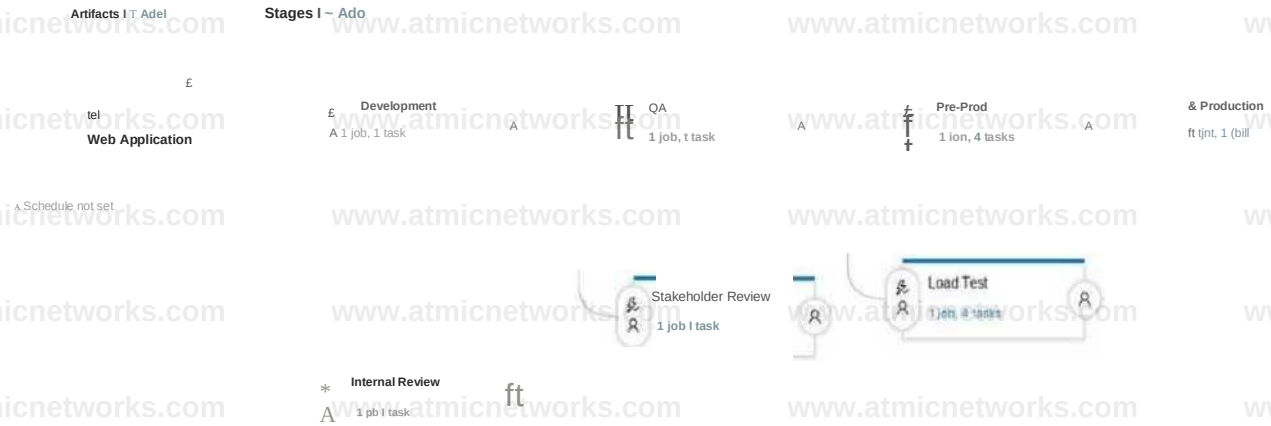

Use thedrop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

NOTE:Each correct selection is worth one point.

How many stages have triggers set?

- 0
- 1
- 2
- 3
- 4
- 5
- 6
- 7

Which component should you modify to enable continuous delivery?

- The Development stage
- The Internal Review stage
- The Production stage
- The Web Application artifact

Answer:

Explanation:

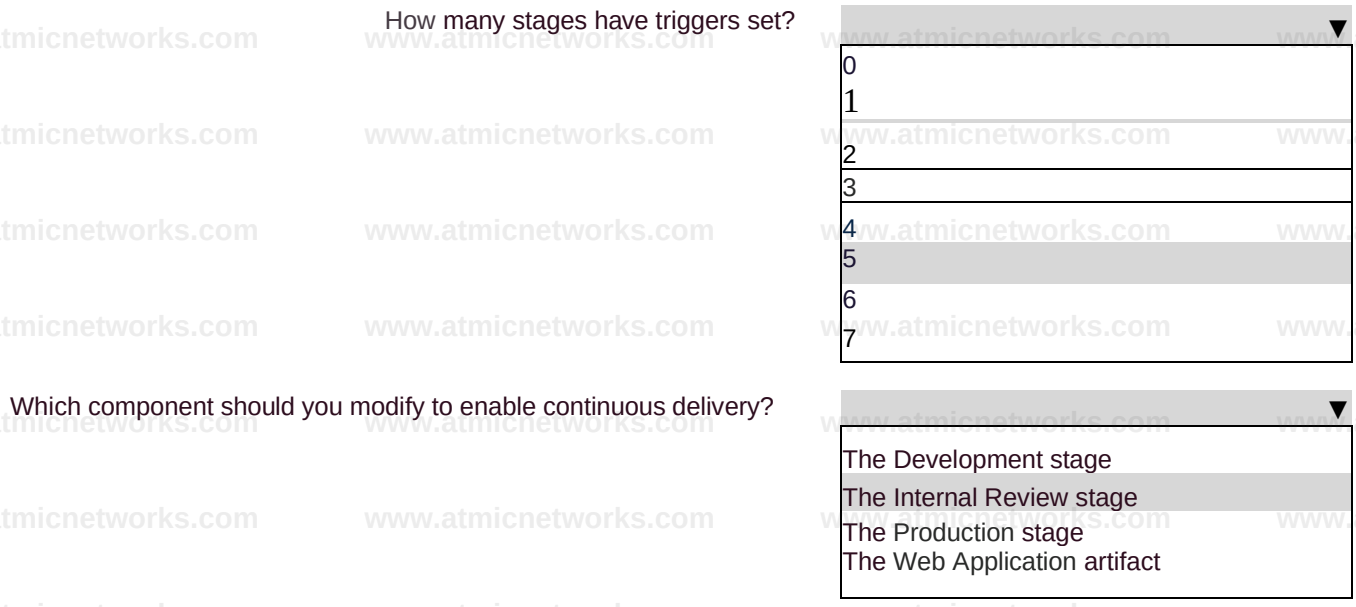

Box 1: 5

There are five stages: Development, QA, Pre-production, Load Test and Production. They all have triggers.

Box 2: The Internal Review stage

Reference:https://docs.microsoft.com/en-us/azure/devops/pipelines/release/triggers

Question: 67

Your companyhas a project in Azure DevOps for a new web application.

The company uses Service Now 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 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 Service Now SOAP API.
B. Define a post deployment gate afterthe deployment to the QA stage.
C. Define a deployment control that invokes the ServiceNow REST API.
D. Define a pre deployment gate before the deployment to the Prod stage.

Answer: A,B

Explanation:

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

Question: 68

Your company has an on-premises Bitbucket Server that is used for Git-based source control. The server is protected by a firewall that blocks inbound Internet traffic.
You plan to use Azure DevOps to manage the build and release processes
Which two components are required to integrate Azure DevOps and Bitbucket?
Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one port.

A. an External Git service connection
B. a Microsoft hosted agent
C. service hooks
D. a self- hosted agent
E. a deployment M group

Answer: A,D

Explanation:
When a pipeline uses a remote, 3rd-party repository host such as Bitbucket Cloud, the repository is configured with webooks that notify Azure Pipelines Server or TFS when code has changed and a build should be triggered. Since on-premises installations are normally protected behind a firewall, 3rd-party webhooks are unable to reach the on-premises server. As a workaround, you can use the External Git repository type which uses polling instead of webhooks to trigger a build when code has changed.

Reference:https://docs.microsoft.com/en-us/azure/devops/pipelines/repos/pipeline-options-for

Question: 69

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 restnct 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 • Changes must be reviewed by at least two project members • Changes must be associated to at least one work team.
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 usefor 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.

Version Control Systems

- Git
- Perforce
- Subversion
- Team Foundation Version Control

Answer Area

Project 1:
Project 2:
Project 3:
Project 4:

Answer:

Explanation:

1 -> TFVS Refer : https://docs.microsoft.com/en-us/azure/devops/repos/tfvc/control-access-team- foundation-version-control?view=azure-devops

2 -> TFVS Refer :https://docs.microsoft.com/en-us/azure/devops/repos/tfvc/add-check- policies?view=azure-devops

3 -> Git Refer : https://docs.microsoft.com/en-us/azure/devops/repos/git/share-your-code-in-git- xcode?view=azure-devops

4 -> TFVS Refer :https://docs.microsoft.com/en-us/azure/devops/organizations/security/permissions?view=azure-devops#tfvc

Question: 70

DRAG DROP

You mc 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 built agent pools to the correct application types. Each butt 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.

Build Agent

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:

Answer:

Explanation:

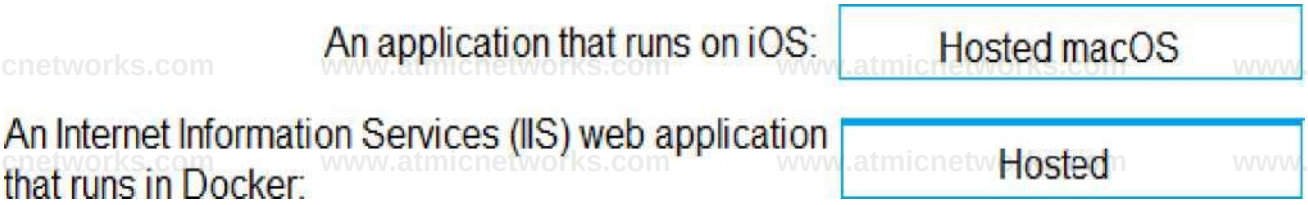

Box 1: Hosted macOS

HostedmacOS pool (Azure Pipelines only): Enables you to build and release on macOS without having to configure a self-hosted macOS agent. This option affects where your data is stored.

Box 2: Hosted

Hosted pool (Azure Pipelines only): The Hosted pool is the built-in pool that is a collection of Microsoft-hosted agents.

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

Question: 71

You have a brand 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 change to the master branch, evenif the code fails to compile. The solution must use the principle of least privilege.

What should you do?

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

Answer: B

Explanation:

In some cases, you need to bypass policy requirements so you can push changes to the branch directlyor 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 thispermission along the with other Git permissions.

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

Question: 72

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. an Azure Resource Manager parameter file
B. an Azure Storage table
C. an Appsettings.json files
D. Azure Key Vault
E. a Web.config file

Answer: D

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: 73

HOTSPOT

Your company is creating a suite of three mobile applications.

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

What should you use? To answer, select the appropriate options m the answer area.

NOTE: Each correct selection is worth one point.

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:

Explanation:

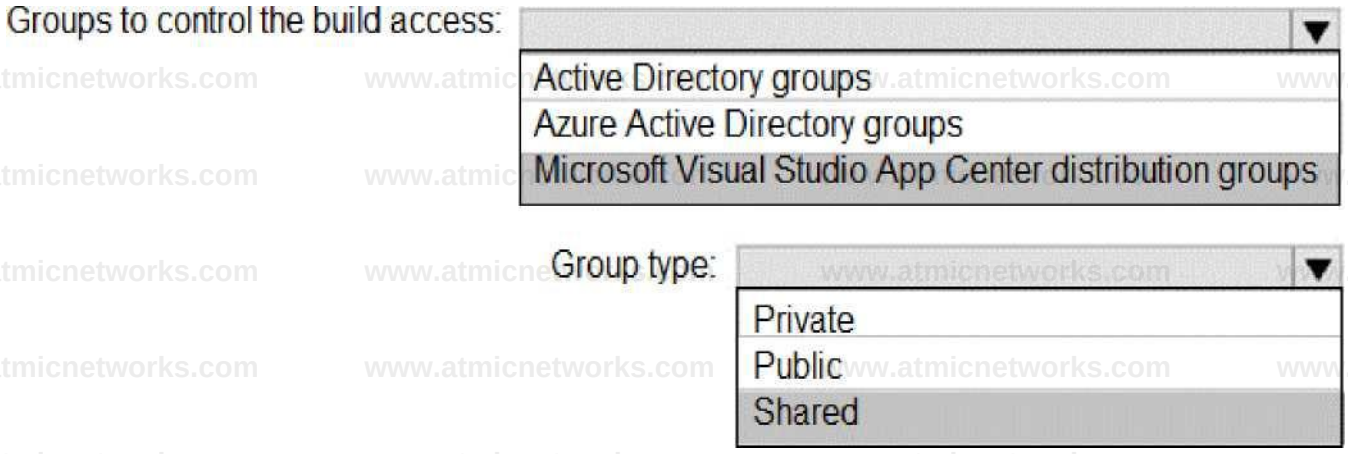

Box 1: Microsoft Visual Studio App Center distribution Groups

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. Exampleof 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.

Reference:https://docs.microsoft.com/en-us/appcenter/distribution/groups

Question: 74

Your company uses a Git repository in Azure Repos lo 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 tie changes into a tingle commit

Which merge strategy should you use in the branch policy?

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

Answer: C

Explanation:

Squash merging is a merge option that allows you to condense the Git history of topic branches when you complete a pull request. Instead of each commit on the topic branchbeing added to the history of the defaultbranch, 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.

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

Question: 75

Your company uses cloud-hosted Jenkins for builds.

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

Which three actions should you perform? Each correct answer presents part of the solution NOTE: Each correct answer selection is worth one point

A. Add the Team Foundation Server (TFS) plug-in to Jenkins.
B. Create a personal access token myour Azure DevOps account.
C. Create a webhook in Jenkins.

D. Add a domain to your Jenkins account.
E. Create a service hook m Azure DevOps.

Answer: A,B,E

Explanation:
Reference:
https://blogs.msdn.microsoft.com/devops/2017/04/25/vsts-visual-studio-team-services-integration- with-jenkins/

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

Question: 76

You are developing an open source solution that uses aGitHub 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. a personal access token
B. SAML
C. GitHub App
D. OAuth

Answer: C

Explanation:

https://docs.microsoft.com/en-us/azure/devops/pipelines/repos/github?view=azure-devops&tabs=yaml
https://developer.github.com/v3/checks/

Question: 77

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 asingle feed. The solution must limit the release of packages that are in development.
What should you use?

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

Answer: C

Explanation:

Views enable you to share subsets of the NuGet, npm, Maven, Python and Universal Packages package-versions in your feed with consumers. A common use for views is to share package versions that have been tested, validated, or deployed but hold back packages still under development and packages that didn't meet aquality bar.
https://docs.microsoft.com/en-us/azure/devops/artifacts/concepts/views?view=azure-devops

Question: 78

You use Azure Artifacts to host NuGet packages that you create.
You need tomake one of the packages available to anonymous users outside your organization. The solution must minimize the number of publication points.
What should you do?

A. Create a new feed for the package
B. Publish the package to a public NuGet repository.
C. Promote the package to a release view.
D. Change the feed URL of the package.

Answer: A

Explanation:

Azure Artifacts introduces the concept of multiple feeds that you can use to organize and control access to your packages.
Packages you host in Azure Artifacts are stored in a feed. Setting permissions on the feed allows you to share your packages with as many or as few people as your scenario requires.
Feeds have four levels of access: Owners, Contributors, Collaborators, and Readers.

Reference:https://docs.microsoft.com/en-us/azure/devops/artifacts/feeds/feed-permissions?view=vsts&tabs=new-nav

Question: 79

Your company « 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. Code Style
B. Microsoft Visual SourceSafe

C. Black Duck
D. Jenkins

Answer: C

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 andcontainer 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: 80

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. Delete Package lock.json.
C. Configure the Artifactory plug-in.
D. Add a devDependencies section to Package-lock.json.

Answer: D

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: 81

You use Azure SQL Database Intelligent Insights and Azure Application Insights foe monitoring.
You need to write ad-hoc Queries against the monitoring data.
Which Querylanguage should you use?

A. PL/pgSQL
B. Transact-SQL
C. Azure Log Analytics

D. PL/SQL

Answer: C

Explanation:

Data analysis in Azure SQL Analytics is based on Log Analytics language for your custom querying and reporting.

Reference:https://docs.microsoft.com/en-us/azure/azure-monitor/insights/azure-sql

Question: 82

Your company uses Service Now for incident management.
You develop an application that runs on Azure.
The company needs to generate a ticket in Service Now when the application fails to authenticate.
Which Azure Log Analytics solution should you use?

A. Automation &Control
B. IT Service Management Connector (ITSM)
C. Application lmiQ.hu Connector
D. insight & Analytics

Answer: B

Explanation:

The IT Service Management Connector (ITSMC) allows you to connect Azure and a supported IT Service Management (ITSM)product/service.

ITSMC supports connections with the following ITSM tools:

ServiceNow
System Center Service Manager
Provance
Cherwell

With ITSMC, you can
Create work items in ITSM tool, based on your Azure alerts (metric alerts, Activity Log alerts andLog Analytics alerts).
Optionally, you can sync your incident and change request data from your ITSM tool to an Azure Log Analytics workspace.

Reference: https://docs.microsoft.com/en-us/azure/azure-monitor/platform/itsmc-overview

Question: 83

HOTSPOT
Your company is building a new web application.
You plan to collect feedback from pilot users on the features being delivered.
All the pilot users have a corporate computer that has Google Chrome and the Microsoft Test & Feedback extension installed. Thepilot users will test the application by using Chrome.
You need to identify which access levels are required to ensure that developers can request and gather feedback from the pilot users. The solution must use the principle of least privilege.
Which access levels m Azure DevOps should you identify? To answer, select the appropriate options in the answer area
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

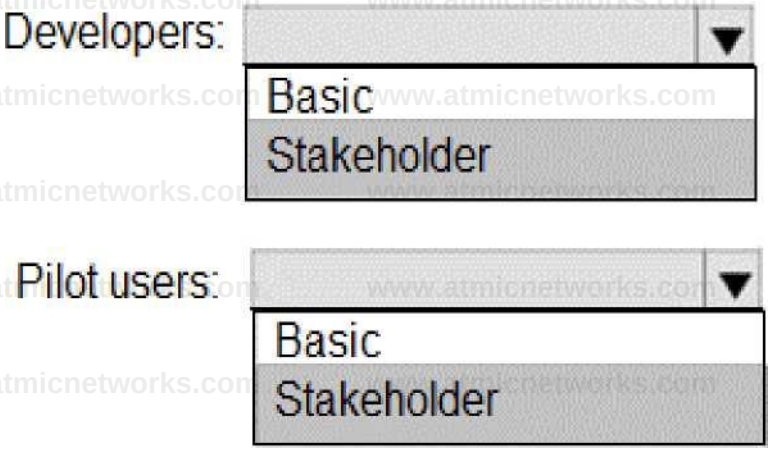

Box 1: Basic

Assign Basic to users with a TFS CAL, with a Visual Studio Professional subscription, and to users for whom you are paying for Azure Boards & Repos in an organization.

Box 2: Stakeholder

AssignStakeholders to users with no license or subscriptions who need access to a limited set of features.

Note:

You assign users or groups of users to one of the following access levels:

Basic: provides access to most features

VS Enterprise: provides access to premium features

Stakeholders: provides partial access, can be assigned to unlimited users for free

Reference:https://docs.microsoft.com/en-us/azure/devops/organizations/security/access- levels?view=vsts

Question: 84

You have multi-tier application that h« an Azure Web Apps front end and art Azure SQL Datable back end.

You need to recommend a solution tocapture and store telemetry dat

a. The solution must meet the following requirements:

- Support using ad-hoc queries to identify baselines.
- Trigger alerts when metrics in the baseline are exceeded.
- Store application and database metrics in a central location.

What should you include in the recommendation?

A. Azure Application Insights
B. Azure SQL Database Intelligent Insights
C. Azure Event Hubs
D. Azure Log Analytics

Answer: A

Explanation:

Azure Platform as a Service (PaaS) resources, like Azure SQL and Web Sites (Web Apps), can emit performance metrics data natively to Log Analytics.
The Premium plan will retain up to 12 months of data, giving you an excellent baseline ability. There are two options available in the Azure portal for analyzing data stored in Log analytics and for creating queries for ad hoc analysis.

Reference:https://docs.microsoft.com/en-us/azure/azure-monitor/platform/collect-azurepass-posh

Question: 85

Your company creates a web application.

You need to recommend a solution that automatically sends to Microsoft Teams a dairy summary of the exceptions that occur m the application.
Which two Azure services should you recommend? Eachcorrect answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Microsoft Visual Studio App Center
B. Azure DevOps Project
C. Azure Logic Apps
D. Azure Pipelines
E. Azure Application Insights

Answer: C,E

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/asp-net-exceptions

https://docs.microsoft.com/en-us/azure/azure-monitor/app/automate-custom-reports

Question: 86

Your company plans to use an agile approach to software development.

You need to recommend an application to provide communication between members of the development team who work in locations around the world. The applications must meet the followingrequirements:

Provide the ability to isolate the members of different project teams into separate communication channels and to keep a history of the chats within those channels.
Be available on Windows 10, Mac OS, iOS, and Android operating systems.
Provide the ability to add external contractors and suppliers to projects.
Integrate directly with Azure DevOps.

What should you recommend?

A. Microsoft Project
B. Bamboo
C. Microsoft Lync

D. Microsoft Teams

Answer: D

Explanation:

Within each team,users can create different channels to organize their communications by topic.
Each channel can include a couple of users or scale to thousands of users.
Microsoft Teams works on Android, iOS, Mac and Windows systems and devices. It also works in Chrome, Firefox, Internet Explorer 11 and Microsoft Edge web browsers.
The guest-access feature in Microsoft Teams allows users to invite people outside their organizations

to join internal channels for messaging, meetings and file sharing. This capability helps tofacilitate business-to-business project management.
Teams integrates with Azure DevOps.
Reference: https://searchunifiedcommunications.techtarget.com/definition/Microsoft-Teams

Question: 87

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 inthis 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 runsautomatically.

Solution: From the Pre-deployment conditions settings of the release pipeline, you selectAfter stage.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead, In Visual Designer you enable continuous integration (CI) by:
Select the Triggers tab.
Enable Continuous integration.

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

Question: 88

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, youselectBatch changes while a build is in progress.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Use a Pull request trigger.

Note: Batch changes
Select this check box if you have a lot of team members uploading changes often and you wantto 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/build/triggers

Question: 89

Note: This question is part of a series of questions that present the same scenario. Eachquestion 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 beable 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 policystating 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: FromPre-deployment conditions, you modify the Time between re-evaluation of gates

option.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Reference:https://docs.microsoft.com/en-us/azure/devops/pipelines/release/approvals/gates

Question: 90

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 createthe 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 two standalone templates, each of which will deploy the resources in its respective group.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

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

Question: 91

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 NOTbe 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

Answer: B

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: 92

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.

Afteryou 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: Createa main template that has two linked templates, each of which will deploy the resource in its respective group.

Does this meet the goal?

A. Yes
B. No

Answer: A

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: 93

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. PDM
C. WhiteSource
D. OWASP ZAP

Answer: C

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.

Azure DevOps integration with WhiteSourceBolt will enable you to:

Detect and remedy vulnerable open source components.
Generate comprehensive open source inventory reports per project or build.
Enforce open source license compliance, including dependencies' licenses.
Identify outdated open source libraries with recommendations to update.

Reference:https://www.azuredevopslabs.com/labs/vstsextend/WhiteSource/

Question: 94

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 integrate a cloud-hosted Jenkins server and a new Azure DevOps deployment.

You need Azure DevOps to senda notification to Jenkins when a developer commits changes to a branch in Azure Repos.

Solution: You create a service hook subscription that uses the build completed event

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

You can create a service hook for Azure DevOps Services and TFS with Jenkins.

However, the service subscription event should use the code pushed event, is triggered when the code is pushed to a Git repository.

Question: 95

Your company develops an app for OS. All users ofthe 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. .pvk
C. .pfx
D. .p12

Answer: D

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 validprovisioning 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: 96

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 fail to meet the performance baseline?

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

Answer: C

Explanation:

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: 97

Your company is building a mobile app that targets Android devices and OS devices. Your team uses Azure DevOps to manage all work items and release cycles. You need to recommend a solution to perform thefollowing tasks:

- Collect crash reports for issue analysis
- Distribute beta releases to your testers.
- Get user feedback on the functionality of new apps.

What should you include in the recommendation?

A. Jenkins integration
B. Azure ApplicationInsights widgets
C. the Microsoft Test & Feedback extension
D. Microsoft Visual Studio App Center integration

Answer: D

Explanation:

The "Exploratory Testing" extension is now "Test & Feedback" and is now Generally Available.

Anyone can now test web appsand give feedback, all directly from the browser on any platform: Windows, Mac, or Linux. Available for Google Chrome and Mozilla Firefox (required version 50.0 or above) currently. Support for Microsoft Edge is in the pipeline and will be enabled once Edge moves to a Chromium-compatible web platform.

Reference:
https://marketplace.visualstudio.com/items?itemName=ms.vss-exploratorytesting-web

Question: 98

Your company is building a new solution inJava.
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. Octopus
B. Chef
C. Maven
D. Grunt

Answer: A

Explanation:

Question: 99

Your company has a project in Azure DevOps.

You need to ensure that when there are multiple builds pending deployment only the most recent build is deployed.

What should you use?

A. deployment queue settings
B. deployment conditions
C. release gates
D. pull request triggers

Answer: A

Explanation:

Reference:

https://docs.microsoft.com/en-us/azure/devops/pipelines/process/stages?tabs=classic&view=azure-devops#queuing-policies

Question: 100

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. build
B. production release
C. staging
D. integration testing

Answer: D

Explanation:

The Secure Development Lifecycle (SDL) Guidelines recommend that teams perform static analysis duringthe implementation phase of their development cycle.

Note: The company should focus in particular on the implementation of DevOps tests to assess the quality of the software from the planning stage to the implementation phase of the project.

Reference:

https://secdevtools.azurewebsites.net/

Question: 101

You have a GitHub repository.
You create a new repository in Azure DevOps.
You need to recommend a procedure to clone the repository from GitHub to Azure DevOps.
What should you recommend?

A. Create a webhook.
B. Create a service connection for GitHub.
C. From Import a Git repository, click Import
D. Create a pull request.
E. Create a personal access token in Azure DevOps.

Answer: C

Explanation:

Question: 102

DRAG DROP
You are implementing a package management solution for a Node.js application by using Azure Artifacts.
You need to configure thedevelopment environment to connect to the package repository. The solution must minimize the likelihood that credentials will be leaked.
Which file should you use to configure each connection? To answer, drag the appropriate files to the correct connections.
Each file 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.

Files

- The npmtc file in the preyed
- The npmrc file in the user s home folder
- The Package json file in the project
- The Projectjson file in the project

Answer Area

jgistry information: File
Credentials: File

Answer:

Explanation:

Feed registry information:	The .npmrc file in the project
Credentials:	The npmrc file in the user's home folder

All Azure Artifacts feeds require authentication, so you'll need to store credentials for the feed before you can install or publish packages. npm uses .npmrc configuration files to store feed URLs and credentials. Azure DevOps Services recommends using two .npmrc files.

Feed registry information: The .npmrc file in the project
One .npmrc should live at the root of your git repo adjacent to your project's package.json. It shouldcontain a "registry" line for your feed and it should not contain credentials since it will be checked into git.

Credentials: The .npmrc file in the user's home folder
On your development machine, you will also have a .npmrc in $home for Linux or Mac systems or $env.HOME for win systems. This .npmrc should contain credentials for all of the registries that you need to connect to. The NPM client will look at your project's .npmrc, discover the registry, and fetch matching credentials from $home/.npmrc or $env.HOME/.npmrc.

Reference:
https://docs.microsoft.com/en-us/azure/devops/artifacts/npm/npmrc?view=azure- devops&tabs=windows

Question: 103

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. PDM
B. OWASPZAP
C. WhiteSource
D. Jenkins

Answer: C

Explanation:

Question: 104

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

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

Whichthree 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 coned order.

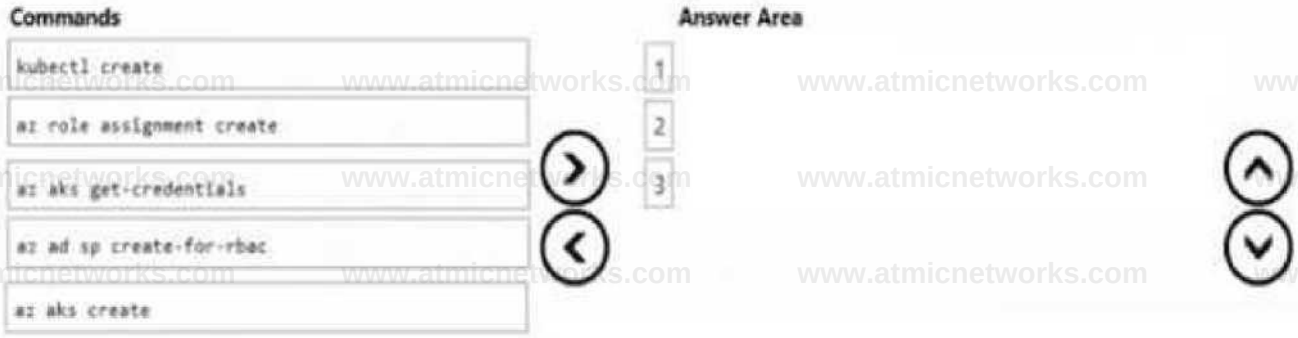

Answer:

Explanation:

```
az aks create

az ad sp create-for-rbac

kubectl create
```

Step 1 : az acr create

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

az acr create
--name <REGISTRY_NAME>
- -resource-group <RESOURCE_GROUP_NAME>
- -sku Basic

Step 2: az ad sp create-for-rbac
Once the ACR has been provisioned, you can either enable administrative access (which is okay for testing) or you create a Service Principal (sp) which will provide a client_id and a client_secret.

az ad sp create-for-rbac
--
scopes/subscriptions/<SUBSCRIPTION_ID>/resourcegroups/<RG_NAME>/providers/Microsoft.Contai
nerRegistry/registries/<REGISTRY_NAME>

- -role Contributor
- -name <SERVICE_PRINCIPAL_NAME>

Step 3: kubectl create

Create a new Kubernetes Secret.

kubectl create secretdocker-registry <SECRET_NAME>
- -docker-server <REGISTRY_NAME>.azurecr.io
- -docker-email <YOUR_MAIL>
- -docker-username=<SERVICE_PRINCIPAL_ID>
- -docker-password <YOUR_PASSWORD>

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

Question: 105

Your company has a project in Azure DevOps for a newapplication. The application will be deployed to several Azure virtual machines that run Windows Server 2016.

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

Answer: C

Explanation:

The Custom ScriptExtension 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.

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

Question: 106

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.

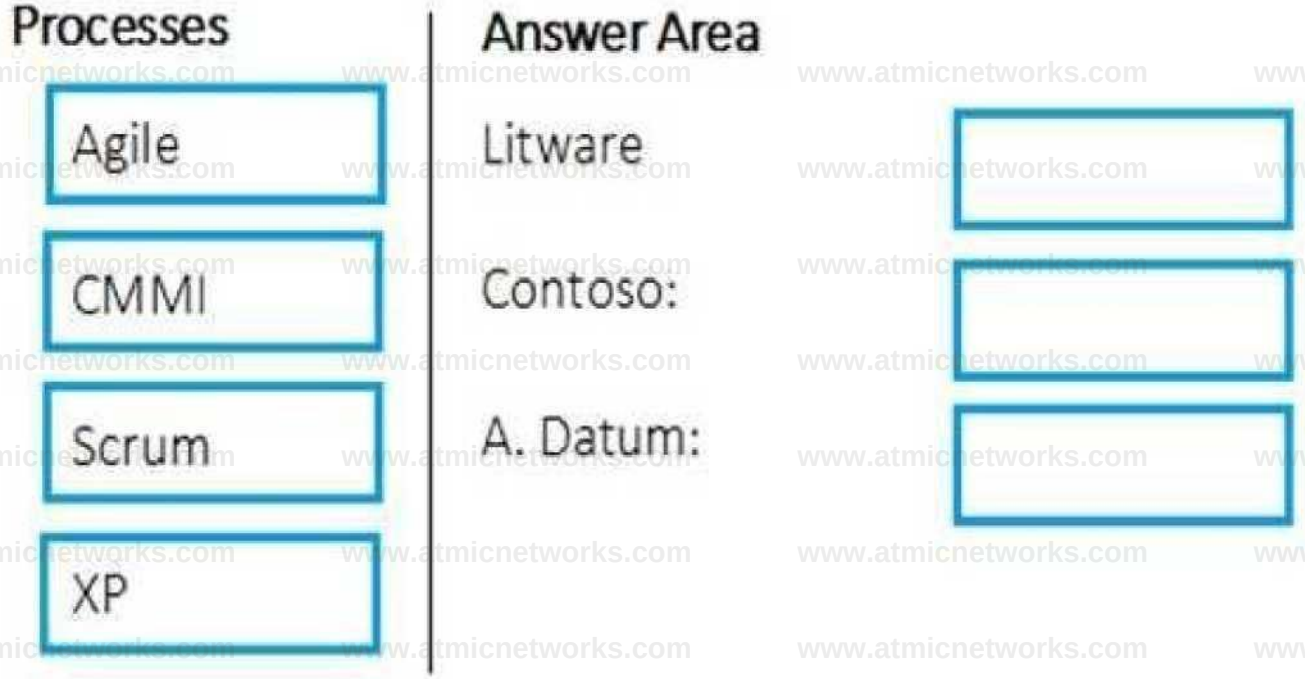

Answer:

Explanation:

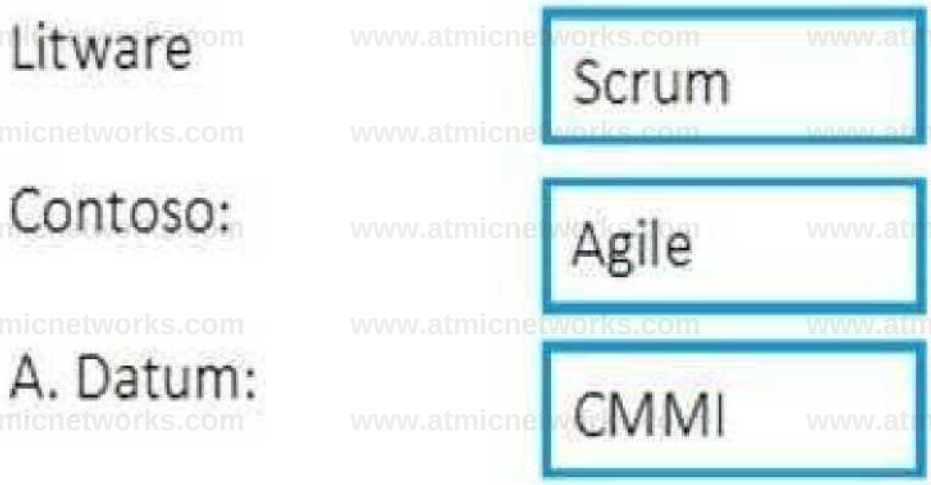

Box 1: Scrum
Choose Scrum when your team practices Scrum. This process works great if you want to track product backlog items (PBIs) and bugs on the Kanban board, or break PBIs and bugs down intotasks on the taskboard.

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.

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

Question: 107

Your development team is building a new web solution by using the Microsoft Visual Studio integrated development environment (IDE).
You need to make a custom package available to all the developers. The package must be managed centrally, and the latest version must be availablefor 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. Publish the package to a feed.
B. Create a new feed in AzureArtifacts.
C. Upload a package to a Git repository.
D. Add the package URL to the Environment settings in Visual Studio.
E. Add the package URL to the NuGet Package Manager settings in Visual Studio.

F. Create a Git repository in Azure Repos.

Answer: A,B,E

Explanation:

8: By using your custom NuGet package feed within your Azure DevOps (previously VSTS) instance, you'll be able to distribute your packages within your organization with ease.

Start by creating a new feed.

A: We can publish, pack and pushthe built project to our NuGet feed.

E: Consume your private NuGet Feed

Go back to the Packages area in Azure DevOps, select your feed and hit "Connect to feed". You'll see some instructions for your feed, but it's fairly simple to set up.

Just copy your package source URL, go to Visual Studio, open the NuGet Package Manager, go to its settings and add a new source. Choose a fancy name, insert the source URL. Done.

Search for your package in the NuGet Package Manager and it should appear there, ready for installation. Make sure to select the appropriate feed (or just all feeds) from the top right select box.

Reference:

https://medium.com/medialesson/get-started-with-private-nuget-feeds-in-azure-devops-8c7b5f022a68

Question: 108

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 AzureActive Directory, configure risky sign-ins.
C. In Azure DevOps, configure Security in Project Settings.
D. In Azure Active Directory, configure conditional access.

Answer: D

Explanation:

Conditional Access is a capability of Azure Active Directory. WithConditional 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: 109

You are automatingthe testing process for your company.
You need to automate UI testing of a web application.
Which framework should you use?

A. JaCoco
B. Playwright
C. Xamarin.UITest
D. Microsoft.CodeAnalysis

Answer: B

Explanation:

Performing user interface (UI) testingas part of the release pipeline is a great way of detecting unexpected changes, and need not be difficult. Selenium can be used to test your website during a continuous deployment release and test automation.
Reference:
https://docs.microsoft.com/en-us/azure/devops/pipelines/test/continuous-test-selenium?view=azure-devops

Question: 110

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 solutionmust 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.

Answer: A

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
https://docs.microsoft.com/en-us/azure/devops/pipelines/library/variable-groups?view=azure-devops&tabs=yaml#link-secrets-from-an-azure-key-vault

Question: 111

DRAG DROP

You are configuring Azure Pipelines for three projectsin Azure DevOps as shown in the following table.

Project name	Project Details
P reject 1	The project team provides preconfigured YAML files that it wants to use to manage future pipeline configuration changes.
Project?	The sensitivity of the project requires that the source code be hosted on the managed Windows server on your company's network.
Projects	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, morethan 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.

Version Control Systems

- Assembla Subversion
- Bitbucket Cloud
- Git in Azure Repos
- GitHub Enterprise

Answer Area

Project1: [ ]

Project2: [ ]

Project3: [ ]

Answer:

Explanation:

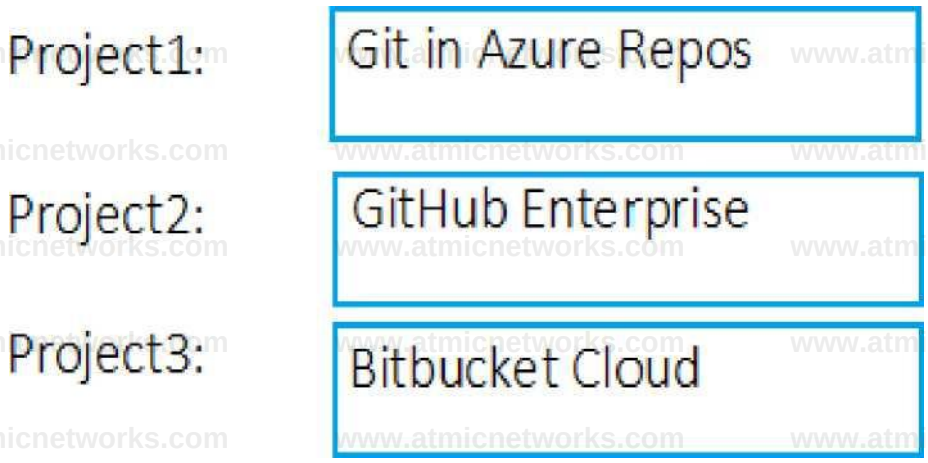

Project1:Git in Azure Repos

Project2: Github Enterprise
GitHub Enterprise is the on-premises version of GitHub.com. GitHub Enterprise includes the same great set of features as GitHub.com but packaged for running on your organization's local network. All repository data is stored on machines that you control, and access is integrated with your organization's authentication system(LDAP, SAML, or CAS).

Project3: Bitbucket cloud
One downside, however, is that Bitubucket does not include support for SVN but this can be easily amended migrating the SVN repos to Git with tools such as SVN Mirror for Bitbucket .

Note: SVN is a centralized version control system.

Reference:
https://www.azuredevopslabs.com/labs/azuredevops/yaml/

https://enterprise.github.com/faq

Question: 112

Your team uses an agile development approach.

You need to recommend a branching strategy for the team's Git repository. The strategymust meet the following requirements.

Provide the ability to work on multiple independent tasks in parallel.

Ensure that checked-in code remains in a releasable state always.

Ensure that new features can be abandoned at any time.

Encourage experimentation.

What should you recommend?

A. a single long-running branch
B. multiple long-running branches
C. a single fork per team member
D. a single-running branch with multiple short-lived topic branches

Answer: D

Explanation:

Topic branches, however, are usefulin projects of any size. A topic branch is a short-lived branch that you create and use for a single particular feature or related work. This is something you've likely never done with a VCS before because it's generally too expensive to create and mergebranches. But in Git it's common to create, work on, merge, and delete branches several times a day. Reference:
https://git-scm.com/book/en/v2/Git-Branching-Branching-Workflows

Question: 113

Your company has a project in Azure DevOps for a new webapplication.

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 taskto the pipeline.
B. Add a PowerShell task to the pipeline and run Set-AzureKeyVaultSecret.
C. Add a Azurre Key Vault task to the pipeline.
D. Add Azure Key Vault references to Azure Resource Manger templates.

Answer: B

Explanation:

Azure Key Vaultprovides 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: 114

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.

Whichthree 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.

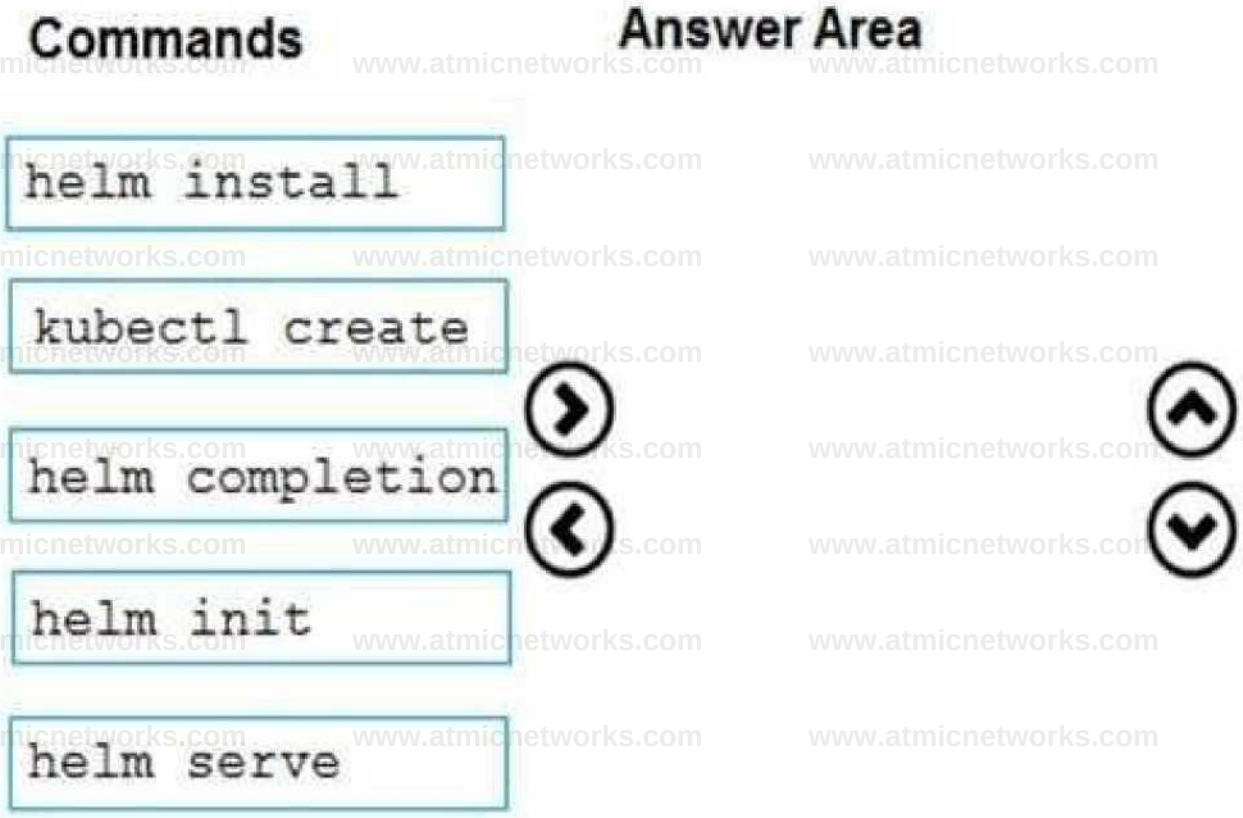

Answer:

Explanation:

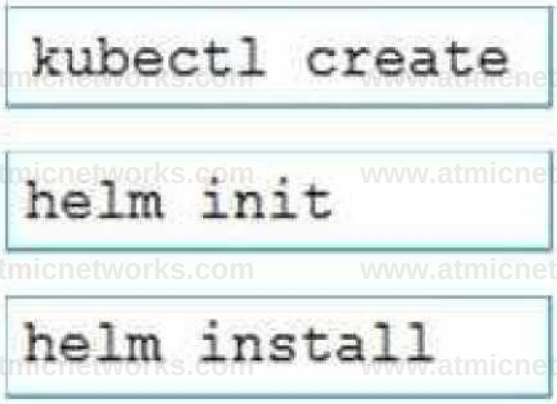

Step 1: Kubectl create

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.

Reference:

https://docs.microsoft.com/en-us/azure/aks/kubernetes-helm

https://docs.helm.sh/using_helm/#tiller-namespaces-and-rbac

Question: 115

You have a project in Azure DevOps. You have an Azure Resource Group deployment project in Microsoft Visual Studio that ischecked 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

Answer: C

Explanation:

There are two different ways to deploy templates to Azure DevOps Services. Both methods provide the same

results, so choosethe 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 deploysthe 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.

S:

https://docs.microsoft.com/en-us/azure/vs-azure-tools-resource-groups-ci-in-vsts

Question: 116

Your company isbuilding 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

Answer: B

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: 117

DRAG DROP

You are developing a fullMicrosoft .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.

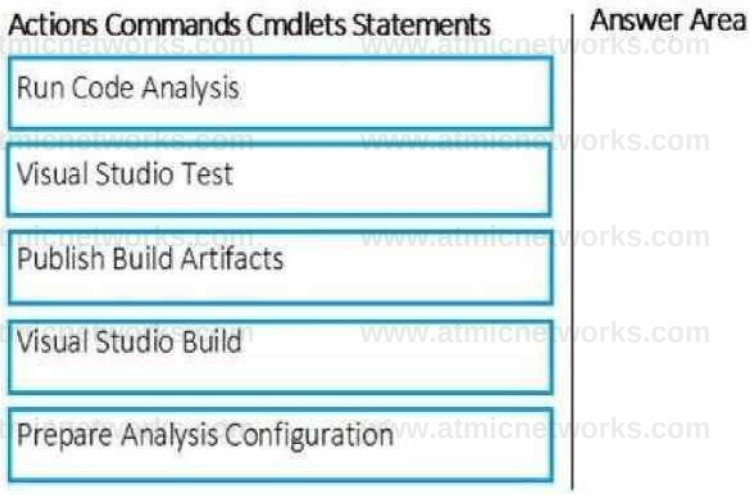

Answer:

Explanation:

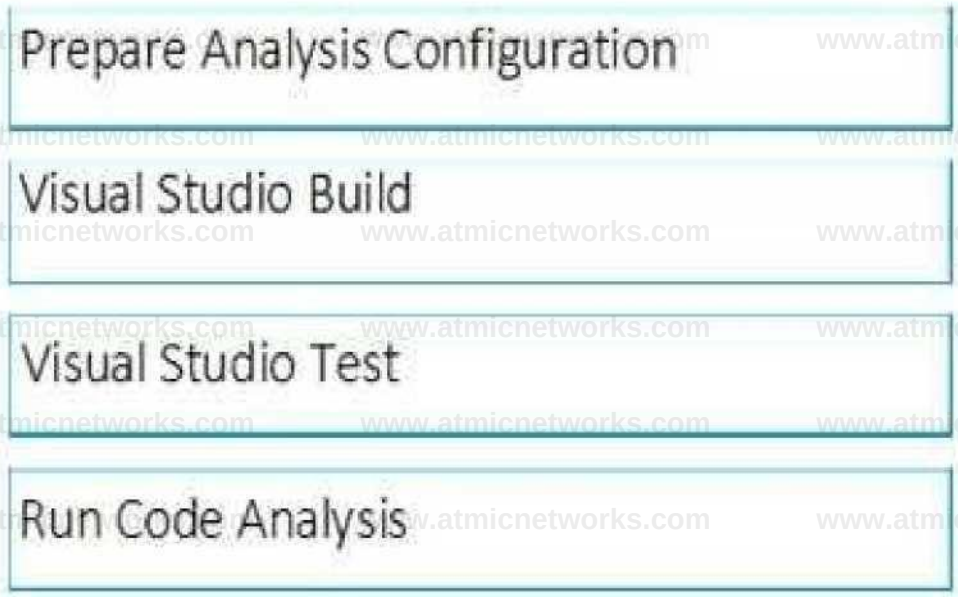

Step 1: Prepare Analysis Configuration

Prepare AnalysisConfiguration 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:

s:https://docs.sonarqube.org/display/SCAN/Analyzing+with+SonarQube+Extension+for+VSTS-TFS

Question: 118

You have an Azure DevOps organization named Contoso and an Azure DevOps project namedProject1.

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

Answer: A,E

Explanation:

https://github.com/microsoft/azure-pipelines-image-generation/blob/d80f81d6c98f8ce2c74b034309bb774ea8d31cfb/images/win/Vs2015-Server2012R2-Readme.md
https://github.com/actions/virtual-environments/blob/master/images/win/Windows2016-Readme.md

Question: 119

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 AzureDevOps 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.

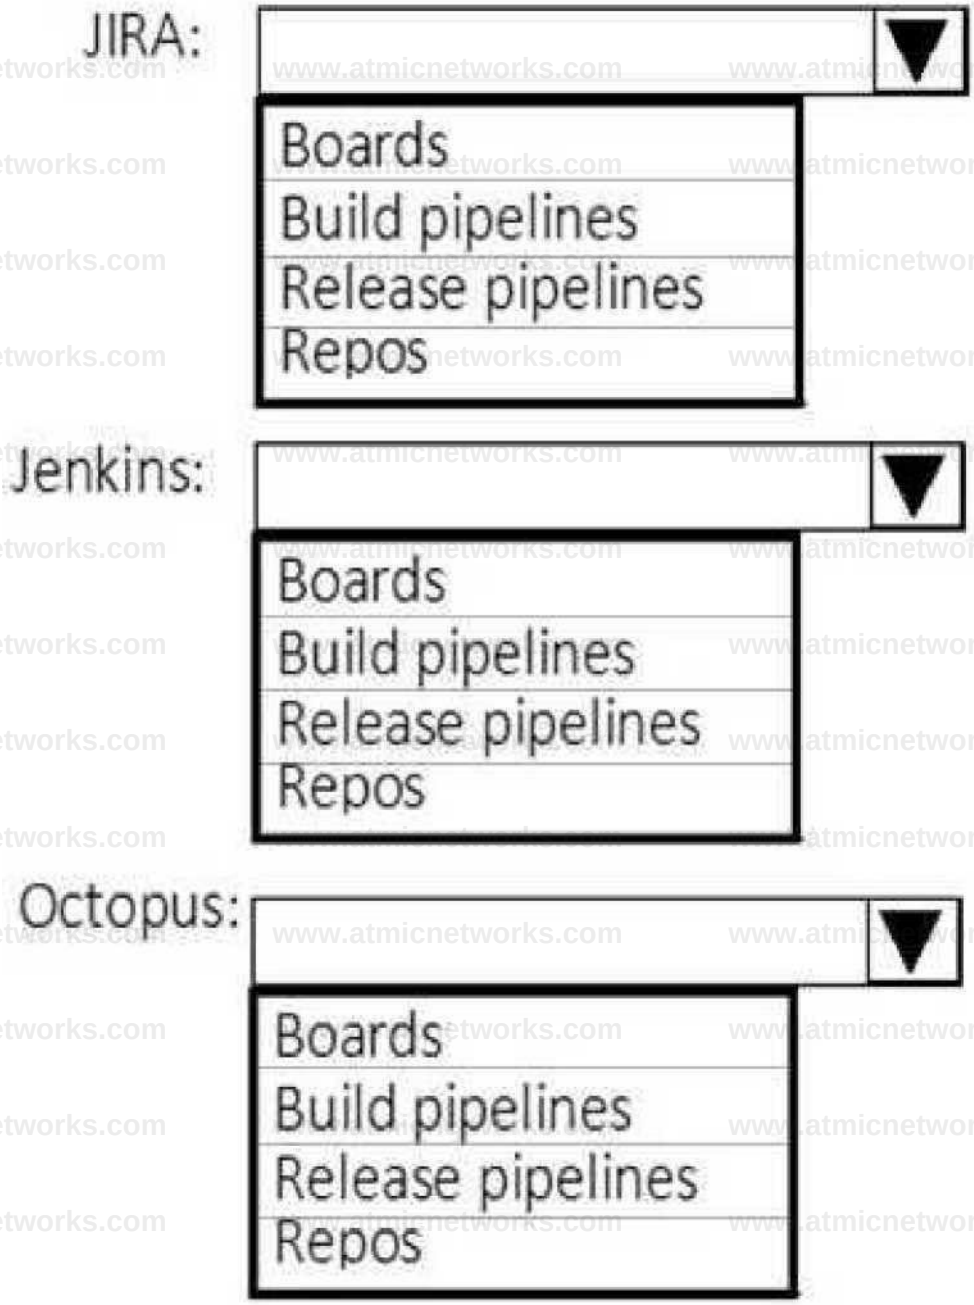

Answer:

Explanation:

JIRA- Board
Jenkins- Build Pipelines
Octopus- Release pipelines

Question: 120

Note: This question is part of a series of questions that present the samescenario. 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 thefollowing 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 single standalone template that willdeploy all the resources.

Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

Use two templates, one for each resource group, and link the templates.

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

Question: 121

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.

```
"policyRule": (
  "if": {
    "allOf": F
```

{

"field": "type" "equals":

"MicrosoftResources/deploymenTs"
"MicrosoftResources/subscriptions"________ _____
"MicrosoftResources/subscriptions/resourceGroups"

"not": {

"field": "tags['organization']",

"equals": "Contoso"

"then": {

"effect":

"details": [“Deny",

"Deploy! fNotExists".

{ --

"field": "tags!'organization']",

"value": "Contoso"

Answer:

Explanation:

"policyRule": (

"if": {

"a11Of": [

{

"field": "type", "equals":

{ "MicrosoftResources/deploym _____
"MicrosoftResources/sub^
"MicrosoftResources/subscriptions/resourceGroups"

"not": {

"field": "tags('organization']",

equals": "Contoso

"then

"effect":

"details"

"field": "tags[’organization']

"value": "Contoso"

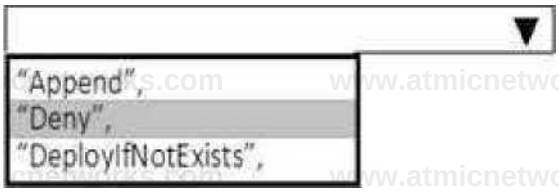

Box 1: " Microsoft.Resources/subscriptions/resourceGroups"

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"
}
}
}
}
```

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

Question: 122

DRAG DROP

You are defining release strategies for two applications as shown in the following table.

Application name	Goal
Appl	Failure of Appl has a major impact on your company. You need a small group of users, who opted in to a testing Appl, 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 releasestrategy 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 scrollto view content.

NOTE:Each correct selection is worth one point.

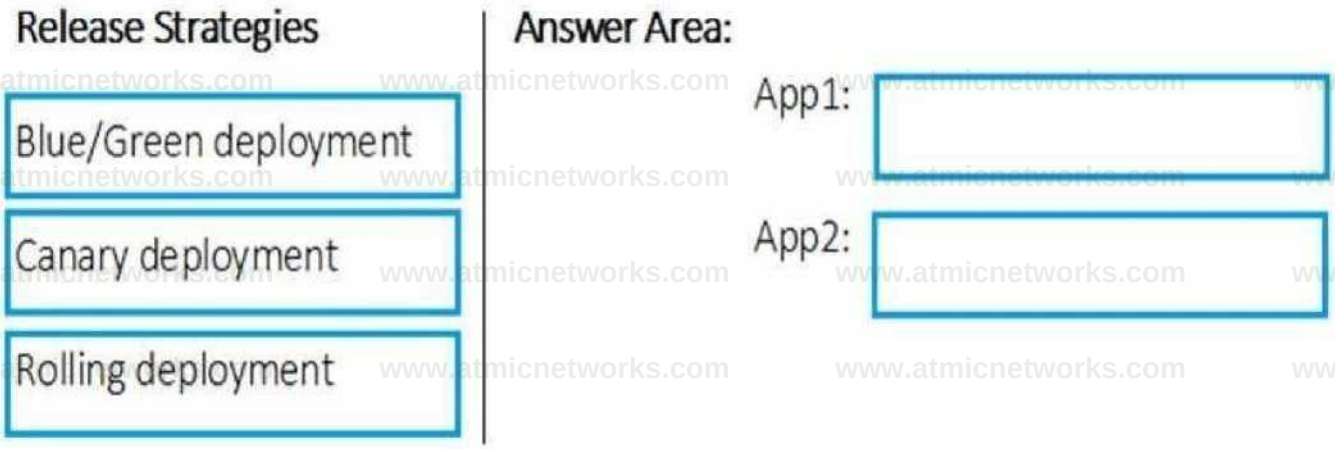

Answer:

Explanation:

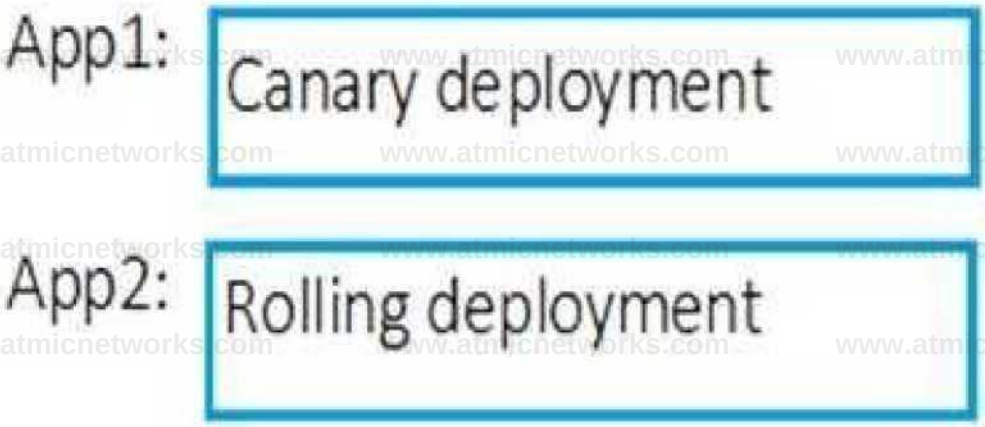

App1: Canary deployment

With canary deployment, you deploy a new application code in a smallpart of the production infrastructure. Once the application is signed off for release, only a few users are routed to it. This minimizes any impact.

With no errors reported, the new version can gradually roll out to the rest of the infrastructure.

App2:Rolling deployment:

In a rolling deployment, an application's new version gradually replaces the old one. The actual deployment happens over a period of time. During that time, new and old versions will coexist without affecting functionality or user experience. This process makes it easier to roll back any new component incompatible with the old components.

Incorrect Answers:

A: blue/green deployment is a change management strategy for releasing software code. Blue/green deployments, which may also be referred to as A/B deployments require two identical hardware environments that are configured exactly the same way. While one environment is active and serving

end users, the other environment remains idle.Blue/green deployments areoften used for consumerfacing applications and

applications with critical uptime requirements. New code is released to the inactive environment, where it is thoroughly tested. Once the code has been vetted, the team makes the idle environment active, typically by adjusting a router configuration to redirect application program traffic. The process reverses when the next software iteration is ready for release. Reference:
https://dev.to/mostlyjason/intro-to-deployment-strategies-blue-green-canary-and-more-3a3

Question: 123

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 nothave 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 requiresthat 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

Answer: B

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: 124

DRAG DROP

Youneed to find and isolate shared code. The shared code will be maintained in a series of packages.

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.

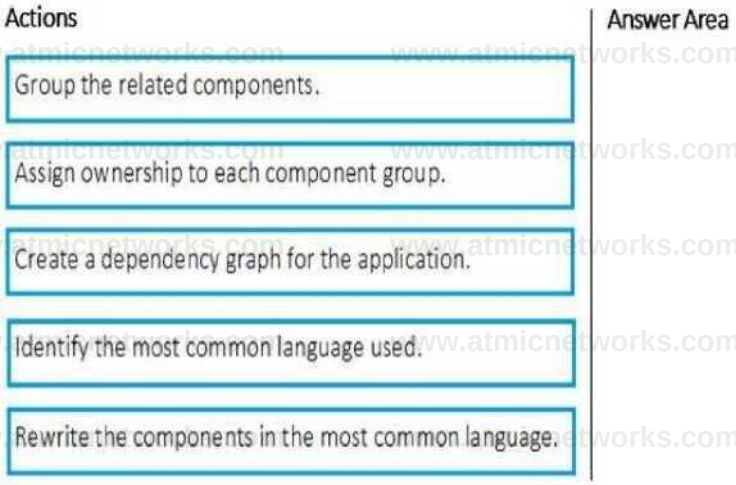

Answer:

Explanation:

Create a dependency graph for the application

Group the related components

Assign ownership to each component group.

Step 1: Create a dependency graph for the application

By linking work items and other objects, you can track related work, dependencies, and changes made over time. All links are defined with a specific link type. For example, you can use Parent/Child links to link work items to support a hierarchical tree structure. Whereas, the Commit and Branch link types support links between work items and commits and branches, respectively.

Step 2: Group the related components.

Packages enable you to share code across your organization: you can compose a large product,

develop multiple products based on a common shared framework, or create and share reusable components and libraries.

Step 3: Assign ownership to each component graph

Reference:
https://docs.microsoft.com/en-us/azure/devops/boards/queries/link-work-items-support-traceability?view=azure-devops&tabs=new-web-form

https://docs.microsoft.com/en-us/visualstudio/releasenotes/tfs2017-relnotes

Question: 125

You have an application that consists of several Azure App Service web apps and Azure functions.
You need to access the security of the web apps and the functions.
Which Azure features 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 inAzure Security Center

Answer: D

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 aset 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: 126

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 VisualStudio 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 organizationin App Center.

Answer: B

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 thedistribution 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 passwordfor 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.

s:

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

Question: 127

DRAG DROP

Your company wants to use Azure Application Insights to understand how userbehaviors affect an application.

Which application Insights tool should you use to analyze each behavior? To answer, drag the appropriate tools to the correct behaviors. Each tool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE:Each correct selection is worth one point.

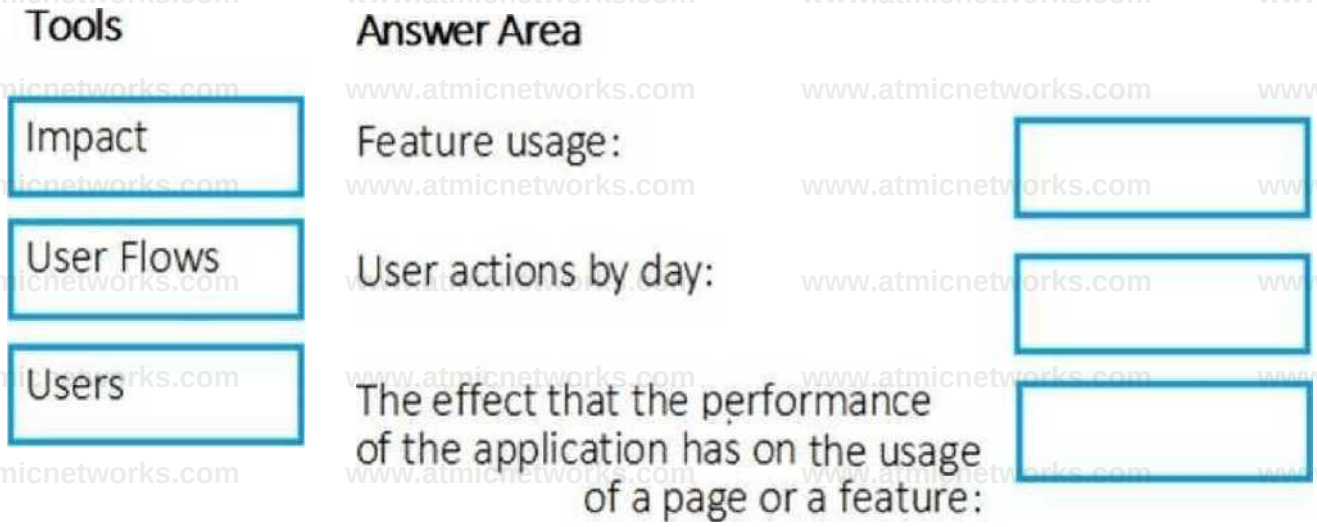

Answer:

Explanation:

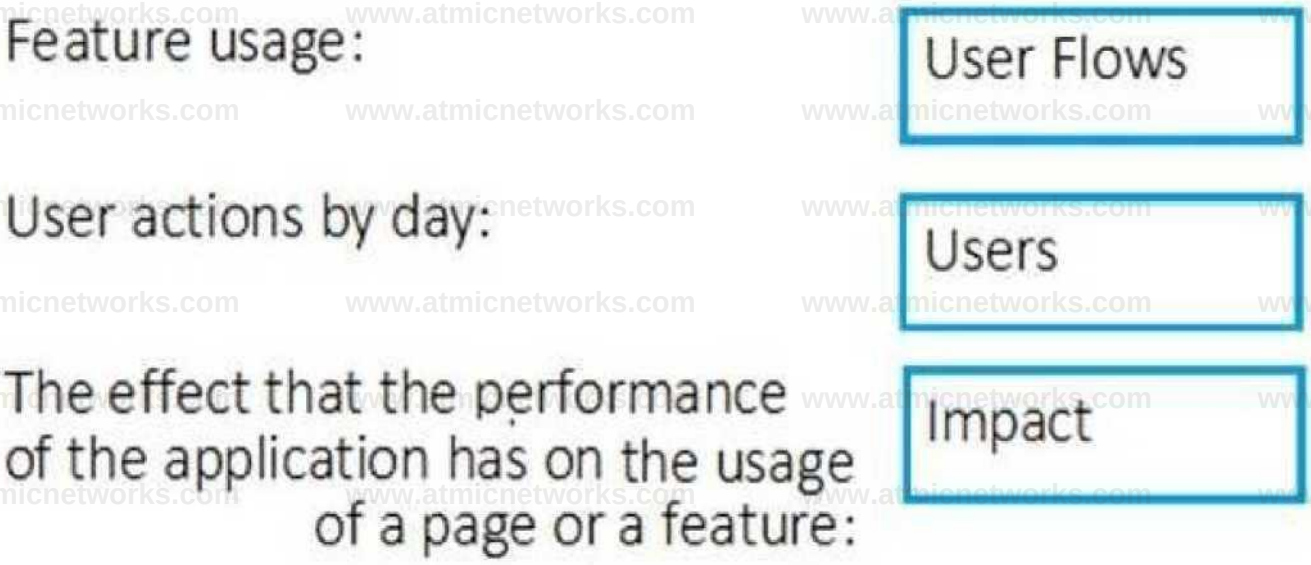

Box 1: User Flows

The User Flows tool visualizes how users navigate between the pages and features of your site. It's great for answering questions like:

How do users navigate away from a page on your site?

What do users click ona page on your site?

Where are the places that users churn most from your site?

Are there places where users repeat the same action over and over?

Box 2: Users

Box 3: Impact

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-flows

Question: 128

You have a GitHub repository.
You create a new repository inAzure DevOps.
You need to recommend a procedure to clone the repository from GitHub to Azure DevOps.
What should you recommend?

A. Create a pull request.
B. Create a webhook.
C. Create a service connection for GitHub.
D. From Import a Git repository, clickImport.
E. Create a personal access token in Azure DevOps.

Answer: D

Explanation:

You can import an existing Git repo from GitHub, Bitbucket, GitLab, or other location into a new or empty existing repo in your project in Azure DevOps.

Import into a newrepo
Select Repos, Files.
From the repo drop-down, select Import repository.
If the source repo is publicly available, just enter the clone URL of the source repository and a name for your new Git repository.

Reference:
https://docs.microsoft.com/en-us/azure/devops/repos/git/import-git-repository?view=azure-devops

Question: 129

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 thebuild 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 storedin a Git 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

Answer: A

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: 130

Note: This question is partof 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 solutionmust 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 Git 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

Answer: B

Explanation:

Instead use Octopus Tentacle.

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

Question: 131

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 Git repository in Azure DevOps.

Solution: Configure the build pipeline to use a Hosted VS 2017 agent pool. Include the Java Tool Installer task in the buildpipeline.

Does this meet the goal?

A. Yes
C. No

Answer: B

Explanation:

Instead use Octopus Tentacle.

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

Question: 132

You are designing the development process for your company.

You need to recommend a solution for continuous inspection of the company's code base to locate common code patterns that are known to be problematic.
What should you include in the recommendation?

A. Microsoft Visual Studio test plans
B. Gradle wrapper scripts
C. SonarCloud analysis
D. the JavaScript task runner

Answer: C

Explanation:

SonarCloud is a cloud service offered by SonarSource and based on SonarQube. SonarQube is awidely adopted open source platform to inspect continuously the quality of source code and detect bugs, vulnerabilities and code smells in more than 20 different languages.
Note: The SonarCloud Azure DevOps extension brings everything you need to have your projects analyzed on SonarCloud very quickly.
Incorrect Answers:
A: Test plans are used to group together test suites and individual test cases. This includes static test suites, requirement-based suites, and query-based suites.
Reference:
https://docs.travis-ci.com/user/sonarcloud/
https://sonarcloud.io/documentation/integrations/vsts/

Question: 133

Your company is building a new solution in Java.
The companycurrently 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

Answer: B

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: 134

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.

Configurations

an Azure Key Vault access policy

a personal access token (PAT)

RBAC

Answer Area

Restrict access 10 deles die key vaut

Restrict access to the secrets in Key Vault by using

Answer:

Explanation:

Resort access to deles lie key vaut | RBAC

Restrict access to the secrets in Key Vaut by using RBAC

Box 1: RBAC
Management plane access control uses RBAC.
The management plane consists of operations that affect the key vault itself, such as:

Creating or deleting a key vault.
Getting a list of vaults in a subscription.
Retrieving Key Vault properties (such as SKU and tags).
Setting Key Vault access policies that control user and application access to keys and secrets.

Box 2: RBAC

Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-manager-tutorial-use- key-vault

Question: 135

Note: This question is part of a series of questions that present the samescenario. 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 a Kubernetes cluster on-premises. You deploy a Helm agent to the cluster. You add a Download Build Artifacts task to the deployment pipeline.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead you should deploy an Azureself-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).

Note 2: As we [Microsoft] are launching this new experience in preview, we are currently optimizing it for Azure Kubernetes Service (AKS) and Azure Container Registry (ACR). Other Kubernetes clusters, for example running on-premises or in other clouds, as well as other container registries, can be used, but require setting up a Service Account and connection manually.

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

Question: 136

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.

Afteryou 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 a Docker build to an on-premises server. You add a Download Build Artifacts task to the deployment pipeline.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead you should deploy anAzure 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: 137

This question is part of a series ofquestions 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

Answer: A

Explanation:

To build your code ordeploy 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: 138

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

Answer: C

Explanation:

Scenarios and use cases for gates include:
Qualityvalidation. 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 thatyou 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: 139

DRAG DROP
You are creating a NuGet package.
You plan to distribute the package to your development team privately.
You needto share the package and test that the package can be consumed.
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.

Actions

- Create a new Azure Artifacts feed.
- Configure a self-hosted agent.
- Publish a package.
- Install a package.
- Connect to an Azure Artifacts feed.

Answer Area

Answer:

Explanation:

Configure a self-hosted agent.

Create a new Azure Artifacts feed

Publish a package

Connect to an Azure Artifacts feed.

Step 1: Configure a self-hosted agent.
The build will run on a Microsoft hosted agent.

Step 2: Create a new Azure Artifacts feed
Microsoft offers an official extensionfor publishing and managing your private NuGet feeds.

Step 3: Publish the package.
Publish, pack and push the built project to your NuGet feed.

Step 4: Connect to an Azure Artifacts feed.
With the package now available, you can point Visual Studio to thefeed, and download the newly published package

Reference:
https://medium.com/@dan.cokely/creating-nuget-packages-in-azure-devops-with-azure-pipelines- and-yaml-d6fa30f0f15e

Question: 140

HOTSPOT

You have an Azure DevOpsproject that contains a build pipeline. The build pipeline uses approximately 50 open source libraries.

You need to ensure that the project can be scanned for known security vulnerabilities in the open source libraries.

What should you do? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

Answer:

Explanation:

Object to create: _____

A build task
A deployment task An artifacts repository

Service to use:

WhiteSource Bolt
Bamboo
CMake
Chef

Box 1: ABuild task

Trigger a build

You have a Java code provisioned by the Azure DevOps demo generator. You will use WhiteSource Bolt extension to check the vulnerable components present in this code.

Go to Builds section under Pipelines tab, select the build definition WhiteSourceBolt and click on Queue to trigger a build.

To view the build in progress status, click on ellipsis and select View build results.

Box 2: WhiteSource Bolt

WhiteSource is the leader in continuous open source software security andcompliance 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://www.azuredevopslabs.com/labs/vstsextend/whitesource/

Question: 141

You have an Azure DevOps project named Project1 and an Azure subscriptionnamed Sub1. Sub1 contains an Azure virtual machine scale set named VMSS1. VMSS1 hosts a web application named WebApp1.

WebApp1 uses stateful sessions.

The WebApp1 installation is managed by using the Custom Script extension. The script resides in an AzureStorage account named sa1.

You plan to make a minor change to a UI element of WebApp1 and to gather user feedback about the change.

You need to implement limited user testing for the new version of WebApp1 on VMSS1.

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

NOTE:Each correct selection is worth one point.

A. Modify the load balancer settings of VMSS1.
B. Redeploy VMSS1.
C. Upload a custom script file to sa1.
D. Modify the Custom Script extension settings ofVMSS1.
E. Update the configuration of a virtual machine in VMSS1.

Answer: B,C,D

Explanation:

Question: 142

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. SourceGear Vault
B. Jenkins
C. Microsoft Visual SourceSafe
D. WhiteSource Bolt

Answer: D

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 securityand 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://www.azuredevopslabs.com/labs/vstsextend/whitesource/

Question: 143

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
Groupl	Application testers who are invited by email
1 Group2	Early release users who use unauthenticated public links
| Groups	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.

Answer Area

Answer:

Explanation:

Group 1: Private

Group2: Public

Group3: Shared

Box1: Private
In App Center, distribution groups are private by default. Only testers invited via email can access thereleases 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: 144

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

Answer: B

Explanation:

Use Azure SQLDatabase 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: 145

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 configurationof 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

Answer: B

Explanation:

Instead implement Continuous Assurance for the project.

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

Question: 146

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

Answer: B

Explanation:

Instead implement Continuous Assurance for the project.

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

Question: 147

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution thatmight 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

willnot 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

Answer: A

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 thechange).

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

Question: 148

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

Answer: D

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: 149

Note: This question is part of a series of questions that present the same scenario. Each question inthe 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 Git 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

Answer: A

Explanation:

Question: 150

You plan to use a NuGet package in a project in Azure DevOps. The NuGet package is in a feed that requires authentication.
You needto 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 principalD18912E1457D5D1DDCBD40AB3BF70D5D

Answer: B

Explanation:

The Azure Artifacts Credential Provider automates the acquisition of credentials needed torestore 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: 151

HOTSPOT
: 144
HOTSPOT

You needto 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 svstem.

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.

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

Answer:

Explanation:

Deployment 1:

azure key va u It-fl exvo 1 u me blobfuse-flexvol
kubernetesJo/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:

azure key va u It-f 1 ex vo 1 u me blobfuse-flexvol
kubernetesJo/azure-disk
kubernetes.io/azure-file
volume.beta.kubernetes.io/storage-provisioner

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.

Reference:
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: 152

You create a Microsoft ASP.NET Core application.

Youplan 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

Answer: B

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: 153

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.

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:

Explanation:

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

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.

Reference:
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: 154

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. B. the number of unit test failures
C. the percentage of unit test failures
D. the percentage of overall time spent on rework

Answer: D

Explanation:

Question: 155

DRAG DROP

You have an Azure Kubernetes Service (AKS) cluster.

You need to deploy an application tothe 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.

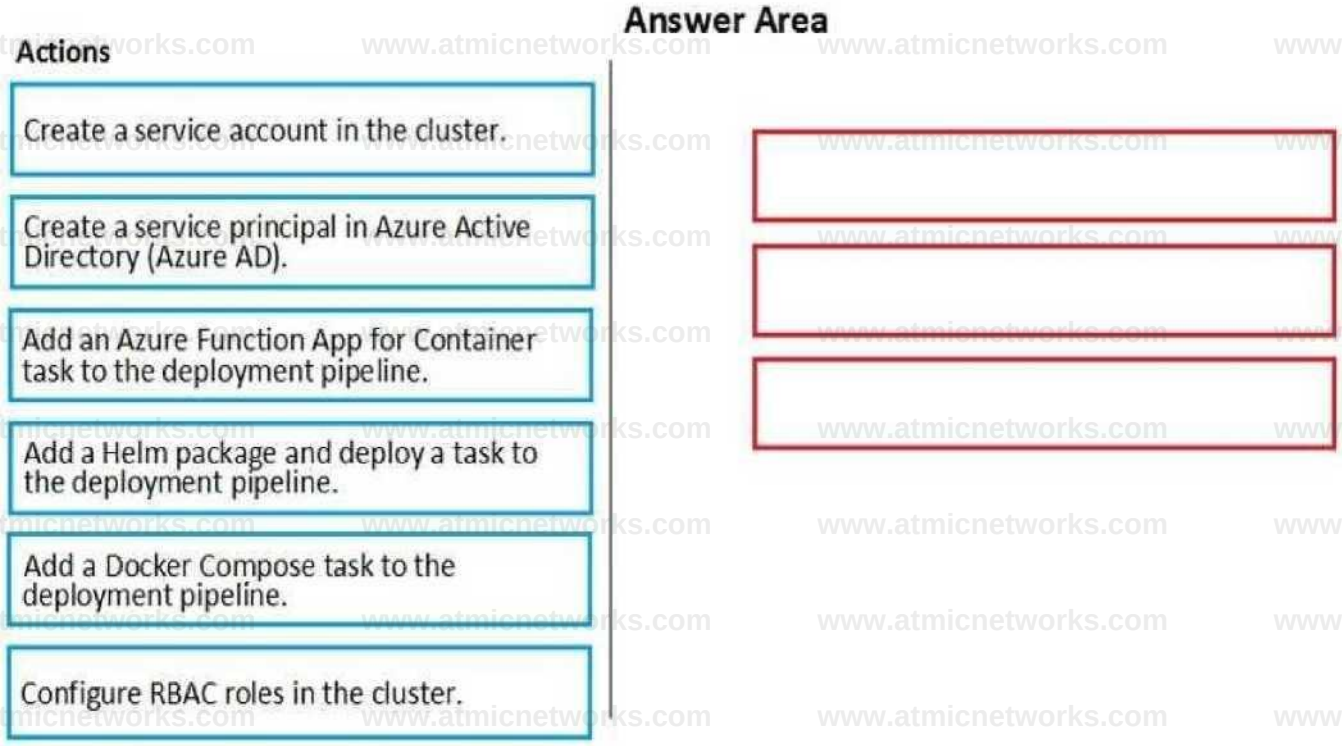

Answer:

Explanation:

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.

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:

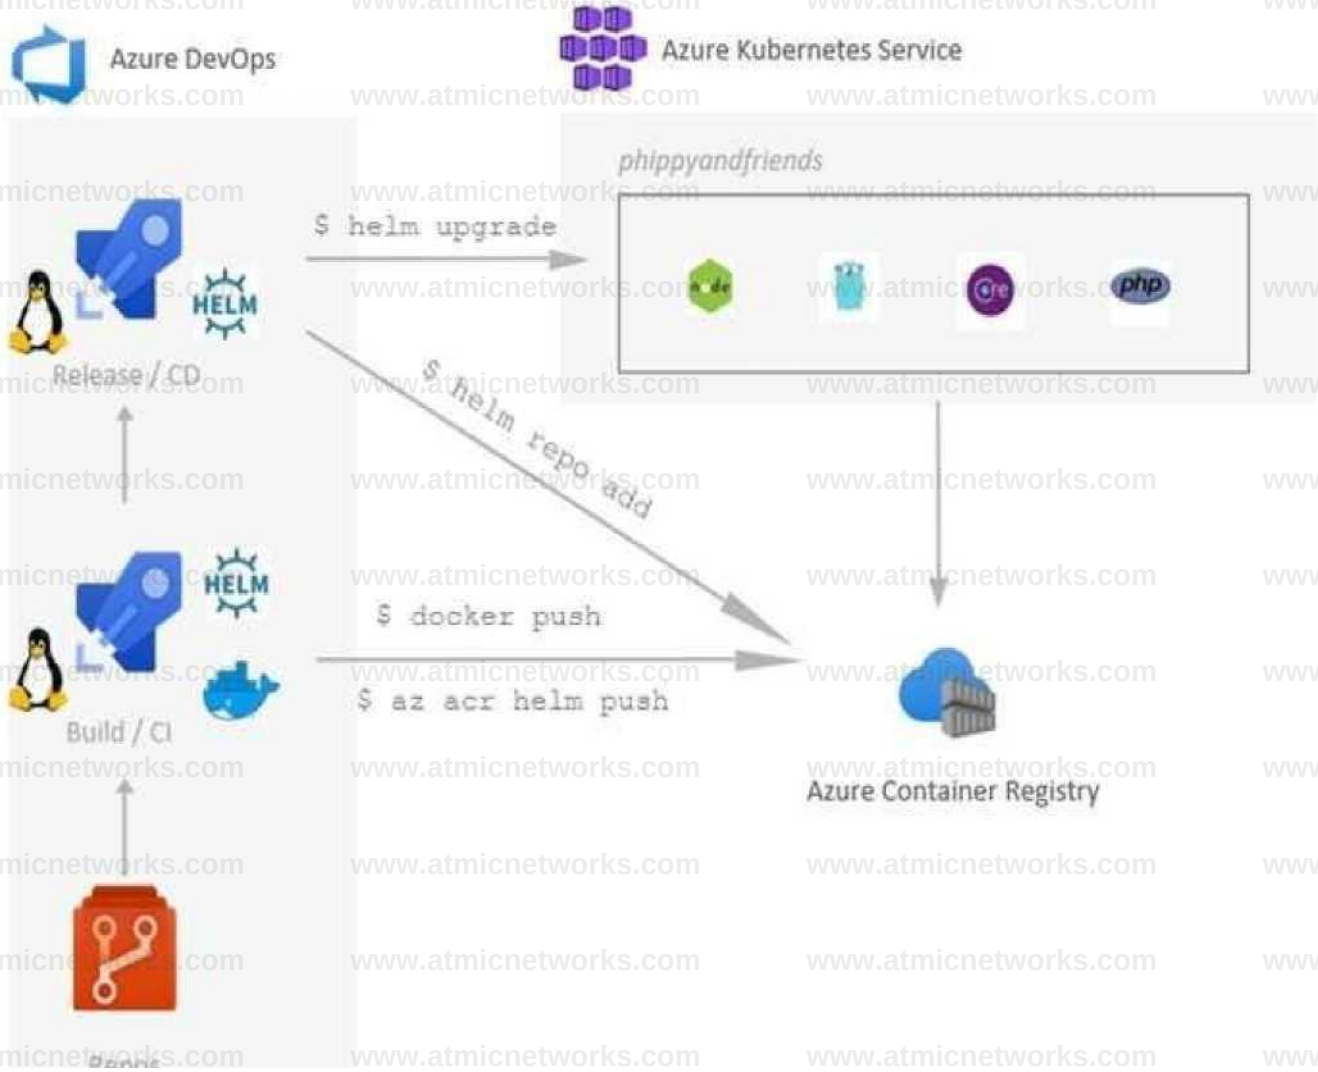

Step 3: Add a Docker Compose task tothe deployment pipeline.
Dockerfile file is a script leveraged by Docker, composed of various commands (instructions) and arguments listed successively to automatically perform actions on a base image in order to create a new Docker image by packaging theapp.

Reference:
https://cloudblogs.microsoft.com/opensource/2018/11/27/tutorial-azure-devops-setup-cicd-pipeline-kubernetes-docker-helm/

Question: 156

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

Answer: C

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 toautomatically 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: WhiteSourcewould 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: 157

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

Answer: B

Explanation:

Instead implement Continuous Assurance for the project.

Note: The Subscription Security health check features in AzSK contains aset 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-Continous-Assurance/Readme.html

Question: 158

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 devicecompliance 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.

Answer: E,F

Explanation:
Reference:
https://help.apple.com/xcode/mac/current/#/dev7ccaf4d3c

Question: 159

You have a privatedistribution 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. SelectRegister devices and sign my app.
B. Generate a new .p12 file for each device.
C. Create an active subscription in App Center Test.

D. Add the device owner to the collaborators group.

Answer: A

Explanation:

The following diagram displays the entire app re-signing flow in App Center.

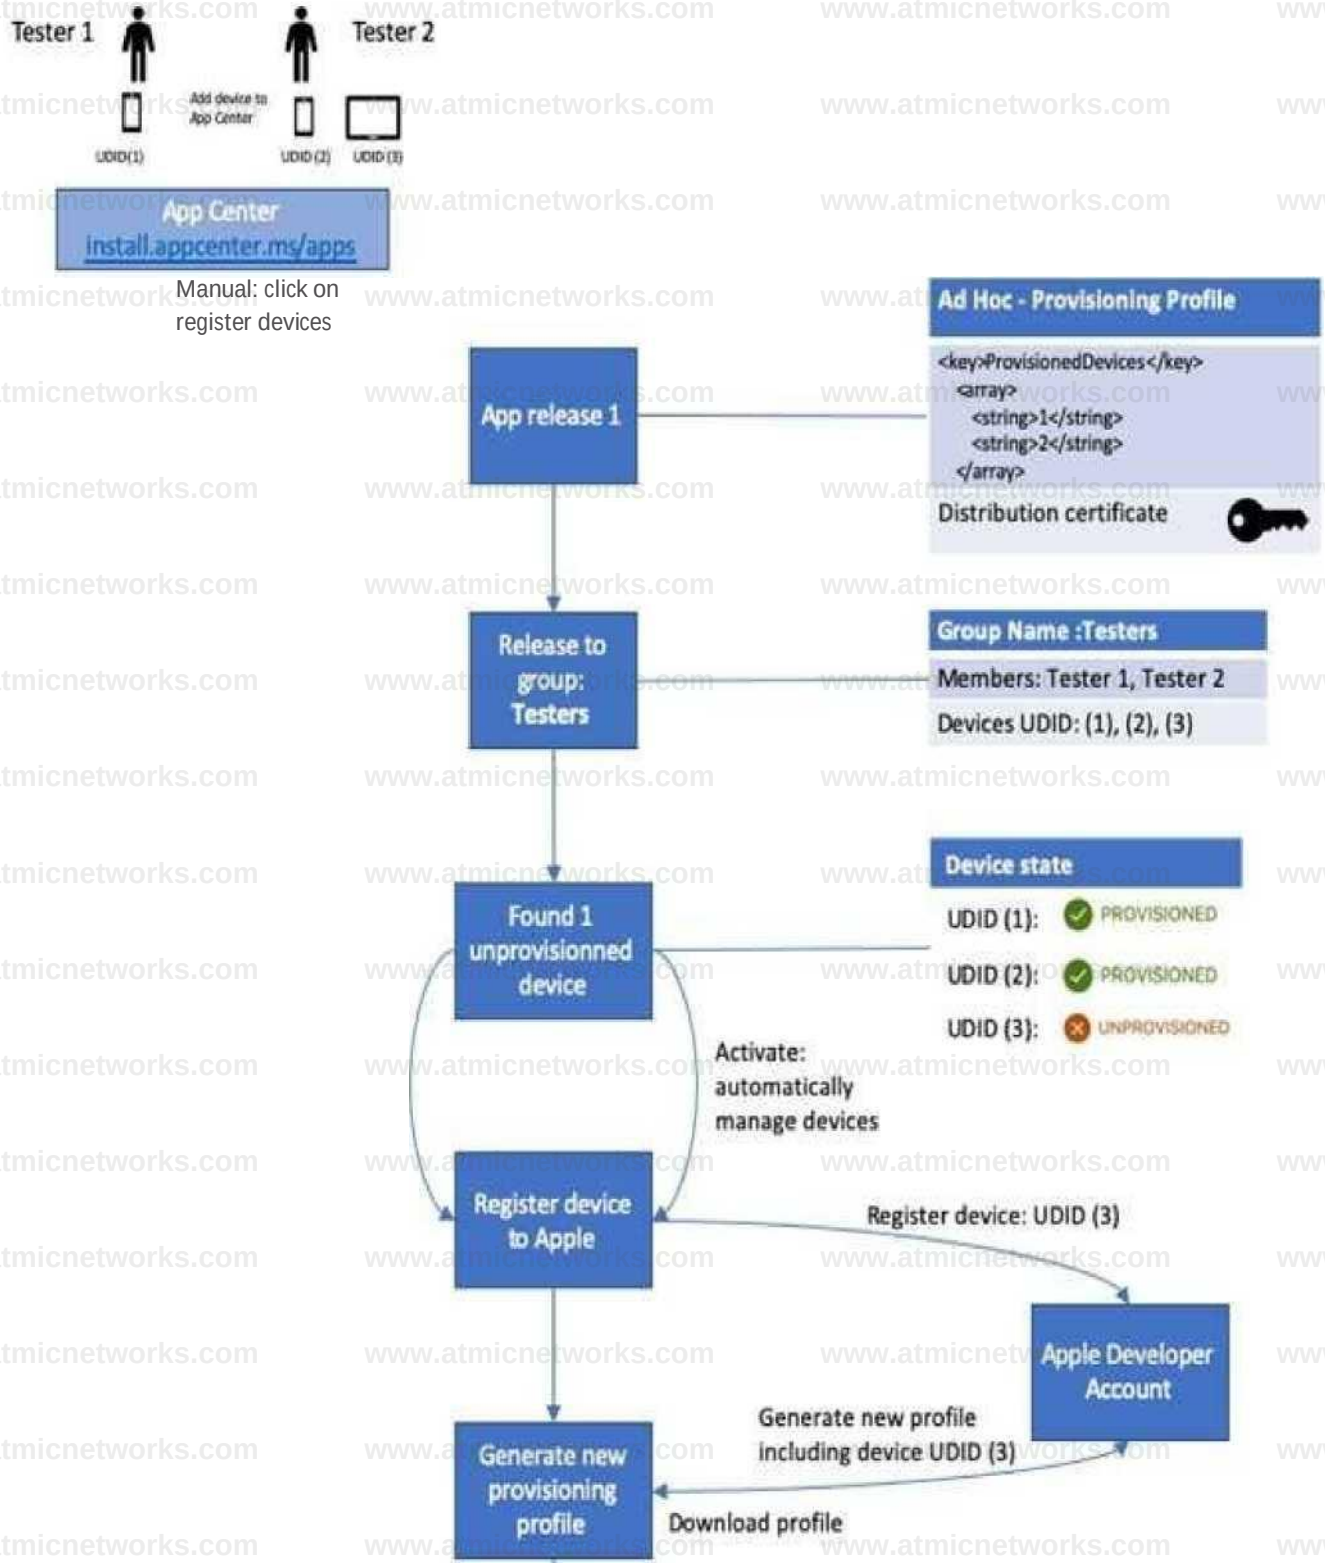
Tester 1
Tester 2
Add device to App Center
UDID (1)
UDID (2)
UDID (3)
App Center
install.appcenter.ms/apps
Manual: click on register devices
App release 1
Ad Hoc - Provisioning Profile
<key>ProvisionedDevices</key>
<array>
<string>1</string>
<string>2</string>
</array>
Distribution certificate
Release to group: Testers
Group Name :Testers
Members: Tester 1, Tester 2
Devices UDID: (1), (2), (3)
Found 1 unprovisionned device
Device state
UDID (1): PROVISIONED
UDID (2): PROVISIONED
UDID (3): UNPROVISIONED
Activate: automatically manage devices
Register device to Apple
Register device: UDID (3)
Apple Developer Account
Generate new provisioning profile
Generate new profile including device UDID (3)
Download profile

Incorrect Answers:

B: Only one .p12 file for the app, not one for each device.

Devicestate

Reference:
https://docs.microsoft.com/hu-hu/appcenter/distribution/auto-provisioning

Question: 160

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.

The lead developer at your company reports that adding new application features takes longer than expected due to a large accumulated technical debt.

You need to recommend changes to reduce the accumulated technical debt.

Solution: You recommend increasing the code duplication.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead reduce the code complexity.

Reference:
https://dzone.com/articles/fight-through-the-pain-how-to-deal-with-technical

Question: 161

Note: This question is part of a series of questions that presentthe 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.

The lead developer at your company reports that adding new application features takes longer than

expected due to a large accumulated technical debt.

You need to recommend changes to reduce the accumulated technical debt.

Solution: You recommend increasing the test coverage.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead reduce the code complexity.

Reference:
https://dzone.com/articles/fight-through-the-pain-how-to-deal-with-technical

Question: 162

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.

The lead developer at your company reports that adding new application features takes longer than expected due to a large accumulated technical debt.

You need to recommend changes to reduce the accumulated technical debt.

Solution: You recommend reducing the code complexity.

Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

Reference:
https://dzone.com/articles/fight-through-the-pain-how-to-deal-with-technical

Question: 163

Your company has 60 developerswho are assigned to four teams. Each team has 15 members.

The company uses an agile development methodology.

You need to structure the work of the development teams so that each team owns their respective work while working together to reach a common goal.

Which parts of the taxonomy should you enable the team to perform autonomously?

A. Features and Tasks
B. Initiatives and Epics
C. Epics and Features
D. Stories and Tasks

Answer: A

Explanation:

A feature typically represents a shippable component of software.

Features, examples:

Add view options to the new work hub

Add mobile shopping cart

Support text alerts

Refresh the web portal with new look and feel

User Stories and Tasks are used to track work. Teams can choose how they track bugs, either asrequirements or as tasks

Reference:

https://docs.microsoft.com/en-us/azure/devops/boards/backlogs/define-features-epics

https://docs.microsoft.com/en-us/azure/devops/boards/work-items/about-work-items

Question: 164

HOTSPOT

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

You plan to add a new functionality to App1.

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

Which branch lifetime and branch 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.

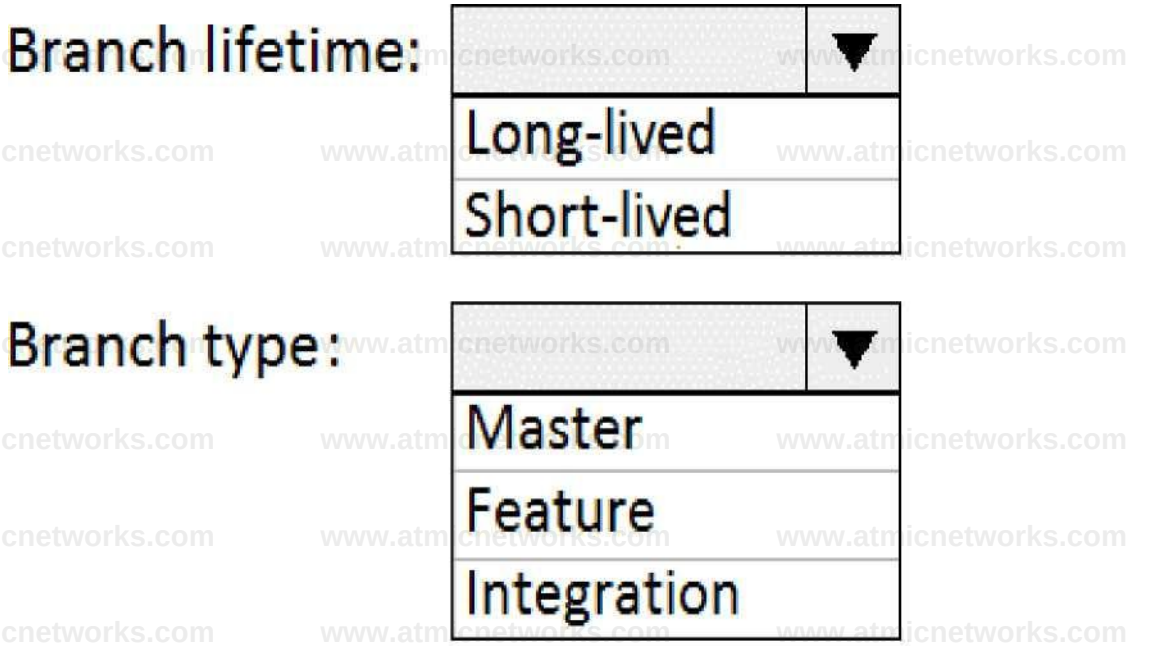

Answer:

Explanation:

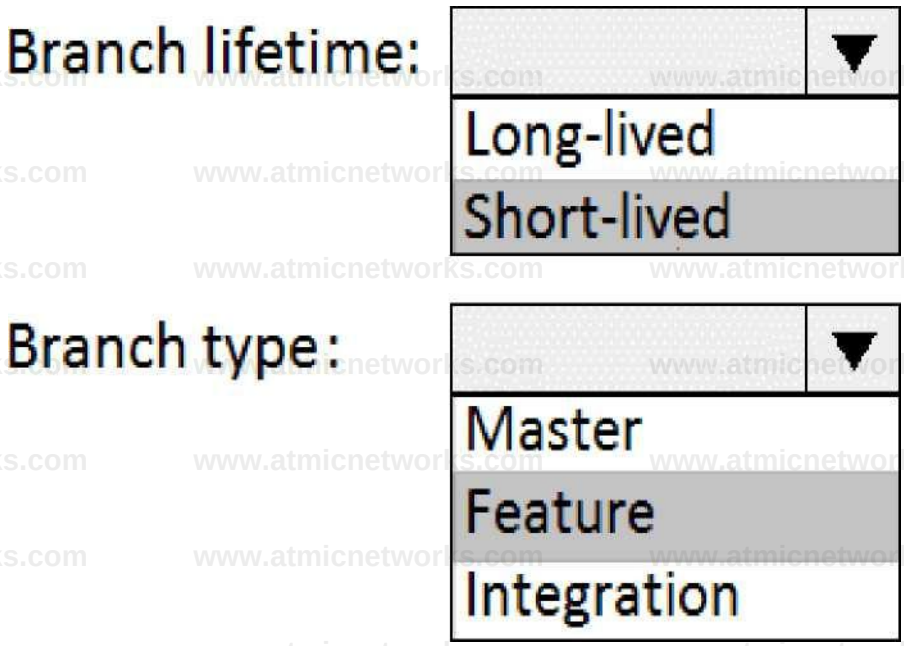

Branch lifetime: Short-lived

Branch type: Feature
Feature branches are used whendeveloping 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 branchwill be finished, it will always be merged back into the master branch.

Reference:
https://gist.github.com/digitaljhelms/4287848

Question: 165

You store source code in a Git 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

Answer: B

Explanation:

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

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

Question: 166

Your company creates a new Azure DevOps team.
D18912E1457D5D1DDCBD40AB3BF70D5D
You plan to use Azure DevOps for sprint planning.
You need to visualize the flow of your work by using an agile methodology.
Which Azure DevOps component shouldyou use?

A. Kanban boards
B. sprint planning
C. delivery plans
D. portfolio backlogs

Answer: A

Explanation:

Customizing Kanban boards

To maximize a team's ability to consistently deliver high quality software, Kanban emphasize two main practices. The first, visualize the flow of work, requires you to map your team's workflow stages and configure your Kanban board to match. Your Kanban board turns your backlog into an interactive signboard, providing a visual flow of work.

Reference:

https://azuredevopslabs.com/labs/azuredevops/agile/

Question: 167

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.

D18912E1457D5D1DDCBD40AB3BF70D5D

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.

Answer: B,C,E

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 theapp.

Select Add Access Policy.

Open Secret permissions and provide the app with Get and List permissions.

Select Select principal and select the registered app by name. Select the Select button.

Select OK.

Select Save.

Deploy the app.

Reference:

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

https://docs.microsoft.com/en-us/azure/key-vault/general/tutorial-net-virtual-machine

Question: 168

HOTSPOT

You manage build and release pipelines by using Azure DevOps. Your entire managed environment residesin 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.

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:

Explanation:

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

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 serviceswith an automatically managed identity in Azure AD. You can use the identity to authenticate to any service that supports Azure AD authentication, including Key Vault, without any credentials in your code.

Reference:

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

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

Question: 169

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 objectsto the answer area and arrange them in the correct order.

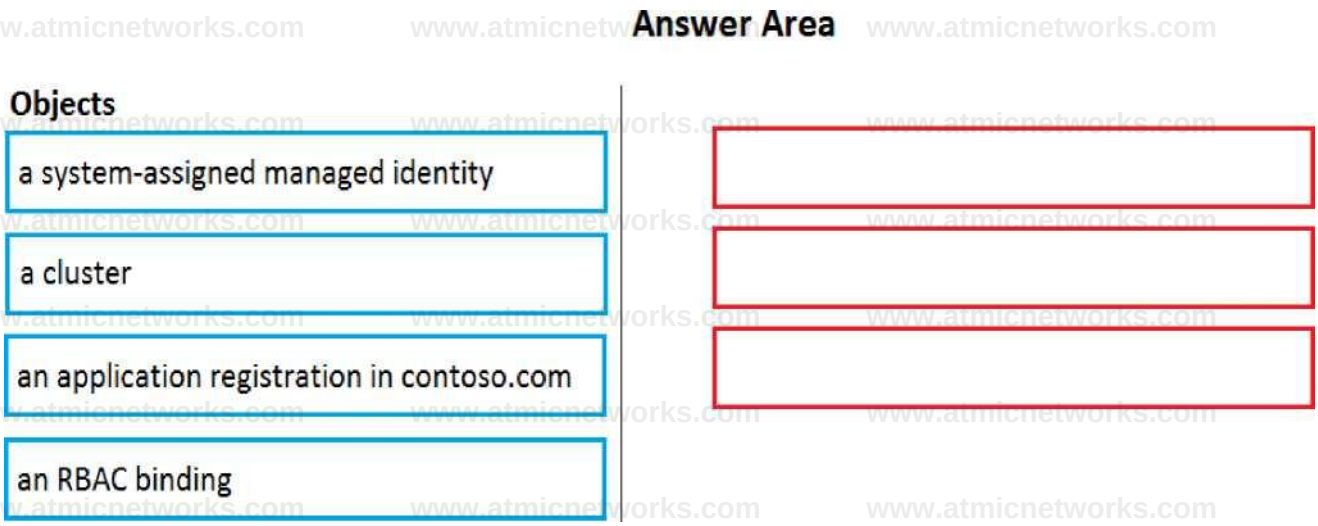

Answer:

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.

Sign in to the Azure portal.

In the search field at the top of the page, enterAzure Active Directory. Click Enter.

In the Manage menu, select Users.

In the name field, search for your account.

In the Name column, select the link to your account.

In the Identity section, copy the Object ID.

Step 3: a RBAC binding

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

Question: 170

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.

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.

The pipeline contains

one stage two
stages three
stage four
stages five
stages

Build vm contains

one job two

jobs three jobs

four jobs five

jobs

Answer:

Explanation:

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

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

Question: 171

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 scansproduction 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 toReassign.
C. Modify the devDependencies section of the project's Package.json file.
D. Configure WhiteSource Bolt to scan the node_modules directory only.

Answer: A,C

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 relevantto 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: 172

DRAG DROP

You are creating a container for an ASP.NET Coreapp.

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

How should you configure the file? To answer, drag the appropriate values to the correct targets. Each value must be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

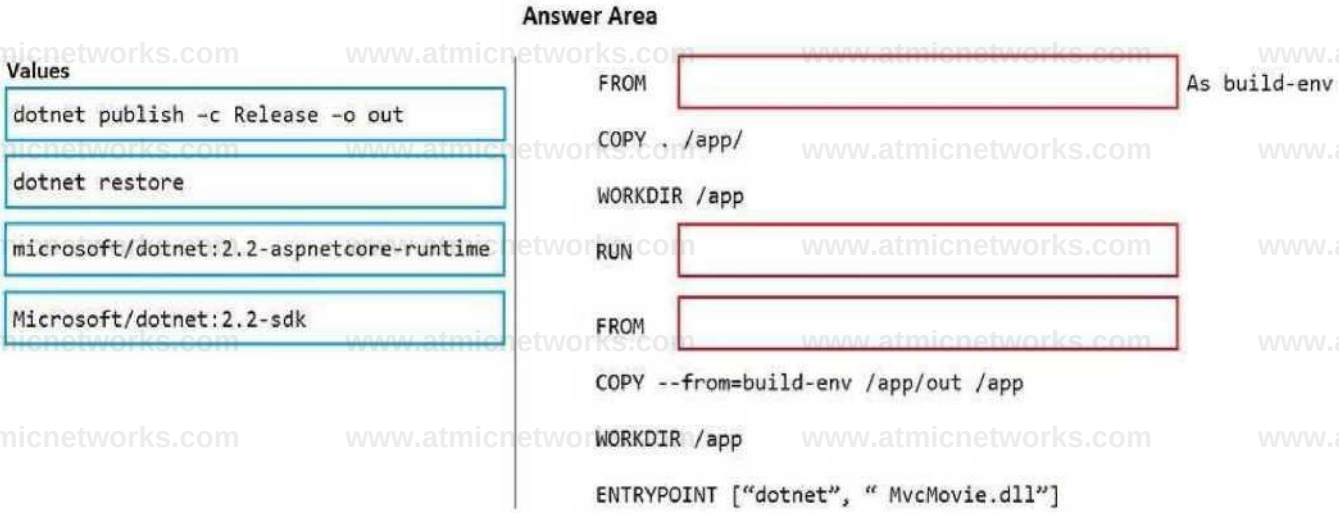

Answer:

Explanation:

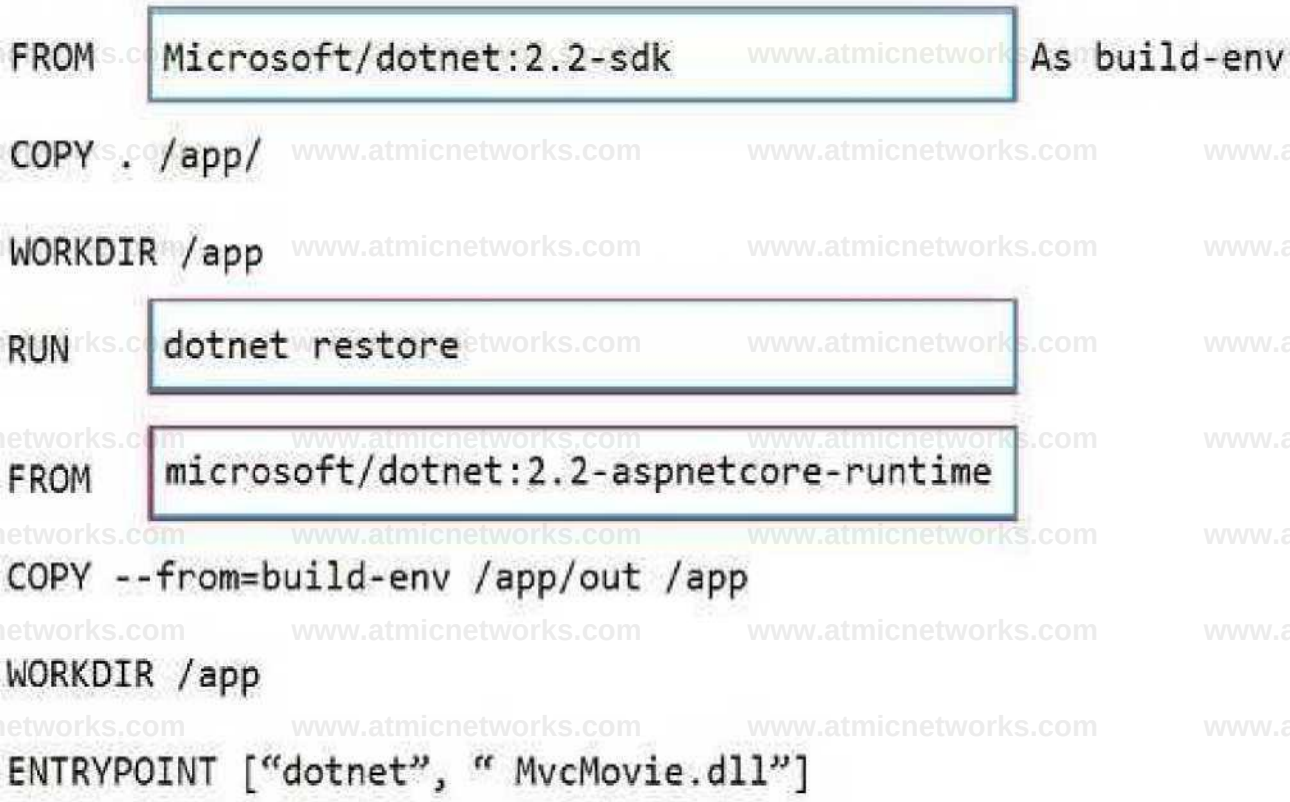

Box 1: microsoft.com/dotnet/sdk:2.3

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

Box 2: dotnet restore

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

COPY *.csproj ./
RUN dotnet restore

Box 3: microsoft.com/dotnet/2.2-aspnetcore-runtime
When building container images, it's good practice to include only the production payload and its dependencies in the container image. We don't want the .NET core SDK included in our final image because we only need the .NET core runtime, so the dockerfile is written to use a temporary container that is packaged with the SDK called build-env to build the app.

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

Question: 173
HOTSPOT

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

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

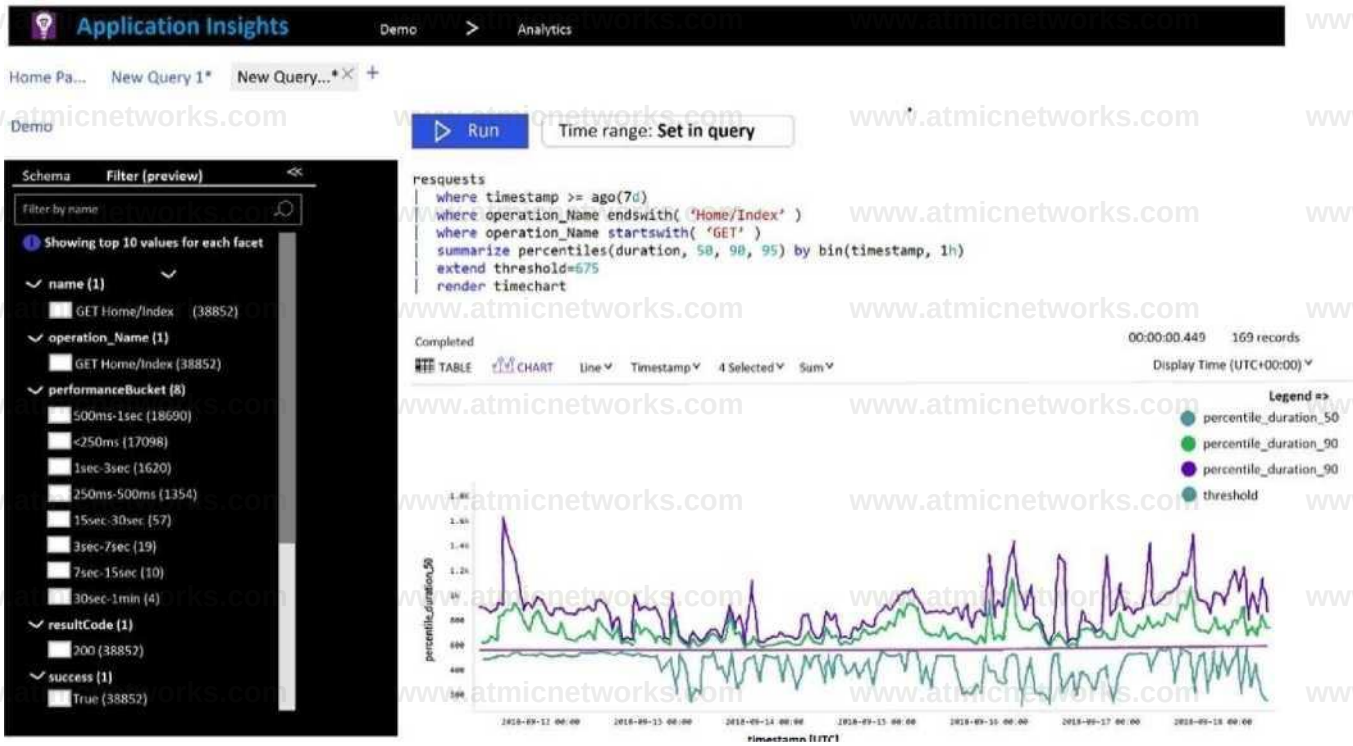

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE:Each correct selection is worth one point.

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

-

percentile_duration_50
percentile_duration_90

percents le_duration_95

threshold

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

item Type resultcode source success

Answer:

Explanation:

Box 1: percentile_duration_95

Box 2:resultCode

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

Question: 174

DRAG DROP

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

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

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

Which policy type should you configure for each requirement? To answer, drag the appropriate policy types to the correct requirements. Each policy type may be used once, more than once, or not at all.
You may need to drag the split bar between panes or scroll to view content.

NOTE:Each correct selection is worth one point.

Policy Types

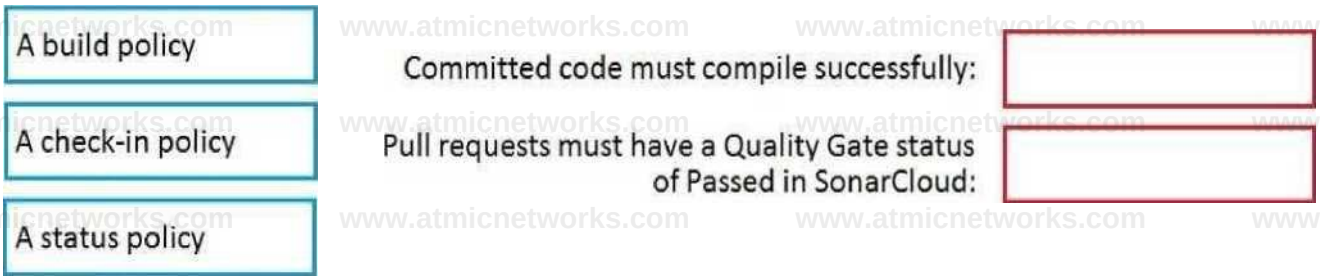

Answer:

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 theyconduct 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 Requiresthat 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/

Question: 175

You use a Git repository in Azure Repos to manage the source code of a web application. Developers commit changes directly to the master branch.

You need to implement achange management procedure that meets the following requirements: The master 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 master branch by using pull requests.

What should you configure in Azure Repos?

D18912E1457D5D1DDCBD40AB3BF70D5D

A. branch policies of the master branch

B. Services in Project Settings
C. Deployment pools in ProjectSettings
D. branch security of the master branch

Answer: A

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: 176

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questionsets 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 planto 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 thismeet the goal?

A. Yes
B. No

Answer: A

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: 177

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 thegoal?

A. Yes
B. No

Answer: B

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: 178

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, youwill 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 developmentprocess:

Licensing violations
Prohibited libraries

Solution: You implement automated security testing.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead use implement continuous integration.

Note: WhiteSource is the leader incontinuous 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 inthe background, checking the security, licensing, and quality of youropen source components against WhiteSource constantly-updated definitive database of open source repositories.

Reference:
https://azuredevopslabs.com/labs/vstsextend/whitesource/

Question: 179

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 questionin 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 Git 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

Answer: A

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: 180

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more thanone 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 tomanage the build and release processes for applications.

You use a Git 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

Answer: B

Explanation:

Question: 181

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution thatmight 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 willnot appear in the review screen.

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

You use a Git repository for applications source control.

You need to implement a pull request strategy that reduces the historyvolume 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

Answer: B

Explanation:

Instead use fast-forward merge.

Note:

No fast-forward merge - This option mergesthe 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: 182

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not

appear in the review screen.

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

You usea Git 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

Answer: B

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: 183

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 Git repository.

Which Git option should you use?

A. Rebase
B. Fetch
C. MergeE1457D5D1DDCBD40AB3BF70D5D
D. Push

Answer: C

Explanation:

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

Question: 184

DRAG DROP

Your companyplans to deploy an application to the following endpoints:

Ten virtual machines hosted in Azure
Ten virtual machines hosted in an on-premises data center environment

All the virtual machines have the Azure Pipelines agent.

You need to implement a release strategy for deploying the application to the endpoints.

What should you recommend using to deploy the application to the endpoints? To answer, drag the appropriate components to the correct endpoints. Each component may be used once, more than once, ornot 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.

Components	Answer Area
A deployment group	Ten virtual machines hosted in Azure:
A management group	Ten virtual machines hosted in an on-premises data center environment:
A resource group	
Application roles	

Answer:

Explanation:

Box 1: A deployment group

When authoring an Azure Pipelines or TFS Release pipeline, you can specify the deployment targets for a job using a deployment group.

If the target machines are Azure VMs, you can quickly andeasily prepare them by installing the Azure Pipelines Agent Azure VM extension on each of the VMs, or by using the Azure Resource Group Deployment task in your release pipeline to create a deployment group dynamically.

Box 2: A deployment group

s:https://docs.microsoft.com/en-us/azure/devops/pipelines/release/deployment-groups

Question: 185

You plan touse 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 isworth one point.

A. Vault
B. Terratest
C. Node.js
D. Yeoman
E. Tiller

Answer: B,D

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 Terraformmodules.

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 Azurespecific 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: 186

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

Answer: B

Explanation:

Feedshave 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/mstall packages	J	/		
Save packages from upstream sources		/		7
Push packages				
Uni ist/dep recate packages				
Promote a package to a view				
Delete/unpubhsh package				
Edit feed permissions				

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

Question: 187

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 Contoso project as shown in the Permissions exhibit (Click the Permissions tab.)

All other permissions have the default values set.

Statements	Yes	No
Web Team can delete DasM.		
Contoso Ipam can view DasM,		
Project administrators can create new dashboards.		

Answer:

Explanation:

Statements	Yes	No
Web Team can delete Dash1.	O	O
Contoso Team can view Dash1.	O	O
Project administrators can create new dashboards.	O	O

Question: 188

You have an Azure DevOps organization named Contoso and anAzure subscription.

You use Azure DevOps to build and deploy a web app named App1. Azure Monitor is configured to generate an email notification in response to alerts generated whenever App1 generates a serverside error.
You need to receive notificationsin Microsoft Teams whenever an Azure Monitor alert is generated. Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Create an Azure logic app that has an HTTP request trigger.
B. Modify the Diagnostics settings in Azure Monitor.
C. Modify an action group in Azure Monitor.
D. Create an Azure Monitor workbook.
E. Create an Azure logic app that has an Azure DevOps trigger.

Answer: A,B

Explanation:

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

Question: 189

You are integrating Azure Pipelines and Microsoft Teams.
You install the Azure Pipelines app in Microsoft Teams.
You have an Azure DevOps organization named Contoso that contains a project name Project1.
You subscribe to Project1 in Microsoft Teams.
You need to ensure that you only receive events about failed builds in Microsoft Teams.
What should you do first?

A. From Microsoft Teams, tun @azure pipelines subscribehttps://dev.azure.com/Contoso/Project1.
B. From Microsoft Teams, run @azure pipelines subscriptions.
C. From Azure Pipelines, enable continuous integration for Project1.
D. From Azure Pipelines, add a Publish Build Artifacts task to Project1.

Answer: A

Explanation:

To start monitoring all pipelines in a project, use the following command inside a channel:

@azure pipelines subscribe [project url]
The project URL can be to any page within your project (except URLs to pipelines).

For example:
@azure pipelines subscribe https://dev.azure.com/myorg/myproject/

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

Question: 190

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

Answer: D

Explanation:

When you connect to a Git repository from your Git 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 credentialsthat 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: 191

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. Project Collection Administrators
B. Reader
C. Project Administrators

D. Contributor

Answer: D

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.

e:

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

https://docs.microsoft.com/en-us/azure/devops/organizations/security/permissions-access-work- tracking?view=azure-devops#queries-and-semantic-search

Question: 192

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.

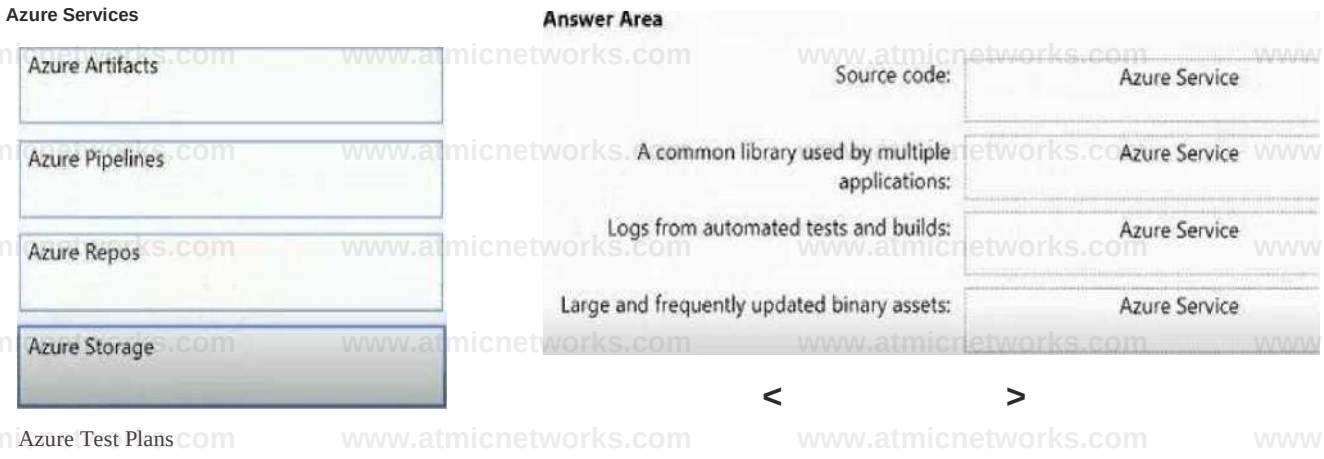

Answer:

Explanation:

Box 1: Azure Repos

Box 2: Azure Artifacts

Use Azure Artifacts to create, host, and share packages withyour 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: 193

DRAG DROP

You manage the Git repository for a large enterprise application.

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

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

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

Actions

Run the git commit command.

Run the git reflog expire command.

Run the git add .gitignore command.

Add Config.json to the .gitignore file.

Delete and recreate the repository.

Answer Area

Answer:

Explanation:

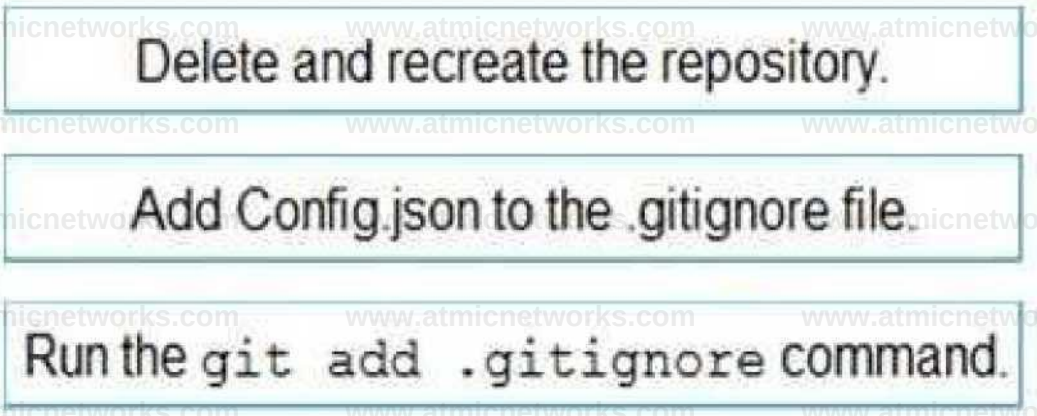

Step 1: Delete and recreate the repository.
Step 2: Add Config.json to the .gitignore file
Each line in the .gitignore excludes a file or set of files that match a pattern.
Example:
ignore a singlefile
Config.json

Step 3: Run the git add .gitignore command
At the initial commit we want basically move from Untracked to Staged, for staging we have to indicate which file we want to move or specify a pattern, as example:

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

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

Question: 194

You have a free tier of anAzure DevOps organization named Contoso. Contoso contains 10 private projects.
Each project has multiple jobs with no dependencies.
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. Purchase self-hosted parallel jobs.
B. Register additional self-hosted agents.
C. Purchase Microsoft-hosted parallel jobs.
D. Configurethe pipelines to use the Microsoft-hosted agents.

Answer: B

Explanation:

https://docs.microsoft.com/en-us/azure/devops/organizations/billing/buy-more-build- vs?view=azure-devops#self-hosted-cicd

Question: 195

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 features 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.

Actions

Load the file to Azure files

Create a PowerShell configuration file.

jcrairta a YAML configuration fife

1 Configure the Custom Script Extension on the virtual machines.

Configure the DSC extension on the virtual machines.

Load the file to Azure Blob storage

Answer Area

Answer:

Explanation:

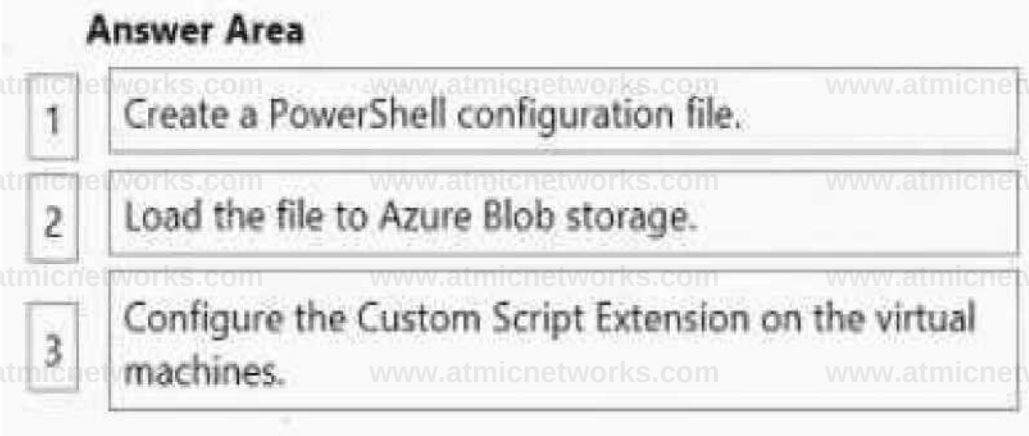

Step 1: Create a PowerShell configuration file

You create a simple PowerShell DSCconfiguration 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 andexecutes 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: 196

You are designing a build pipeline in Azure Pipelines.

The pipeline requires a self-hosted agent. The build pipeline will runonce 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. Azure virtual machines
B. an Azure virtual machine scale set
C. an Azure Kubernetes Service (AKS) cluster
D. Azure Container Instances

Answer: D

Explanation:

https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/agents?view=azure-devops&tabs=browser#faq

Question: 197

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 mustbe able to enable the service in App1 once App2 is deployed.

What should you do?

A. Create a branch in the build.
B. Implement a branch policy.
C. Create a fork in the build.
D. Implement a feature flag.

Answer: D

Explanation:

Reference:

https://docs.microsoft.com/en-us/azure/devops/migrate/phase-features-with-feature-flags

Question: 198

DRAG DROP

You need to deploy Internet Information Services (IIS) to an Azure virtual machinethat runs Windows Server 2019.

How should you complete the Desired State Configuration (DSQ configuration script? To answer, drag the appropriate values to the correct locations. Each value may be used once, more than once, or not at all. You may need todrag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

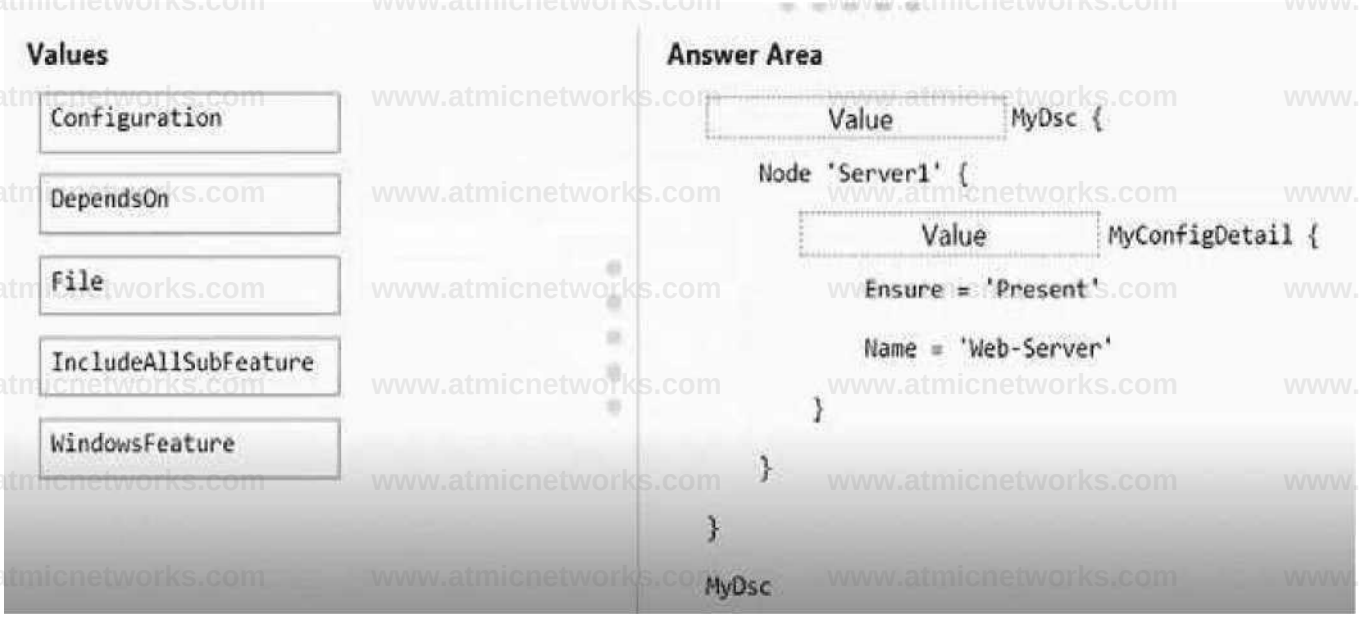

Answer:

Explanation:

Answer Area

| Conf i^urdtion %DSi {

Node *$eryerr {

WWows Feature HyConfigDetail {

Croupe = [1] Present'

Name 'Web-Server'

Box 1: Configuration
The following example shows a simple example of a configuration.

```
configuration IISInstall
{
node "localhost"
{
WindowsFeature IIS
{
Ensure = "Present"
```

```
Name = "Web-Server"
}
}
}
```

Box 2: WindowsFeature

Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/dsc-overview

Question: 199

Your company uses Azure Artifacts for package management.
You need to configure an upstream source in Azure Artifacts for Pythonpackages.
Which repository type should you use as an upstream source?

A. PyPI
B. npmjs.org
C. Maven Central
D. third-party trusted Python

Answer: A

Explanation:

Get started with Python packages in Azure Artifacts

Create a feed
Select Artifacts (in the left navigation of your Azure DevOps project).
On the Artifacts page, select Create Feed.
In the Create new feed dialog box:
In the Name field, give the feed a name.
PyPI is the default repository name for twine, which is a tool for publishingPython packages.

Reference:
https://docs.microsoft.com/en-us/azure/devops/artifacts/quickstarts/python-packages

Question: 200

You have an Azure DevOps organization named Contoso that contains a project named Project 1.
You provision an Azure key vault name Keyvault1.
You need to reference Keyvault1 secrets in a build pipeline of Project1.
What should you do first?

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

Answer: C

Explanation:

Before this will work, the build needs permission toaccess 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: 201

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

Answer: B

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?
DependaBot regularly checks dependencies for updates
If an update is found, DependaBot creates a new branch with this upgrade and Pull Request for approval
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/

Question: 202

You have a Microsoft ASP.NET Core web app in Azure that is accessed worldwide.
You need to run a URL ping test once every five minutes and create an alert when the web app is unavailable from specific Azure regions. The solution must minimize development time.

What should you do?

A. Create an Azure Application Insights availability test and alert.
B. Create an Azure Service Health alert for the specific regions.
C. Create an AzureMonitor Availability metric and alert
D. Write an Azure function and deploy the function to the specific regions.

Answer: A

Explanation:

There are three types of Application Insights availability tests:
URL ping test: a simple test that you can createin the Azure portal.
Multi-step web test
Custom Track Availability Tests

Note: After you've deployed your web app/website, you can set up recurring tests to monitor availability and responsiveness. Azure Application Insights sends web requests to yourapplication at regular intervals from points around the world. It can alert you if your application isn't responding, or if it responds too slowly.

You can set up availability tests for any HTTP or HTTPS endpoint that is accessible from the public internet. You don't have to make any changes to the website you're testing. In fact, it doesn't even have to be a site you own. You can test the availability of a REST API that your service depends on.

Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability#create-a- url-ping-test

Question: 203

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

Answer: A,D

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
Environments supported	Azure	Azure	Azure Other cloud On-premises	Azure Other cloud On-premises
Agent requirements	None	None	None	Requires Log Analytics agent
Data collected	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
Data sent to	Azure Monitor Logs Azure Monitor Metrics	Azure Storage Azure Monitor Metrics	Azure Monitor Logs	Azure Monitor Logs (through Log Analytics agent)

Event 1—11 ih

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

Question: 204

You configure an Azure Application Insights availability test.

You need to notify thecustomer services department at your company by email when availability is degraded.

You create an Azure logic app that will handle the email and follow up actions. Which type of trigger should you use

to invoke the logic app?

A. an ApiConnectiontrigger
B. a Request trigger
C. an HTTPWebhook trigger
D. an HTTP trigger

Answer: C

Explanation:

You can use webhooks to route an Azure alert notification to other systems for post-processing or custom actions. You can use a webhook on an alert to route it to services that send SMS messages, to log bugs, to notify a team via chat or messaging services, or for various other actions.

Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-webhooks

Question: 205

You have an Azure DevOps organization named Contoso and an Azure subscription.
You use Azure DevOps to build a containerized app named Appl and deploy App1 to an Azure container instance named ACM.
You need to restart ACI1 when App1 stops responding.
What should you do?

A. Add a liveness probe to the YAML configuration of App1.
B. Use Connection Monitor in Azure Network Watcher.
C. Add a readiness probe to the YAML configuration of App1.
D. Use IP flow verify in Azure Network Watcher.

Answer: A

Explanation:

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

Question: 206

DRAG DROP
You have an Azure DevOps release pipeline as shown in the following exhibit.

You need to complete the pipeline to configure OWASP ZAP for security testing.

Which five Azure CLI tasks should you add insequence? To answer, move the tasks from the list of tasks to the answer area and arrange them in the correct order.

Tasks

Build machine image

Convert Report Format

Download the file

Publish Test Results

Docker CLI installer

Destroy OWASP Container

Call the Baseline Scan

Answer Area

Answer:

Explanation:

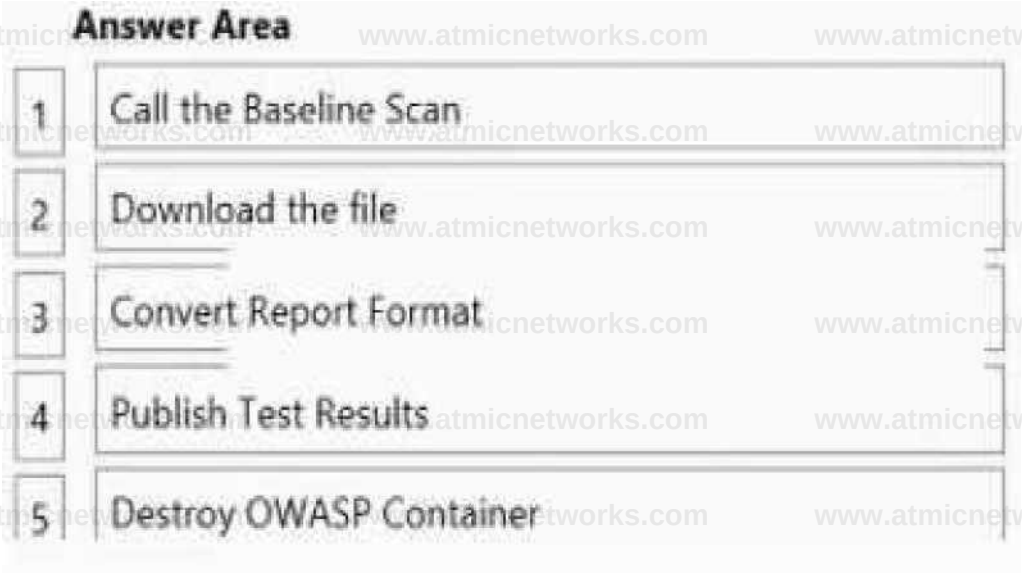

Defining the Release Pipeline

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.

Create Resource Group (if not created)
Azure QI

Create Storage Account (if not created) Azure QI

., Create OWASP Container
Azure QI

a Call the Baseline Scan
Azure QI

Download the file
Azure QI

rw Convert Report Format
PowerShell

n Publish Test Results
Publish Test Result

Destroy OWASP Container
Azure QI

Reference:
https://devblogs.microsoft.com/premier-developer/azure-devops-pipelines-leveraging-owasp-zap- in-the-release-pipeline/

Question: 207

DRAG DROP

You have several Azure virtual machines that run Windows Server 2019.
You need to identify the distinct event IDs of each virtual machine as shown in the following table.

How should you complete the Azure Monitor query? To answer, drag the appropriate values to the correct locations. Each value may be used once, more than once, or not at all. You may need todrag the split bar between panes or scroll to view content.

Answer Area

```
Event
| where TimeGenerated > ago(12h)
| order by TimeGenerated desc
                    Value
                    Value        by Computer
```

Answer:

Explanation:

```
Event
I where TimeGenerated > ago(12h)
I order by TimeGenerated desc
I I summarize makelist(EventlD)j by Computer
```

You can use makelist to pivot data by the order of values in a particular column. For example, you may want to explore the most common order events take place on your machines. You can essentially pivot the data by the order of EventIDs on each machine.

Example:

```
Event
| where TimeGenerated > ago(12h)
| order by TimeGenerated desc
| summarize makelist(EventID) by Computer
```

Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/advanced-aggregations

Question: 208

You have a project in Azure DevOps named Project1. Project! contains a pipeline that builds acontainer image named Image! and pushes Image1 to an Azure container registry named ACR1. Image! 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. Create and run an Azure Container Registry task.
B. Add a Docker Hub service connection to Azure Pipelines.
C. Enable the Azure Event Grid resource provider and subscribe to registry events.
D. Create a service hook in Project1.

Answer: A

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: 209

Your company uses the following resources:
Windows Server 2019 container imageshosted 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.

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

Answer: C,E

Explanation:

https://docs.microsoft.com/en-us/azure/security-center/features-paas

https://docs.microsoft.com/en-us/azure/security-center/container-security

Question: 210

: 213

You have an Azure subscription that contains an Azure Active Directory (Azure AD) tenant.

You are configuring a build pipeline in Azure Pipelines that will include a task named Task1. Task1 will authenticate by using an Azure AD service principal.

Which three values should you configure for Task1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. the object ID
B. the tenant ID
C. the app ID
D. the client secret
E. the subscription ID

Answer: B,C,D

Explanation:

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

Question: 211

DRAG DROP

You are configuring an Azure DevOps deployment pipeline. The deployed applicationwill 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 actionsto the answer area and arrange them in the correct order.

Actions

- Export a certificate from the key vault
- Add an app registration in Microsoft Azure Active Directory (Azure AD), part of Microsoft Entra.
- Generate a self signed certificate.
- Create a service principal in Microsoft Azur Entra.
- Configure an access policy in the key vault
- Add an Azure Resource Manager service connection to the pipeline.

Answer Area

Active Directory (Azure AD), part of Microsoft

Answer:

Explanation:

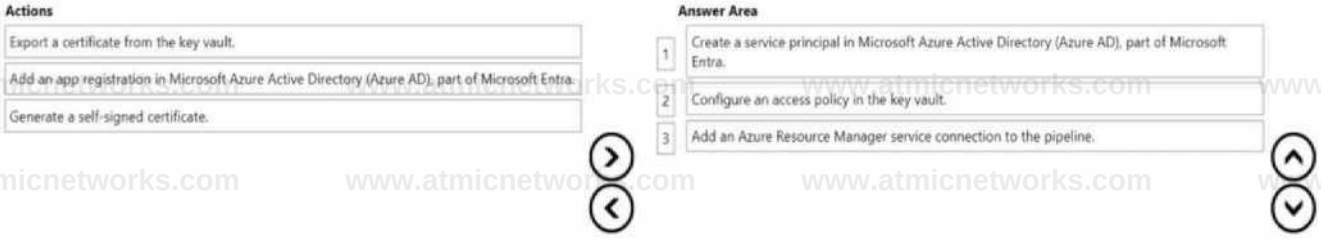

Question: 212

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

Answer: C

Explanation:

The Azure Pipelines app uses the OAuth authentication protocol, and requires Third-party application access via OAuthfor 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: 213

HOTSPOT

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

MTM PartsUnlimitodE2E

&*ww tufa V*r.»We» T*>9gc4 ♦^'emtsn H^isty

Pipeline

Arai V Gel sources

Cloud Agent

k-i a > - N*

NuGet revor#

- -s •

Compile Apptcabon

ft

Copy Cite

Pubitih Artdea

JI

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

Answer Area

Answer:

Explanation:

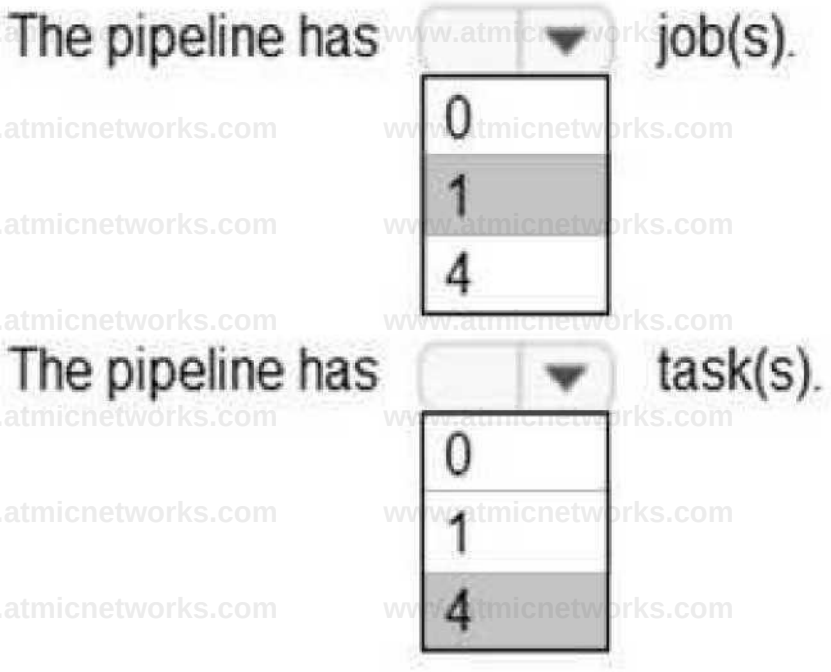

Box 1: 1
The Cloud agent job only.

Box 2: 4
The pipelines has the four tasks: NuGetrestore, Compile Application, Copy Files, and Publish Artifact.

Reference:
https://azuredevopslabs.com/labs/azuredevops/continuousintegration/

Question: 214

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

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

Which two resources should you include in the solution? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. an Azure Monitor alert that has a dynamic threshold
B. an Azure Automation runbook
C. an Azure Monitor alert that uses an action group that has an email action
D. the Azure Monitor autoscale settings
E. an Azure Monitor alert that has a static threshold

Answer: B,C

Explanation:

Question: 215

You have an Azure DevOps organization named Contoso and an Azure subscription. The subscription contains an Azure virtual machine scale set named VMSS1 and an Azure Standard Load Balancer named LB1. LB1 distributes incoming requests across VMSS1 instances.

You use Azure DevOps to build a web app named Appl and deploy App1 to VMSS1. App1 is accessible via HTTPS only and configured to require mutual authentication by using a client certificate.

You need to recommend a solution for implementing a health check of App1. The solution must meet the following requirements:

- Identify whether individual instances of VMSSl are eligible for an upgrade operation.
- Minimize administrative effort.

What should you include inthe recommendation?

A. the Custom Script Extension
B. the Application Health extension
C. Azure Monitor autoscale
D. an Azure Load Balancer health probe

Answer: B

Explanation:

https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets- health-extension#when-to-use-the-application-health-extension

Question: 216

You are monitoring the health and performance of an Azure web app by using Azure Application Insights. You need to ensure that an alert is sent when the web app has a sudden rise in performance issues and failures. What should you use?

A. Application Insights Profiler
B. Continuous export
C. Smart Detection
D. custom events
E. usage analysis

Answer: C

Explanation:

Smart Detection automatically warns you of potential performance problems and failure anomalies in your web application. It performs proactive analysis of the telemetry that yourapp sends to Application Insights. If there is a sudden rise in failure rates, or abnormal patterns in client or server performance, you get an alert.

Reference:

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

Question: 217

HOTSPOT

You have a project in Azure DevOps named Contoso App that contains pipelines in Azure Pipelinesfor GitHub repositories.

You need to ensure that developers receive Microsoft Teams notifications when there are failures in a pipeline of Contoso App. What should you run in Teams? To answer, select the appropriate options in the answer are a. NOTE: Each correct selection is worth one point.

Answer Area

Answer:

Explanation:

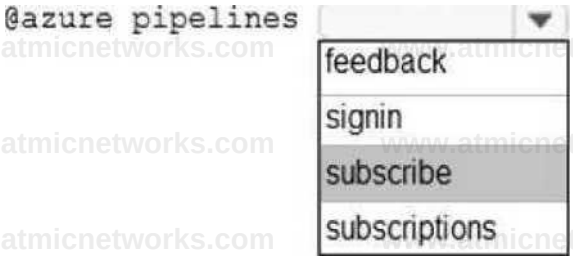

https://devazure.com/contoso/contoso-app/

https://dev azure com/contoso/contoso-app/_bui Id

https://devazure.com/contoso/contoso-app/_packaging

https://devazure.com/contoso/contoso-app/_work-items

Box 1: subscribe

To start monitoring all pipelines in a project, use the following command inside a channel:

@azure pipelines subscribe [project url]

Box 2: https://dev.azure.com/contoso/contoso-app/

Subscribe to a pipeline or all pipelines in a project to receive notifications:

@azure pipelines subscribe [pipeline url/ projecturl]

Question: 218

You have a private GitHub repository.

You need to display the commit status of the repository on Azure Boards. What should you do first?

A. Create a GitHub action in GitHub.

B. Add the Azure Pipelines app to the GitHub repository.

C. Configure multi-factor authentication (MFA) for your GitHub account.

D. Add the Azure Boards app to the repository.

Answer: D

Explanation:

To connect Azure Boards to GitHub.com, connect and configure from Azure Boards. Or, alternatively, install andconfigure the Azure Boards app from GitHub. Both methods have been streamlined and support authenticating and operating via the app rather than an individual.

Note (see step 4 below):
Add a GitHub connection:
Sign into Azure Boards.
Choose (1) Project Settings, choose (2) GitHub connections and then (3) Connect your GitHub account.
If this is your first time connecting to GitHub from Azure Boards, you will be asked to sign in using your GitHub credentials.
Choose an account for which you are an administrator for the repositories you want to connect to.
The Add GitHub Repositories dialog automatically displays and selects all GitHub.com repositories for which you are an administrator. Unselect any repositories that you don't want to participate in the integration.

Add GitHub repositories

Add the GitHub repositories you want to use with your Azure Boards.

V Filter by keywords

Viewing 4. 4 selected

© JamalHart/fabrikam-apps-2

© JamalHart/fabrikam-demo

© ^ JamalHart/fabrikam- open-source

© JamalHart/fabrikam-suite

Question: 219

You have a build pipeline in Azure Pipelines.

You create a Slack App Integration.

You need to send build notifications to a Slack channel named #Development.

What should you do first?

A. Configure a service connection.
B. Create a service hook subscription.
C. Create a project-level notification.
D. Create a global notification.

Answer: B

Explanation:

Create a service hook for Azure DevOps with Slack to post messages to Slack in response to events inyour Azure DevOps organization, such as completed builds, code changes, pull requests, releases, work items changes, and more.

Note:

1. Go to your project Service Hooks page: https://{orgName}/{project_name}/_settings/serviceHooks

Select CreateSubscription.

3. Choose the types of events you want to appear in your Slack channel.
4. Paste the Web Hook URL from the Slack integration that you created and select Finish.
5. Now, when the event you configured occurs in your project, a notification appears in your team's Slack channel.

Reference:
https://docs.microsoft.com/en-us/azure/devops/service-hooks/services/slack

Question: 220

You are building an ASP.NET Core application.
You plan to create an application utilization baseline by capturing telemetry data.
You need to add code to the application to capture the telemetry dat
a. The solution must minimize the costs of storing the telemetry data.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Add the <InitialSamplingPercentage>99</InitialSamplingPercentage> parameter to the ApplicationInsights.config file.
B. From the code of the application, enable adaptive sampling.
C. From the codeof the application, add Azure Application Insights telemetry.
D. Add the <MaxTelemetryItemsPerSecond>5</MaxTelemetryItemsPerSecond> parameter to the ApplicationInsights.config file.
E. From the code of the application, disable adaptive sampling.

Answer: C,E

Explanation:

"Fixed-rate sampling reduces the volume of telemetry sent from both your ASP.NET or ASP.NET Core or Java server and from your users' browsers. You set the rate. The client and server will synchronize their sampling so that, in Search, you can navigatebetween related page views and requests."https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling https://docs.microsoft.com/en-us/azure/azure-monitor/app/asp-net-core
https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling#configuring-adaptive- sampling-for-aspnet-core-applications

Question: 221

You have an Azure subscription that containsresources in several resource groups.
You need to design a monitoring strategy that will provide a consolidated view. The solution must support the following requirements:

- Support role-based access control (RBAC) by using Azure Active Directory (Azure AD) identities.
- Include visuals from Azure Monitor that are generated by using the Kusto query language.
- Support documentation written in markdown.
- Use the latest data available for each visual.

What should you use to create the consolidated view?

A. Azure Data Explorer
B. Azure dashboards

C. Azure Monitor

D. Microsoft Power BI

Answer: A

Explanation:

There are several tools available for running queries in Azure Data Explorer, including Kusto.

Kusto uses a role-based access control (RBAC) model,under which authenticated principals are mapped to roles, and get access according to the roles they're assigned.

Note: Azure Data Explorer is a highly scalable and secure analytics service that enables you to do rich exploration of structured and unstructured data for instant insights. Optimized for ad-hoc queries, Azure Data Explorer enables rich data exploration over raw, structured, and semi-structured data delivering fast time to insight. Query with a modern, intuitive query language that offers fast,ad-hoc, and advanced query capabilities over high-rate data volumes and varieties

Reference:
https://docs.microsoft.com/en-us/azure/data-explorer/tools-integrations-overview

Question: 222

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. Enable Test Impact Analysis (TIA).

B. Enable test slicing.
C. Clear Flaky tests included in test pass percentage
D. Set Flaky test detection to Off
E. Manually mark the test as flaky.

Answer: C,E

Explanation:

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

Question: 223

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 AzurePipelines. Each virtual machine has the Web Server (IIS) role installed and configured.
You need to ensure that the web server configurations pin 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.

Actions	Answer Area
Install the custom Desired State Configuration (DSC) extension on the virtual machines.	
Compile the Desired State Configuration (DSC) configuration.	
Import a Desired State Configuration (DSC) configuration into the Azure Automation account	
Create an Azure Automation account	
Onboard the virtual machines to the Azure Automation account	

®

Answer:

Explanation:

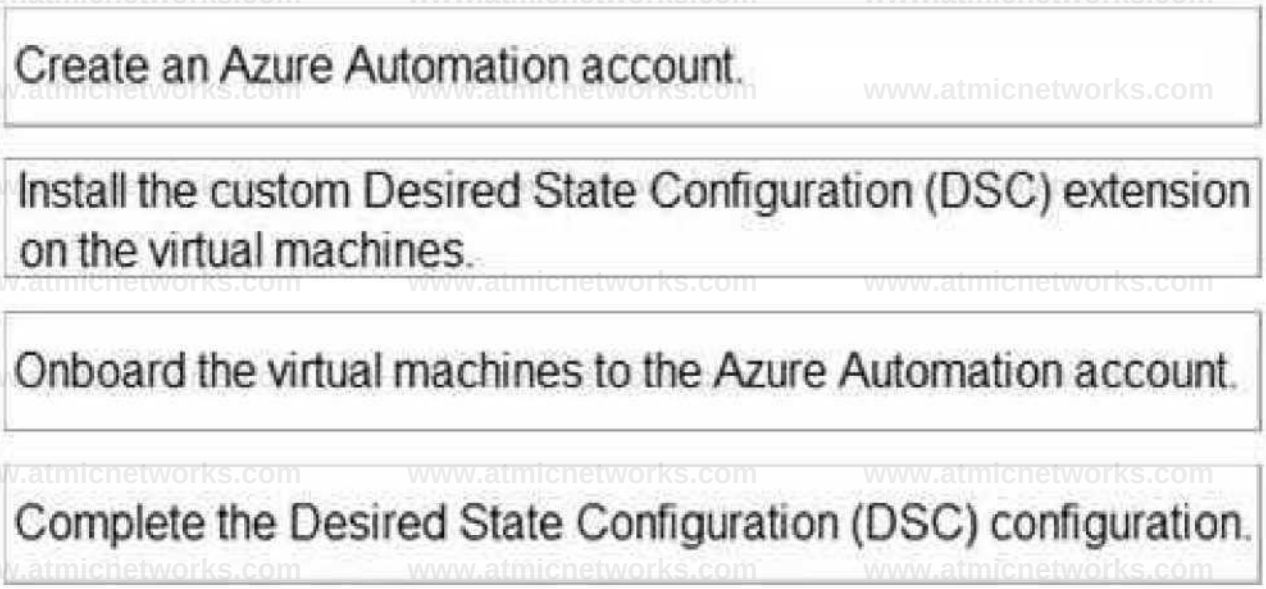

Step1: Create an Azure Automation account. An Azure Automation account is required.

Step 2: Install the custom Desired State Configuration (DSC) extension on the virtual machines Underthe 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.

Step4: Complete the Desired State Configuration (DSC) configuration.
Create a DSC configuration.

Reference:
https://docs.microsoft.com/en-us/azure/automation/automation-dsc-onboarding

Question: 224

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 threeactions 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.

Actions	Answer Area
[Create an npm package.	N
Create an Azure Artifacts feed that uses upstream sources,	
Modify the configuration files to reference the Azure Artifacts feed,	
Run an initial package restore,	^^^
Create a NuGet package.	z**v
Create a Microsoft Visual Studio project that includes all the packages.	v^x

Answer:

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 multiplefeeds together into a single feed.

Step 3: Create a Micrtosoft 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: 225

DRAG DROP

As part of your application build process, you need to deploy a group of resources to Azure by using an Azure Resource Manager template located on GitHub.

Which three action 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.

Actions	Answer Area
Create a package.	
Add an Azure Resource Group Deployment task.	
Create a job agent	
Create a release pipeline.	—
Set the template parameters.	

Answer:

Explanation:

Create a release pipeline.

Add an Azure Resource Group Deployment task

Set the template parameters

Step 1: Create a release pipeline

You need to create a new pipeline.

You can integrate Azure Resource Manager templates (ARM templates) with Azure Pipelines for continuous integration and continuous deployment (CI/CD).

Step 2: Add an Azure Resource Group Deployment task

Step3: Set the template parameters

Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/add-template-to- azure-pipelines

Question: 226

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. Create a pipeline variable.
B. Add a deployment gate.
C. Configure a deployment trigger.
D. Modify the Deploymentqueue settings.

Answer: B

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 toapprove or reject a deployment. You can use gates to ensure that the release meets a wide range or criteria, without requiring user intervention.

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

Question: 227

You have an Azure DevOps project that uses many package feeds.
You need to simplify theproject 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. views in AzureArtifacts
C. upstream sources
D. a symbol server

Answer: C

Explanation:

Question: 228

During a code review, you discover quality issues in a Java application.
You need to recommend a solution to detect quality issues including unused variables and empty catchblocks.
What should you recommend?

A. In an Xcode build task, select Use xcpretty from Advanced.
B. In a Maven build task, select Run PMD.
C. In a Grunt build task, select Enabled from Control Options.
D. In a Gulp build task, specify a custom condition expression.

Answer: B

Explanation:

PMD is a source code analyzer. It finds common programming flaws like unused variables, empty catch blocks, unnecessary object creation, and so forth.

There is an Apache Maven PMD Plugin which allows you to automaticallyrun the PMD code analysis tool on your project's source code and generate a site report with its results.

Question: 229

You have a project in Azure DevOps named Project1. Project1 contains a build pipeline named Pipe1 that builds an application namedAppl.
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 namedApp2.
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. Create two container jobs.
B. Changethe self-hosted agent to use Red Hat Enterprise Linux (RHEL) 9.
C. Add another self-hosted agent
D. Add a Docker Compose task to the build pipelines.

Answer: A

Explanation:

Question: 230

Your company implements an Agile development methodology.
You plan toimplement retrospectives at the end of each sprint.
Which three questions should you include? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Who performed well?
B. Who should have performed better?
C. What could have gone better?

D. What went well?
E. What should we try next?

Answer: C,D,E

Explanation:

https://www.scrum.org/resources/what-is-a-sprint-retrospective

Question: 231

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

You use a Git 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 asdefault branch
B. Approve with suggestions
C. Cherry-pick
D. Reactivate
E. Revert

Answer: C

Explanation:

Cherry-pick a pull request

To copy changes made in a pull request to another branch in your repo, follow these steps:

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.

In Target branch, enter the branch you want to merge the copied changes.

In Topic branch name, enter a new branch to contain the copied changes, then select Cherry-pick. Select Create pull request to merge the topic branch into the target branch to complete the cherrypick.

nce:

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

Question: 232

Note: This question is part of a series of questions that present thesame 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

Answer: B

Explanation:

Instead implement continuous integration.

Note: WhiteSource is the leader incontinuous 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 inthe 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: 233

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 toimplement 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 buildserver, 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.

Answer: D,E

Explanation:

https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/v2-windows?view=azure-devops (Get PAT, run config)

Question: 234

You need to execute inline testing of an Azure DevOps pipeline that uses a Docker deploymentmodel.
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

Answer: D

Explanation:
"Build and test with a multi-stage Dockerfile: build and tests execute inside the container using a multi-stage Docker file, as such test results are not published back to the pipeline."https://docs.microsoft.com/en-us/azure/devops/pipelines/tasks/test/publish-test-results?view=azure-devops&tabs=trx%2Cyaml

Question: 235

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

Answer: C

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 integratesinto 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: 236

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

You need toidentify 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

Answer: D

Explanation:

With enterprise level of SonarQube you can use OWASP that runs the securityscans for known vulnerabilities.
https://www.sonarqube.org/features/security/https://www.sonarqube.org/features/security/owasp
/?gclid=Cj0KCQiAzZL-BRDnARIsAPCJs70Teq0-efI2Hd_h-
kykCB7I_C7L88Q7kpiuTzuD6Xw1jUb6ZqIP7O0aApVzEALw_wcB

Question: 237

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

Answer: A

Explanation:

An action group is a collection ofnotification 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: 238

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

Answer: B

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: 239

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 Git repository in AzureDevOps.

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

Answer: A

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: 240

HOTSPOT
Your company uses GitHub for source control. GitHub repositories store source code and store process documentation.
The process documentation is saved as Microsoft Word documents that contain simple flow charts stored as .bmp files.
You need to optimize the integration and versioning of the process documentation and the flow charts. The solution must meet the following requirements:

- Store documents as plain text.
- Minimize the number of files that must be maintained.
- Simplify the modification, merging, and reuse of flow charts.
- Simplify the modification, merging, and reuse of documents.

What should you include in the solution? To answer, select the appropriate options in the answer are a. NOTE: Each

correct selection is worth one point.

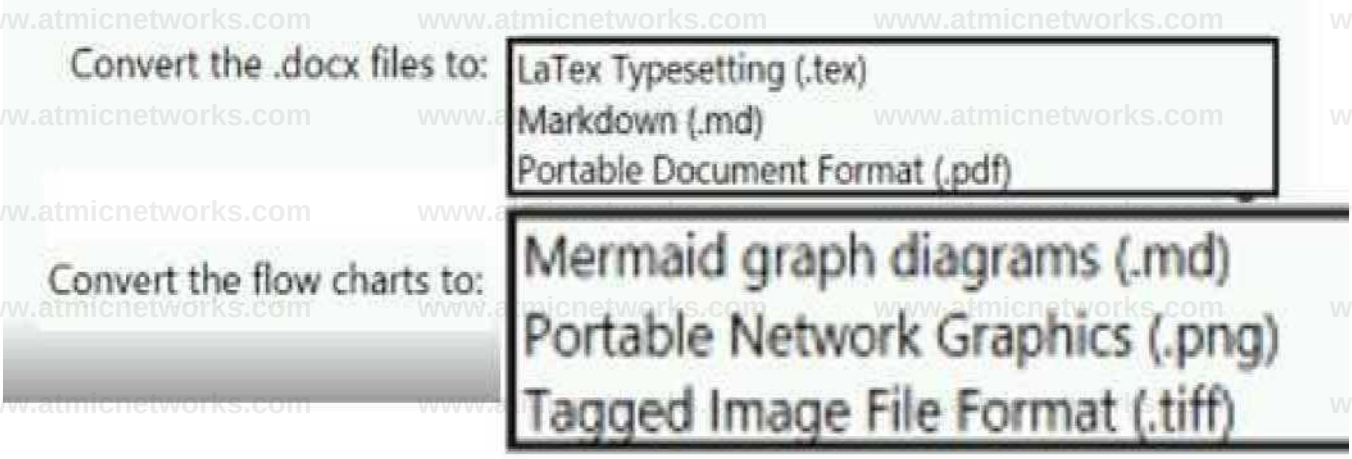

Answer:

Explanation:

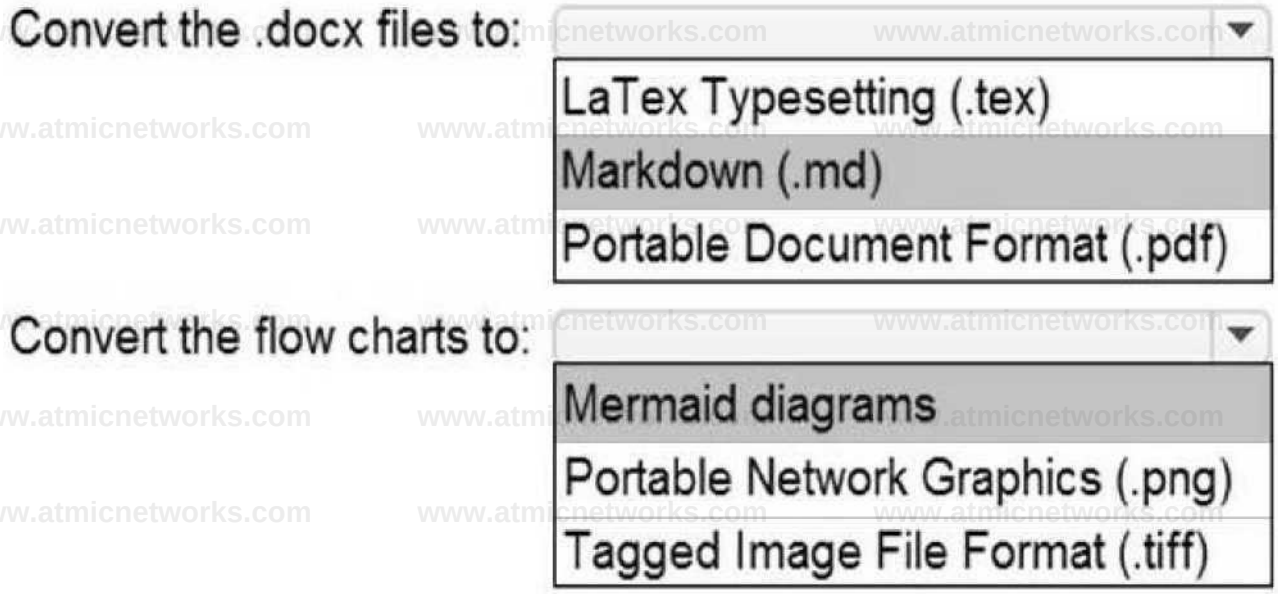

Box 1: Markdown (.md)
Github understands several text formats, including .txt and .md. .md stands for a file written in Markdown.

Box 2: Mermaid diagrams
Mermaid lets you create diagrams and visualizations using text and code.
It is a Javascript based diagramming and charting tool that renders Markdown-inspired text definitions to create and modify diagrams dynamically.

Reference:
https://ourcodingclub.github.io/tutorials/git/

https://mermaid-js.github.io/mermaid/#/

Question: 241

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

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

Answer: C

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

https://docs.microsoft.com/en-us/azure/active-directory/develop/quickstart-v2-aspnet-webapp

Question: 242

HOTSPOT
You use Azure Pipelines lo manage the build and deployment of apps.

You are planning the release strategies for a new app. You need to choose strategies for the following scenarios:

- Releases will be made available to users who are grouped by their tolerance for software faults.
- Code will be deployed to enable functionality that will be available in later releases of the app.
- When a new release occurs, the existing deployment will remain active to minimize recovery time if a return to the previous version is required.

Answer Are*

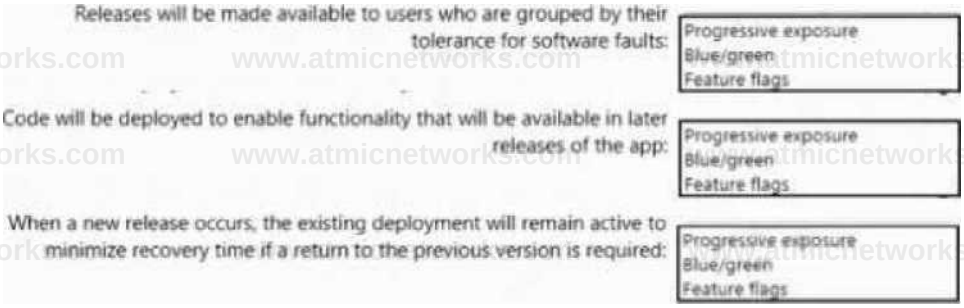

Answer:

Explanation:

Answer Are*

Releases will be made available to users who are grouped by their _ tolerate for software faults ^^ exposure

Code will be deployed to enable functionality that will be available m later _ Feature flags releases of the *pp

When a new release occurs, the existing deployment will remain active to minimize recovery time if a return to the previous version is required ^ areen

Box 1: Progressive exposure

Continuous Delivery may sequence multiple deployment “rings” for progressive exposure (also known as “controlling the blast radius”). Progressive exposure groups users who get to try new releases tomonitor their experience in “rings.” The first deployment ring is often a “canary” used to test new versions in production before a broader rollout. CD automates deployment from one ring to the next and may optionally depend on an approval step, in whicha decision maker signs off on the changes electronically. CD may create an auditable record of the approval in order to satisfy regulatory procedures or other control objectives.

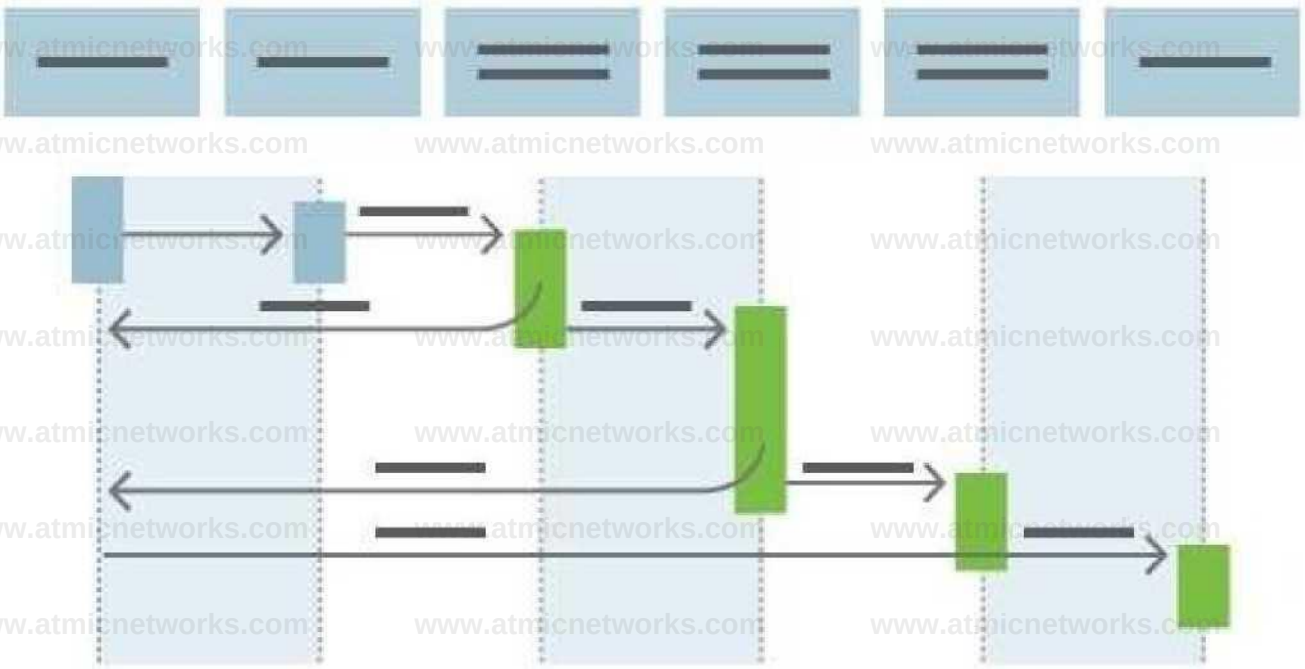

Box 2: Feature flags

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.

Box 3: Blue/green

Blue/green deployments which means that instead of replacing the previous version (here we refer to this version as blue), we bring up the new version (here referred to as the green version) next to the existing version, but not expose it to the actual users right away. On the condition of having successfully validated that the green version works correctly, we will promote this version to the public version by changing the routing configuration without downtime. If something is wrong with the green version we can revert back without users every noticing interruptions.

Reference:

https://docs.microsoft.com/en-us/azure/devops/learn/what-is-continuous-delivery

https://docs.microsoft.com/en-us/azure/devops/migrate/phase-features-with-feature-flags

https://medium.com/@denniszielke/continuous-kubernetes-blue-green-deployments-on-azure- using-nginx-appgateway-

or-trafficmanager-4490bce29cb

Question: 243

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 machinesand 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.

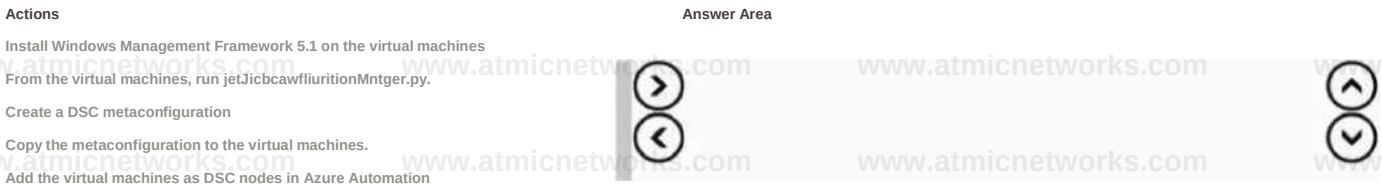

Answer:

Explanation:

Step 1: Create a DSC metaconfiguration

Load up the DSC Configuration into Azure Automation.

Step 2: Copy the metaconfiguration to thevirtual 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: 244

You have an Azure DevOps organization named Contoso.

Youneed to receive Microsoft Teams notifications when work items are updated.

What should you do?

A. From Azure DevOps. configure a service hook subscription.
B. From Microsoft Teams, configure a connector.
C. From Microsoft Teams, add a channel.
D. From Azure DevOps. install an extension.
E. From the Microsoft Teams admin center configure external access.

Answer: A

Explanation:

https://azuredevopslabs.com/labs/vstsextend/teams/

Question: 245

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.

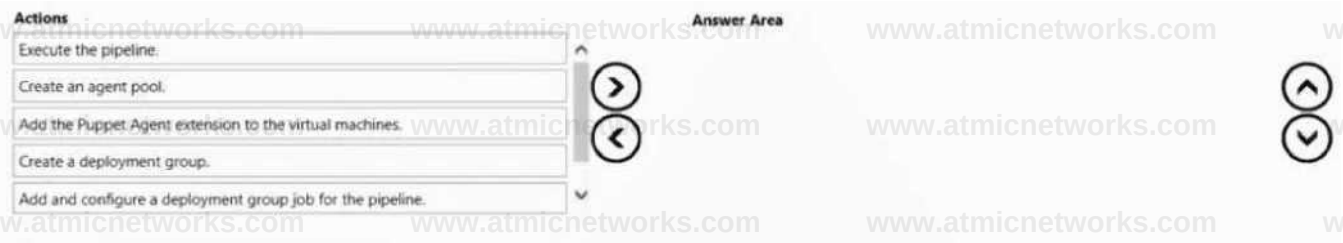

Answer:

Explanation:

Create an agent pool

Create a deployment group

Execute the Azure Pipelines Agent extension to the virtual machines

Add and configure a deployment group job for the pipeline

Step 1: Create an agent pool

AzurePipelines 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 oneach machine.

Step 3: Execute 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: 246

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 tocomplete.

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.

Answer: B,E

Explanation:

Questio n: I need more hosted build resources. What can I do?

Answer: The Azure Pipelinespool provides all Azure DevOps organizations with cloud- Explanation: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: Host your own agents on infrastructure that you manage.

Buy additional parallel jobs.

Reference:

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

Question: 247

You plan to provision a self-hosted Linux agent
Which authentication mechanism should you use to register the self-hosted agent?

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

Answer: B

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

https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/v2-linux?view=azure-devops

Question: 248

DRAG DROP
You have a protect in Azure DevOps.
You need to associate an automated testto 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.

Action# **Answer Area**

Create a test project Create a work rtem. Debug the project Check m a project to the Azure DevOps repository Add the automated test lb a build pipeline

Answer:

Explanation:

Answer Area

1 I Create a test project

2 Check tn a project to the Azure DevOps repository

3 Add the automated test to a build pipeline

The process to associate an automated test with a test case is:

Create a test project containing your automated test. What types of tests are supported?

Check your test project into an Azure DevOps or Team Foundation Server (TFS) repository.

Create a build pipeline for your project, ensuring thatit contains the automated test. What are the differences if I am still using a XAML build?

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:
https://docs.microsoft.com/en-us/azure/devops/test/associate-automated-test-with-test-case

Question: 249

HOTSPOT

You use Azure DevOps to manage the build and deployment of an app named App1.

You have a release pipeline that deploys a virtual machine named VM1.

You plan to monitor the release pipelineby using Azure Monitor

You need to create an alert to monitor the performance of VM1. The alert must be triggered when the average CPU usage exceeds 70 percent for five minutes. The alert must calculate the average once every minute.

How should you configure the alert rule? To answer, select the appropriate options in the answer area.

Answer Are*

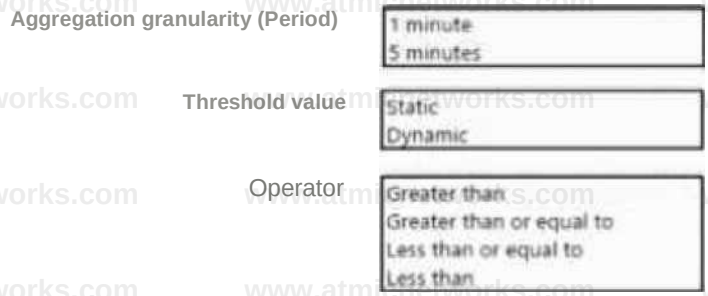

Answer:

Explanation:

Answer Arce

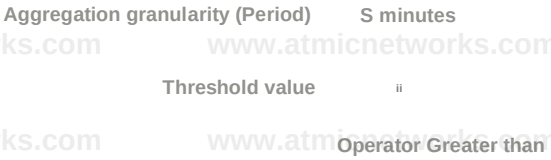

Box 1: 5 minutes
The alert must calculate the average once every minute.

Note: We [Microsoft] recommend choosing an Aggregation granularity (Period) that is larger than the Frequency of evaluation, to reduce the likelihood of missing the first evaluation of added time series

Box 2: Static

Box 3: Greater than

Example, say you have an App Service plan for your website. You want to monitor CPU usage on multiple instances running your web site/app. You can do that using a metricalert rule as follows: Target resource: myAppServicePlan
Metric: Percentage CPU
Condition Type: Static
Dimensions
Instance = InstanceName1, InstanceName2
Time Aggregation: Average
Period: Over the last 5 mins
Frequency: 1 min
Operator: GreaterThan
Threshold: 70
Like before, this rule monitors if the average CPU usage for the last 5 minutes exceeds 70%.
Aggregation granularity

Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-metric-overview

Question: 250

You have a web app hosted on Azure App Service. The web app stores data in an Azure SQL database. You need to generate an alert when there are 10,000 simultaneous connections to the database. The solution must minimize deve4opment effort.
Which option should you select in the Diagnostics settings of the database?

A. Send to Log Analytics
B. Archive to m storage account
C. Stream to an event hub

Answer: A

Explanation:

ENABLE DIAGNOSTICS TO LOG ANALYTICS
Thisconfiguration is done PER DATABASE

1. Click on Diagnostics Settings and then Turn On Diagnostics

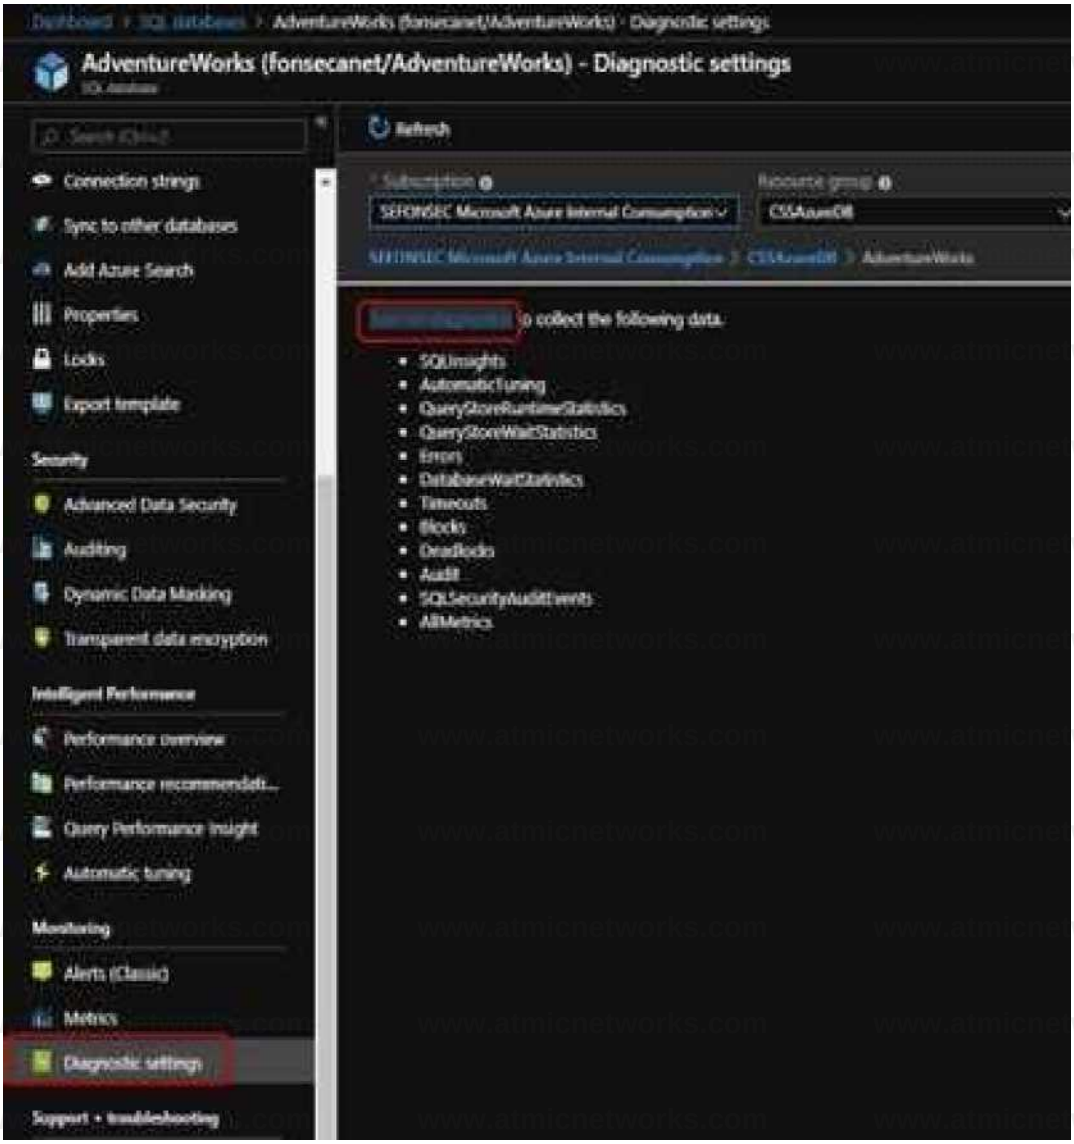

2. Select to Send to Log Analytics and select the Log Analytics workspace. For this sample I will selected only Errors

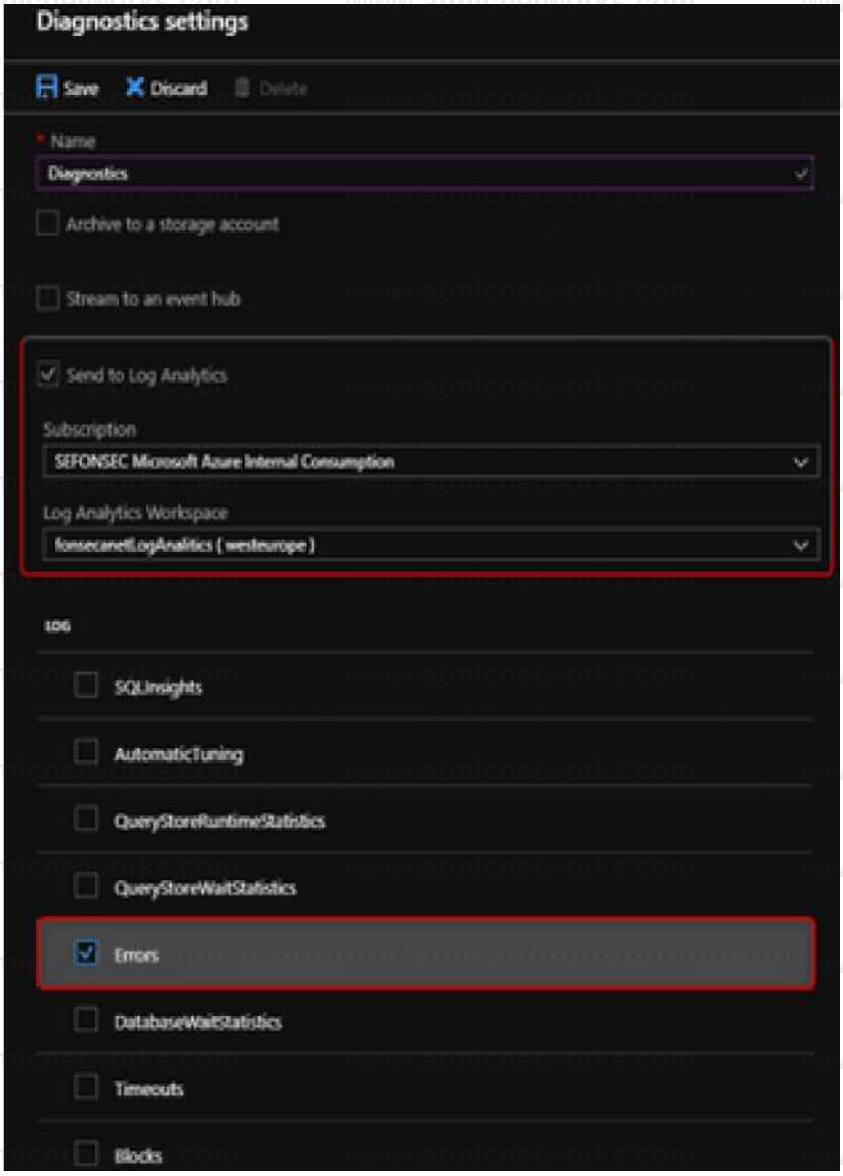

Reference:
https://techcommunity.microsoft.com/t5/azure-database-support-blog/azure-sql-db-and-log- analytics-better-together-part-1/ba-p/794833

Question: 251

You create an alert rule in AzureMonitor as shown in the following exhibit.

Which action will trigger an alert?

A. a failed attempt to delete the ASP-9bb7 resource
B. a change to a rote assignment for the ASP-9bb7 resource
C. a successful attempt to delete the ASP-9bt>7 resource
D. afailed attempt to scale up the ASP-9bb7 resource

Answer: A

Explanation:

Question: 252

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 Appl. What should you configure?

A. a liveness probe
B. an Azure Load Balancer health probe
C. a readiness probe
D. a performancelog

Answer: C

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.

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

Question: 253

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

availability

Learn more about configuring tests against applications hosted behind a firewall

Test type

URL ping test

* URL

https://app.contoso.com

Parse dependent requests

Enable retries for availability test failures.

Test frequency

5 minutes

Test locations
4 location(s) configured

Success criteria

Test Timeout

30 seconds

HTTP response

Status code must equal

200

Content match

Content must contain

Copyright Contoso

[redacted]

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

NOTE:Each correct selection is worth one point.

Answer:

Explanation:

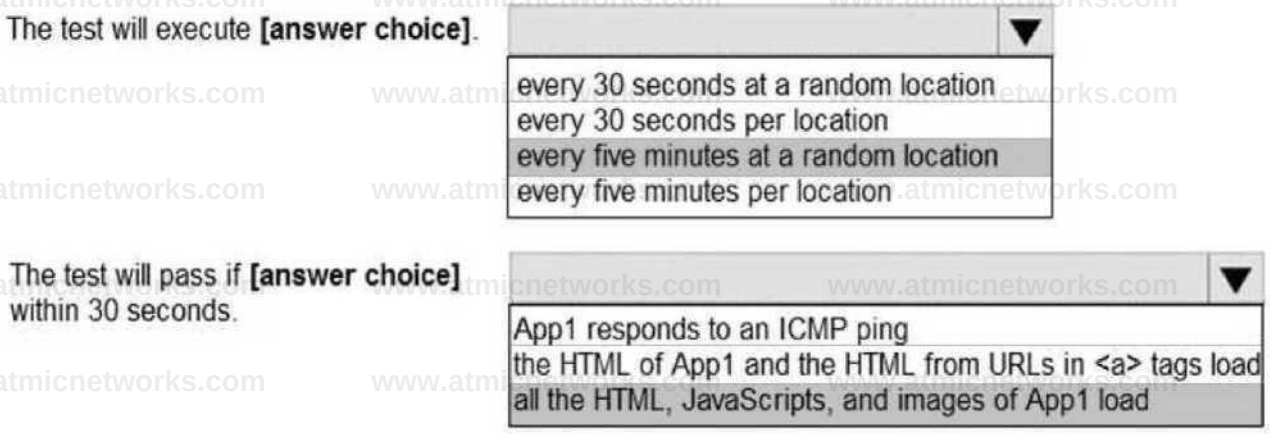

Box 1: every five minutes at a random location

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: 254

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

You use Azure DevOps to build a web app named Appl and deploy Appl to VMSS1. Appl is used heavily and has usage patterns that vary on a weekly basis.

You need to recommend a solution to detect an abnormal rise in the rate of failed requests to Appl.

The solution must minimize administrative effort.

What should you include in the recommendation?

A. an Azure Service Health alert
B. the Failures feature in Azure Application Insights
C. the Smart Detection feature m Azure Application Insights
D. an Azure Monitor alert that uses an Azure Log Analytics query

Answer: C

Explanation:

After setting up Application Insights for your project, and if your app generates a certain minimum amount of data, Smart Detection of failure anomalies takes 24 hours to learn the normal behavior of your app, before it is switched on and can send alerts.

Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/proactive-failure-diagnostics

Question: 255

HOTSPOT
Your company uses a Git 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.

Answer:

Explanation:

Answer Area

Production code Master

Preproduction code Develop

Box 1: Master
The Master branch contains production code. All development code is merged into master in sometime.

Box 2: Develop
The Develop branch containspre-production code. When the features are finished then they are merged into develop.

Reference:
https://medium.com/@patrickporto/4-branching-workflows-for-git-30d0aaee7bf

Question: 256

You plan to create a project in Azure DevOps. Multiple developers will work on the project. The developers will work ofline frequentlyand will require access to the full project history while they are ofline.
Which version control solution should you use?

A. TortotseSVN

B. Team Foundation Version Control
C. Subversion
D. Git

Answer: D

Explanation:

Git history: File history is replicated on the client dev machine and can be viewed even when not connected to the server. You can view history in Visual Studio and on the web portal.

Note: Azure Repos supports two types of version control: Git and Team Foundation Version Control (TFVC).

Reference:
https://docs.microsoft.com/en-us/azure/devops/repos/tfvc/comparison-git-tfvc

Question: 257

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. GitHubApp
B. OAuth
C. personal access token (PAT)
D. Azure Active Directory (Azure AD)

Answer: A

Explanation:

GitHub App uses the Azure Pipelines identity.

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

Question: 258

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 thaiinstalling 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

Answer: B

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: 259

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 thatinstalling JavaScript packages from npm lakes approximately five minutes each time you run the pipeline.

You need to recommend a solution to reduce the pipeline execution time.

Solution: You recommend enabling pipeline caching.

Does this meet the goal?

A. Yes

B. No

Answer: A

Explanation:

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 buildprocess. 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: 260

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 enabling parallel jobs for the pipeline.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead enable pipeline caching.

Note:

npm-cache is a command line utility that cachesdependencies 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: 261

You have the followingAzure policy.

allot: [{ "♦laid*: 'type", "equals' "Microsoft.Storage/storageAccounts*

(

“field*: “Microsoft.Storage/stor ageAccount»/suppert»Mt tpstrafflcOnly "notiquals*: “true*

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

Answer: A

Explanation:

Question: 262

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 Project 1.

You need to link GitHub commits, pull requests, and issues to the work items of Project 1. 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.

Actions

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

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

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

From Project Settings in Azure DevOps, add a GitHub connectiorii

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

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

Answer Area

Answer:

Explanation:

Step 1: From Developer settings in GitHub Enterprise Server, registera 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:
Sign into the web portal for Azure DevOps Services.
Add the GitHub Enterprise Oauth configuration to your organization.
Open Organization settings>Oauth configurations, and choose Add Oauth configuration.
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 openProjects, 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: 263

HOTSPOT
You manage the Git repository for a largeenterprise 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.

Answer Area

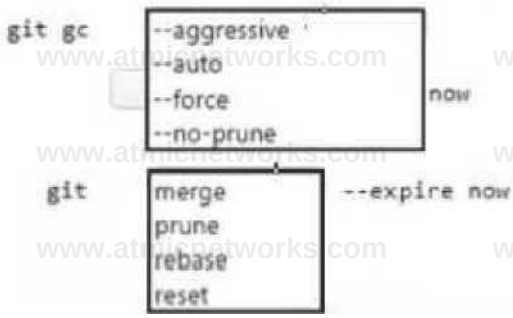

Answer:

Explanation:

Answer Area git gc -aggressive ˅ git prune ^ --expire now

Box 1: --aggressive
Cleanup unnecessary files and optimize the local repository:

get gc --aggressive

Box 2: prune

Prune all unreachableobjects from the object database: git prune

Reference:
https://gist.github.com/Zoramite/2039636

Question: 264

You have an Azure DevOpsorganization named Contoso.
You need to recommend an authentication mechanism that meets the following requirements:

- Supports authentication from Get
- Minimizes the need to provide credentials during authentication What should you recommend?

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

Answer: B

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.

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

Question: 265

You have an Azure subscription. The subscription contains virtual machines that run either

WindowsServer or Linux.
You plan to use Prometheus to monitor performance metrics.
You need to integrate Prometheus and Azure Monitor.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worthone point.

A. Install a Prometheus server on a Windows virtual machine in Azure.
B. On each virtual machine, expose the metrics endpoint.
C. On each virtual machine, enable the Azure Diagnostics extension.
D. On each virtual machine, enable the containerized agent for Azure Monitor.
E. Expose a virtual network service endpoint for Azure Storage.
F. Install a Prometheus server on a Linux virtual machine in Azure.

Answer: A,B

Explanation:

Question: 266

You have an existing build pipeline in Azure Pipelines.
You needto use incremental builds without purging the environment between pipeline executions. What should you use?

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

Answer: B

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.

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

Question: 267

You have the following Azure policy.

```
if: { allof: [ ( "field": "type", "equals": "Microsoft.Storage/storageAccounts"

        { "field": "Microsoft.Storage/stcrageAccounts/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

Answer: B

Explanation:

Denies non HTTPS traffic.

Question: 268

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 areunavailable.
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

Answer: D

Explanation:

Question: 269

You are designing a configuration management solution to support five apps hosted on Azure App Service.
Each app is available inthe 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) permission
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

Answer: A

Explanation:

The Feature Manager in the Azure portal for App Configuration provides a UI for creating and managing the

featureflags that you use in your applications.

App Configuration offers the following benefits:

A fully managed service that can be set up in minutes

Flexible key representations and mappings

Tagging with labels

Point-in-time replay of settings

Dedicated UI forfeature flag management

Comparison of two sets of configurations on custom-defined dimensions

Enhanced security through Azure-managed identities

Encryption of sensitive information at rest and in transit

Native integration with popular frameworks

App Configuration complements Azure Key Vault, which is used to store application secrets.

Reference:

https://docs.microsoft.com/en-us/azure/azure-app-configuration/overview

Question: 270

HOTSPOT

Your company uses Azure DevOps for Get source control.

You have a project in Azure DevOps named Contoso App that contains the following repositories:

https://dev.azure.com/contoso/contoso-app/core-api https://dev.azure.com/contoso/contoso-app/core-spa https://dev.azure.com/contoso/contoso-app/core-db

You need to ensurethat developers receive Slack notifications when there are pull requests created for Contoso App.

What should you run in Slack? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

/azrepos ▼ ▼

feedback signin subscribe subscriptions	https//dev azure com/contoso/contoso-app https://dev.azure.com/contoso/contoso-app/core-api https://dev.azure.com/contoso/contoso-app/core-db https://devazure.com/contoso/contoso-app/core-spa

Answer:

Explanation:

/azrepos ▼ ▼

feedback signin subscribe subscriptions	https//dev azure com/contoso/contoso-app https://dev azure com/contoso/contoso-app/core-api https://dev azure com/contoso/contoso-app/core-db https://dev azure com/contoso/contoso-app/core-spa

Box 1: subscribe
To start monitoring all Get repositories in a project, use the following slash command inside a channel:

/azrepos subscribe [project ur1]

Box 2: https://dev.azure.com/contoso/contoso-app
You can also monitor a specific repository using the following command:

/azrepos subscribe [repository ur1]
The repository URL can be to any page within your repositorythat has your repository name.

For example, for Get repositories, use:

/azrepos subscribe https://dev.azure.com/myorg/myproject/_git/myrepository

Reference:
https://docs.microsoft.com/en-us/azure/devops/repos/integrations/repos-slack

Question: 271
HOTSPOT

You have an Azure Kubernetes Service (AKS) pod.

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

Confirm that the pod is responding to service requests.
Check the status of the pod four times a minute.
Initiate a shutdown if the pod is unresponsive.

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

NOTE:Each correct selection is worth one point.

Answer Area

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

      livenessProbe:
      readinessProbe:
      ShutdownProbe:
      startupProbe:

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

```
initialDelaySeconds: 15
periodSeconds: 15
timeoutSeconds: 15
```

Answer:

Explanation:

```
spec:
     containers:
     - name: container1
       image:    k8s.ger.io/readiness-and—lieveness
       args:
       - /server


        livenessProbe:
        readinessProbe:
        ShutdownProbe:
        startupProbe:
           httpGet:
                path: /checknow
                port:  8123  httpHeaders:  -  name:
                Custom-Header
                  value: CheckNow
                                    ▼
          initialDelaySeconds: 15
            periodSeconds: 15
          timeoutSeconds: 15
```

Box 1: readiness Probe:

For containerized applications that serve traffic, you might want to verify that your container is ready to handle incomingrequests. Azure Container Instances supports readiness probes to include configurations so that your container can't be accessed under certain conditions.

Reference:

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

Question: 272

You plan to onboard 10 new developers.

You need to recommend a development environment that meets the following requirements:

Integrates with GitHub

Provides integrated debugging tools

Supports remote workers and hot-desking environments

Supports developers who use browsers, tablets, and Chromebooks

What should you recommend?

A. VS Code
B. Xamarin Studio
C. Mono Develop
D. Github/Visual Studio Code spaces

Answer: D

Explanation:

Visual Studio Code spaces is built to accommodate the widest variety of projects or tasks, including GitHub and integrating debugging.

Visual Studio Code spaces conceptually and technically extends the Visual Studio Code Remote Developmentextensions.

In addition to "backend" environments, Visual Studio Code spaces supports these "frontend" editors: Visual Studio Code Visual Studio Code-based editor in the browser

Reference:
https://docs.microsoft.com/sv-se/visualstudio/codespaces/overview/what-is-vsonline

Question: 273

Note: This questionis 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.

The lead developer at your company reports that adding new application features takes longer than expected due to a large accumulated technical debt.

You need to recommend changes to reduce the accumulated technical debt.

Solution: You recommend reducing the code coupling and the dependency cycles?

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead reduce the code complexity.

Note: Technical debt is the accumulation of sub-optimal technical decisions made over the lifetime of an application. Eventually, it gets harder and harder to change things: it's the 'sand in the gears'that sees IT initiatives grind to a halt.

Reference:
https://dzone.com/articles/fight-through-the-pain-how-to-deal-with-technical

https://www.devopsgroup.com/blog/five-ways-devops-helps-with-technical-debt/

Question: 274

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

Answer: B

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 nothave 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: 275

HOTSPOT

You have an Azure web app named Webapp1.

You need to use an Azure Monitor query to create a report that details the top 10 pages of Webapp1 that failed.

How should youcomplete the query? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

exceptions
pageViews
requests
traces

I where ▼ duration == 0
itemType == “availabilityResult"
resultcode == "200"
success == false

```
I summarize failedCount=sum(itemCount) by name, resultCode I
top 10 by failedCount desc
I render barchart
```

Answer:

Explanation:

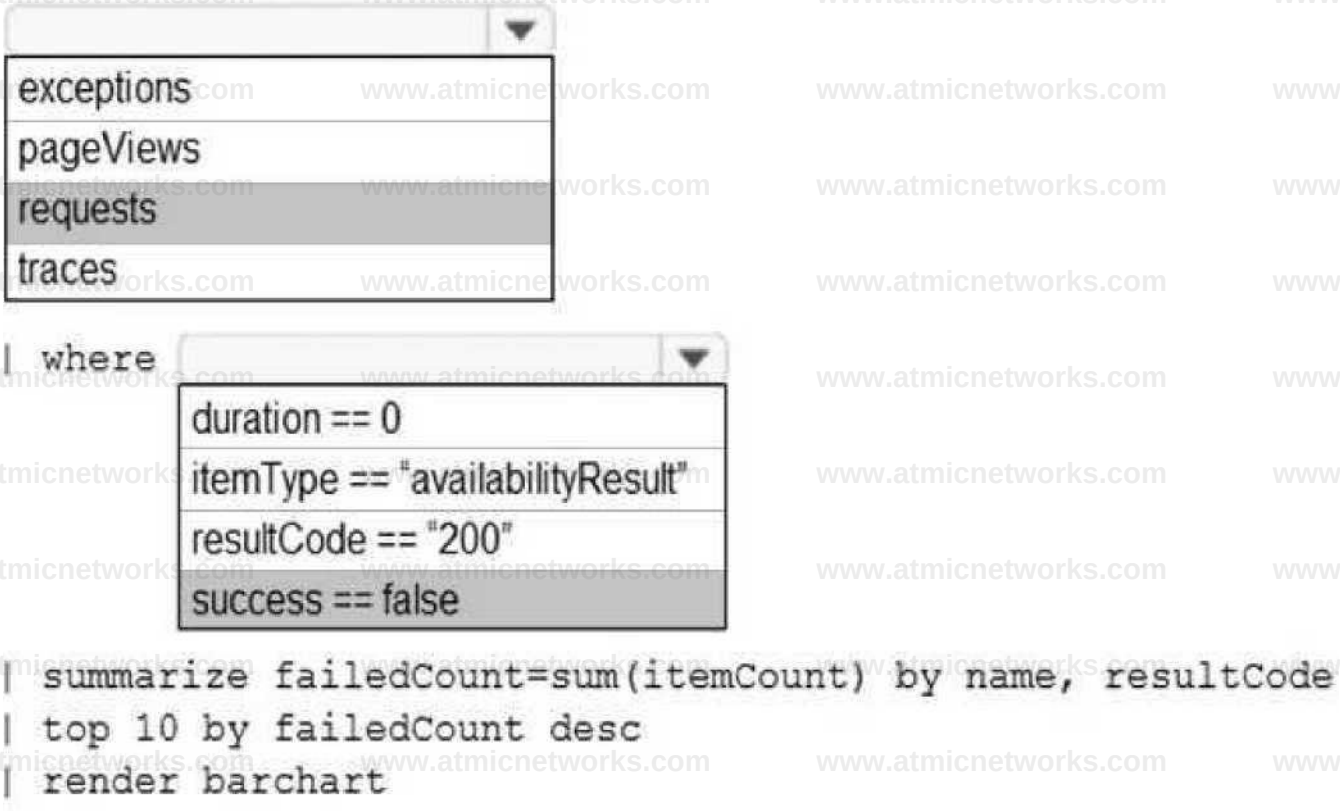

Box 1: requests
Failed requests (requests/failed):
The count of tracked server requests that were marked as failed.
Kusto code:
requests
| where success == 'False'

Box 2: success == false

Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/platform/app-insights-metrics

Question: 276

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 tomeet 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 tothe 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.

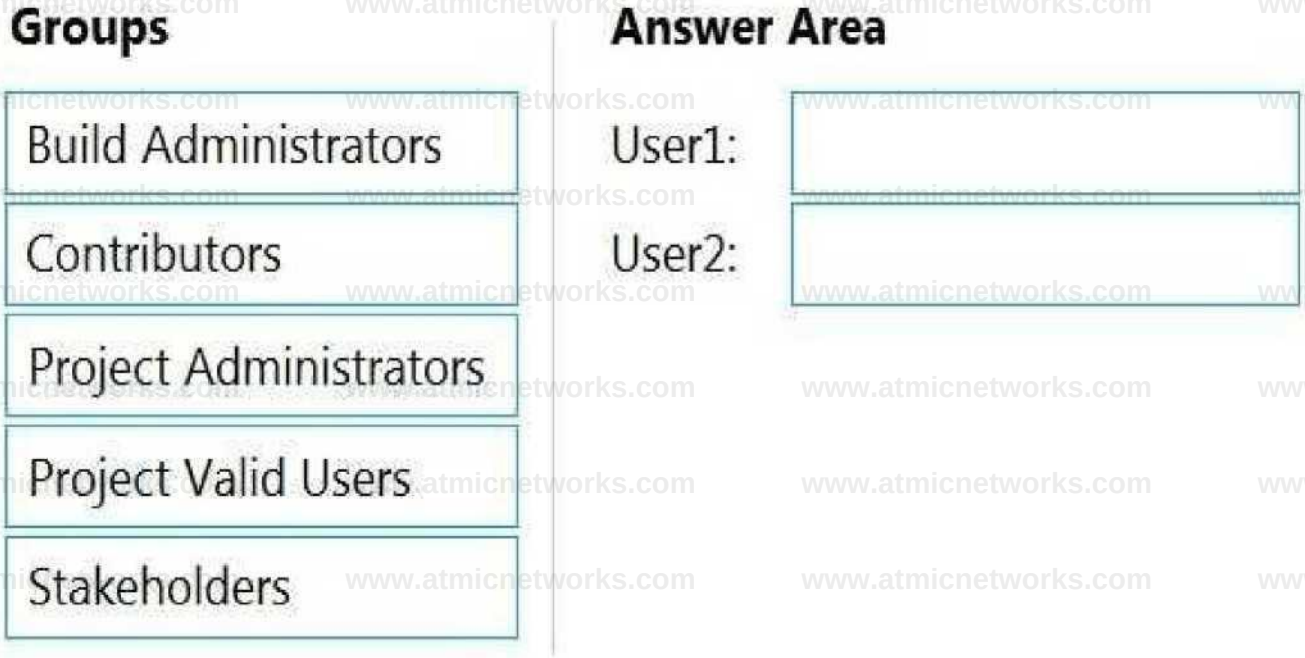

Answer:

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: 277

HOTSPOT

You are finalizing a release in GitHub.

You need to apply the following labels to the release:

Name
Email
Release v3.0
Release date

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

NOTE:Each correct selection is worth one point.

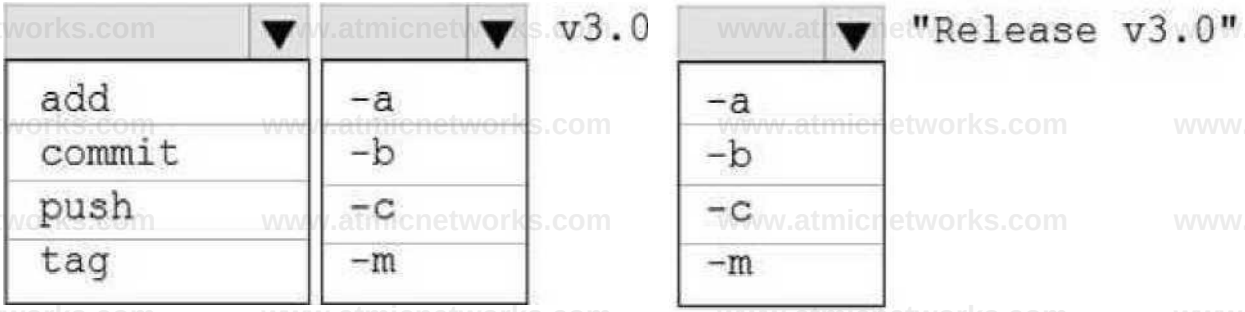

Answer:

Explanation:

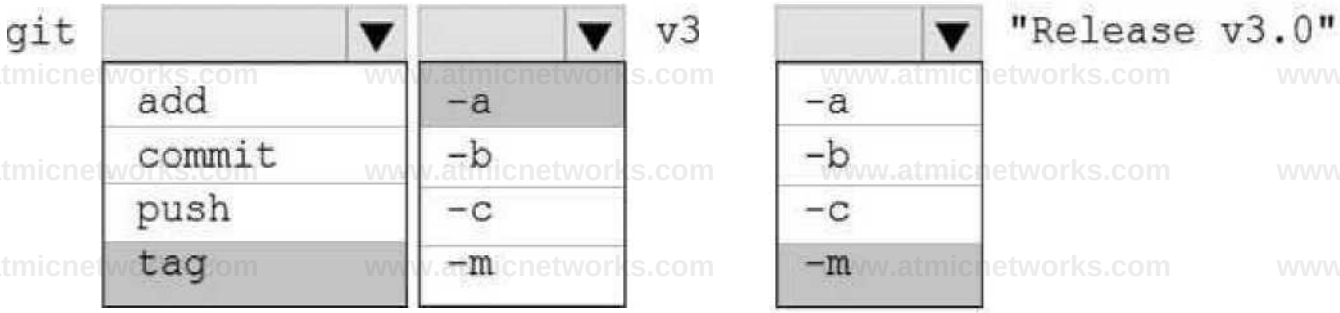

Box 1; tag
Tagging. Like most VCSs, Get has the ability to tag specific points in a repository's history as being important. Typically, people use this functionality to mark release points (v1.0, v2.0 and so on).

Box 2: -a
Creating an annotated tag in Get is simple. The easiest way is to specify -a when you run the tag command:

Example:
$ get tag -a v1.4 -m "my version 1.4"

Box3: -m

Reference:
https://git-scm.com/book/en/v2/Git-Basics-Tagging

Question: 278

HOTSPOT

You are designing YAML-based Azure pipelines for the apps shown in the following table.

You need to configure the YAML strategy value for each app. The solution must minimize app downtime.

Which value shouldyou configure for each app? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

Answer:

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 moreconfidence in the new version, you can release it to more servers in your infrastructure and route more traffic to it.

Reference:
https://docs.microsoft.com/en-us/azure/devops/pipelines/process/deployment-jobs

Question: 279

You have an Azure DevOps project that contains a release pipeline and a Git 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.

Answer: A

Explanation:

Service hooks in Azure DevOps allow you to trigger actions in other tools based on events that occur in your Azure DevOps project. To automatically add release information to work items associated with a Get commit, you would need to configure a service hook that listens for commit events in your Get repository, and then sends the release information to the appropriate work items.
Here's the steps you can follow to set up a service hook for this purpose:
In your Azure DevOps project, navigate to the projectsettings by clicking on the gear icon in the top right corner of the page.
Select "Service Hooks" from the left-hand menu.
Click on the "New Subscription" button to create a new service hook.
In the "Event" drop-down menu, select "Code pushed" to triggerthe service hook when a new code revision is committed to the repository.
In the "Actions" section, select the action that you want to take place when the service hook is triggered. For example, you might use the "Link work items to commits" action to automatically associate work items with the relevant commits.
Configure the remaining settings as needed, and then click on the "Create" button to create the service hook.
You can find more information on Service hooks in Azure DevOps by following this

linkhttps://docs.microsoft.com/en-us/azure/devops/service-hooks/overview?view=azure-devops

Question: 280

Your company uses Azure DevOps for the build pipelines and deployment pipelines of Java-based projects.

You need to recommend a strategy for managing technical debt.

Which action should you include in the recommendation?

A. Configure post-deployment approvals in the deployment pipeline.
B. Integrate Azure DevOps and Sonar Rube.
C. IntegrateAzure DevOps and Azure Dev Test Labs.

Answer: B

Explanation:

You can manage technical debt with Sonar Rube and Azure DevOps.

Note: Technical debt is the set of problems in a development effort that make forward progress on customer value inefficient. Technical debt saps productivity by making code hard to understand, fragile, time-consuming to change, difficult to validate, and creates unplanned work that blocks progress. Unless they are managed, technical debt can accumulate and hurt the overall qualityof the software and the productivity of the development team in the long term

SonarQube an open source platform for continuous inspection of code quality to perform automatic reviews with static analysis of code to:

Detect Bugs
Code Smells
Security Vulnerabilities
Centralize Quality
What's covered in this lab

Reference:
https://azuredevopslabs.com/labs/vstsextend/sonarqube/

Question: 281

You have an Azure DevOps organizationthat contains a project named Project1.
You need to create a published wiki in Project1.
What should you do first?

A. Modify the Storage settings of Project1.
B. In Project1, create an Azure DevOps pipeline.
C. In Project1, create an Azure DevOps repository.
D. Modify the Team configuration settings of Project1.

Answer: C

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/devops/project/wiki/publish-repo-to-wiki?view=azure-devops&tabs=browser

Question: 282

You have a project in Azure DevOps named Project1. Project1 contains a published wiki.

You need to change the order of pages in the navigation pane of the published wiki in the Azure DevOps portal.

What should you do?

A. At the root of the wiki, create a file named order that defines the page hierarchy.
B. At the root of the wiki, create a file named wiki.md that defines the page hierarchy.
C. Rename the pages in the navigation pane.
D. Drag and drop the pages in the navigation pane.

Answer: B

Explanation:

Reorder a wiki page

You can reorder pages withinthe wiki tree view to have pages appear in the order and hierarchy you want.

You can drag-and-drop a page title in the tree view to do the following operations: Change the parent-child relationship of a page

Change the order of the page within the hierarchy

Reference:
https://docs.microsoft.com/en-us/azure/devops/project/wiki/add-edit-wiki

Question: 283

DRAG DROP

You need to recommend project metrics for dashboards in Azure DevOps.

Which chart widgets should you recommend for each metric? To answer,drag the appropriate chart widgets to the correct metrics. Each chart widget 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.

Burndown

Cycle Time

Lead Tiime

Velocity

The elapsed time from the creation of work items to their completion: [ ]

The elapsed time to complete work items once they are active: [ ]

The remaining work: [ ]

Answer:

Explanation:

Burndown

Cycle Time

Lead Time

Velocity

The elapsed time from the creation of work items to their completion: Lead Time

The elapsed time to complete work items once they are active: Cycle Time

The remaining work: Burndown

Box 1: Lead time

Lead time measures the total time elapsed from the creation of work items to their completion.

Box 2: Cycle time

Cycle time measures the time it takes for your team to complete work items once they begin actively working on them.

Box 3: Burn down

Burn down charts focus on remaining workwithin a specific time period.

Incorrect Answers:

A: guide for determining how well the team estimates and meets their planned commitments

Reference:

https://docs.microsoft.com/en-us/azure/devops/report/dashboards/velocity-guidance?view=vsts

https://docs.microsoft.com/en-us/azure/devops/report/dashboards/cycle-time-and-lead- time?view=vsts

https://docs.microsoft.com/en-us/azure/devops/report/dashboards/configure-burndown-burnup- widgets?view=vsts

Question: 284

HOTSPOT

You plan to create alerts that will be triggered based on the page load performance of a home page.
You have the Application Insights log queryshown in the following exhibit.

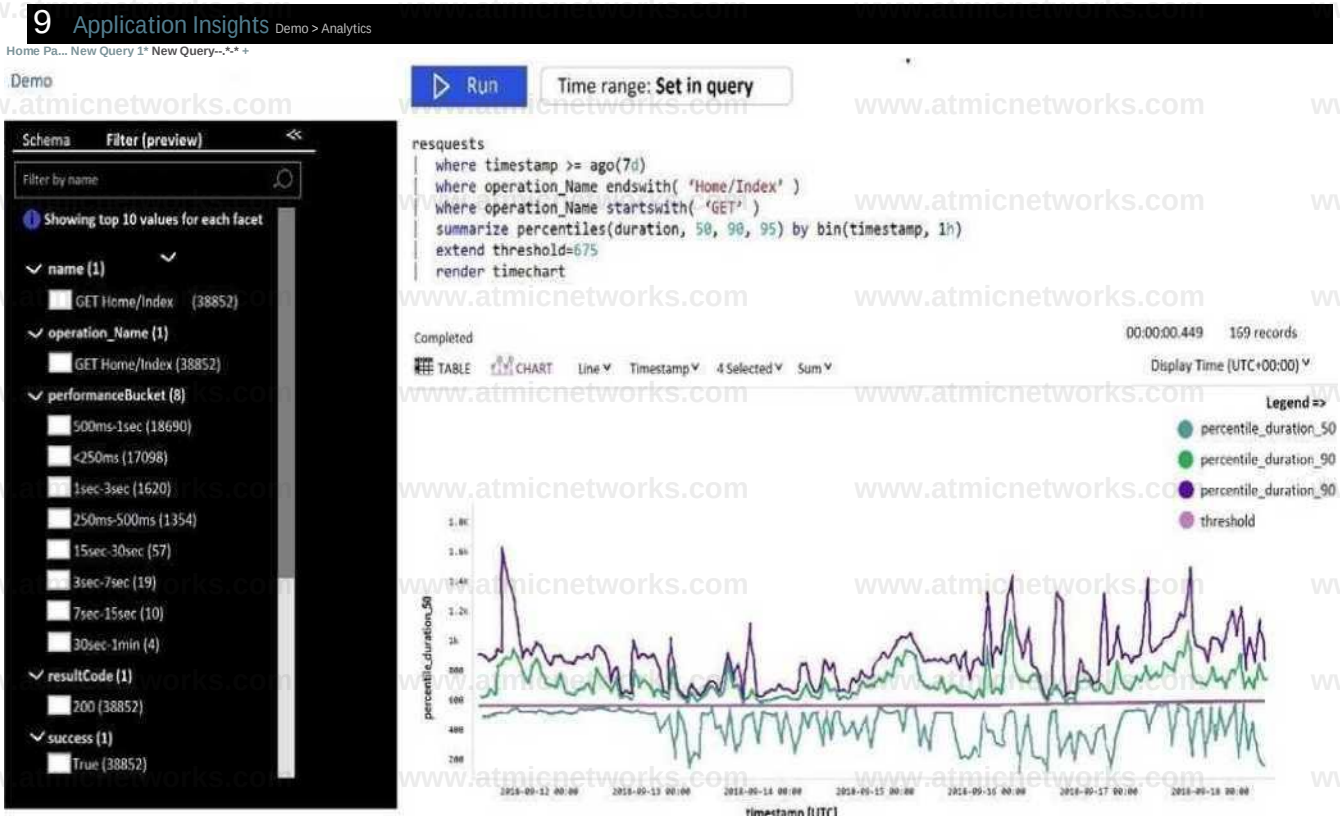

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

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

- percentile_duration_50
- percentile_duration_90
- percentile_duration_95
- threshold

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

- item Type
- resultCode
- source
- success

Answer:

Explanation:

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

percentile_duration_5O
percentile duration 90
percentile_duration_95
threshold

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

item Type
resultcode
source
success

Box 1: percentile_duration_95
Box 2: success
For example – requests
| project name, url, success
| where success == "False"
This will return all the failed requests in my App Insights within the specified time range. Reference:
https://devblogs.microsoft.com/premier-developer/alerts-based-on-analytics-query-using-custom- log-search/

Question: 285

You manage an Azure web app that supports an e-commerce website.
You need to increase the logging level when the web app exceeds normal usage patterns. The solution must minimize administrative overhead.
Which two resources should you include in the solution? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. an Azure Automation run book
B. an Azure Monitor alert that has a dynamic threshold
C. an Azure Monitor alert that has a static threshold
D. the Azure Monitor auto scale settings
E. an Azure Monitor alertthat uses an action group that has an email action

Answer: A,B

Explanation:

A: You can use Azure Monitor to monitor base-level metrics and logs for most services in Azure. You can call Azure Automation run books by using action groups or by using classicalerts to automate tasks based on alerts.
B: Metric Alert with Dynamic Thresholds detection leverages advanced machine learning (ML) to learn metrics' historical behavior, identify patterns and anomalies that indicate possible service issues. It provides support of both a simple UI and operations at scale by allowing users to configure alert rules through the

Azure Resource Manager API, in a fully automated manner.

Reference:

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

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

Question: 286

HOTSPOT

You have an Azure Kubernetes Service (AKS) pod.

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

Confirm that thepod is responding to service requests.

Check the status of the pod four times a minute.

Initiate a shutdown if the pod is unresponsive.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

livenessProbe:
readinessProbe:
ShutdownProbe:
startupProbe:

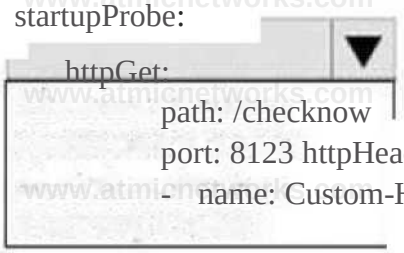

initialDelaySeconds: 15
periodSeconds: 15
timeoutSeconds: 15

Answer:

Explanation:

```
apiVersion: vl kind: Pod metadata: labels:
                test: readiness-and-liveness
      name: readiness-http
spec:
      containers:
      - name: containerl
        image: k8s.ger.io/readines s-and-livenes s args : -
        /server
```

```
        livenes s Probe:
        readiness Probe:
        ShutdownProbe:
        startupProbe:

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

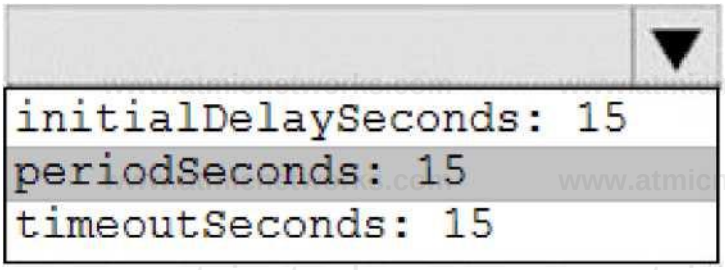

Box 1: readiness Probe:

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.

Reference:

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

Question: 287

You have a Microsoft ASP.NET Core web app in Azure that is accessed worldwide.

You need to run a URL ping test once every five minutes and createan alert when the web app is unavailable from specific Azure regions. The solution must minimize development time.

What should you do?

A. Create an Azure Monitor Availability metric and alert.
B. Create an Azure Application Insights availability test andalert.
C. Write an Azure function and deploy the function to the specific regions.
D. Create an Azure Service Health alert for the specific regions.

Answer: B

Explanation:

There are three types of Application Insights availability tests:
URL ping test: a simple test that you can create in the Azure portal.
Multi-step web test
Custom Track Availability Tests
Note: After you've deployed your web app/website, you can set up recurring tests to monitor availability and responsiveness. Azure Application Insightssends web requests to your application at regular intervals from points around the world. It can alert you if your application isn't responding, or if it responds too slowly.
You can set up availability tests for any HTTP or HTTPS endpoint that is accessible from the public internet. You don't have to make any changes to the website you're testing. In fact, it doesn't even have to be a site you own. You can test the availability of a REST API that your service depends on. Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability#create-a- url-ping-test

Question: 288

You have a multi-tier application. The front end of the application is hosted in Azure App Service.
You need to identify the average load times of theapplication pages.
What should you use?

A. Azure Application Insights
B. the activity log of the App Service
C. the diagnostics logs of the App Service
D. Azure Advisor

Answer: A

Explanation:

Application Insights will tell you about any performance issuesand exceptions, and help you find and diagnose the root causes.
Application Insights can monitor both Java and ASP.NET web applications and services, WCF services.
They can be hosted on-premises, on virtual machines, or as Microsoft Azure websites.
On theclient side, Application Insights can take telemetry from web pages and a wide variety of devices including iOS, Android, and Windows Store apps.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/web-monitor-performance

Question: 289

Your company uses Service Now for incident management.
You develop an application that runs on Azure.
The company needs to generate a ticketin Service Now when the application fails to authenticate.
Which Azure Log Analytics solution should you use?

A. Application Insights Connector
B. Automation & Control
C. IT Service Management Connector (ITSM)
D. Insight & Analytics

Answer: C

Explanation:

The IT Service Management Connector (ITSMC) allows you to connect Azure and a supported IT Service Management (ITSM) product/service.
ITSMC supports connections with the following ITSM tools:
ServiceNow
System Center Service Manager
Provance
Cherwell
WithITSMC, you can
Create work items in ITSM tool, based on your Azure alerts (metric alerts, Activity Log alerts and Log Analytics alerts).
Optionally, you can sync your incident and change request data from your ITSM tool to an Azure Log Analytics workspace.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/platform/itsmc-overview

Question: 290

HOTSPOT

Your company is building a new web application.
You plan to collect feedback from pilot users on the features being delivered.
All the pilotusers have a corporate computer that has Google Chrome and the Microsoft Test & Feedback extension installed. The pilot users will test the application by using Chrome.
You need to identify which access levels are required to ensure that developers can request and gather feedback from the pilot users. The solution must use the principle of least privilege.
Which access levels in Azure DevOps should you identify? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Developers:

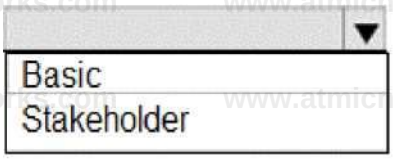

Pilot

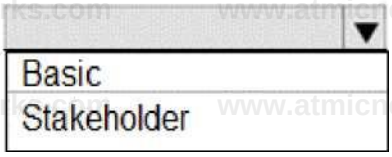

Answer:

Explanation:

Box 1: Basic

Assign Basic to users with a TFS CAL, with a Visual Studio Professional subscription, and to users for whom you are paying for Azure Boards & Repos in an organization.

Box 2: Stakeholder

Assign Stakeholders to users with no license or subscriptions who need access to a limited set of features.

Note:

You assign users or groups of users to oneof the following access levels:

Basic: provides access to most features

VS Enterprise: provides access to premium features

Stakeholders: provides partial access, can be assigned to unlimited users for free

Reference:

https://docs.microsoft.com/en-us/azure/devops/organizations/security/access-levels?view=vsts

Question: 291

You use Azure SQL Database Intelligent Insights and Azure Application Insights for monitoring.

You need to write ad-hoc queries against the monitoring data.

Which query language shouldyou use?

A. Kusto Query Language (KQL)
B. PL/pgSQL
C. PL/SQL
D. Transact-SQL

Answer: A

Explanation:

Azure Monitor Logs is based on Azure Data Explorer, and log queries are written using the same Kusto query language (KQL). This is a rich language designedto be easy to read and author, and you

should be able to start using it with minimal guidance.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/log-query-overview

Question: 292

Your company creates a web application.

You needto recommend a solution that automatically sends to Microsoft Teams a daily summary of the exceptions that occur in the application.

Which two Azure services should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Logic Apps
B. Azure Pipelines
C. Microsoft Visual Studio App Center
D. Azure DevOps Project
E. Azure Application Insights

Answer: A,E

Explanation:

F. Exceptions in your live web app are reported by Application Insights.

Note: Periodical reports help keep a team informed on how their business critical services are doing. Developers, DevOps/SRE teams, and their managers can be productive with automated reports reliably

delivering insights without requiring everyone to signin the portal. Such reports can also help identify gradual increases in latencies, load or failure rates that may not trigger any alert rules. A: You can programmatically query Application Insights data to generate custom reports on a schedule. The following options can help you get started quickly:

Automate reports with Microsoft Flow

Automate reports with Logic Apps

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/asp-net-exceptions
https://docs.microsoft.com/en-us/azure/azure-monitor/app/automate-custom-reports

Question: 293

DRAG DROP

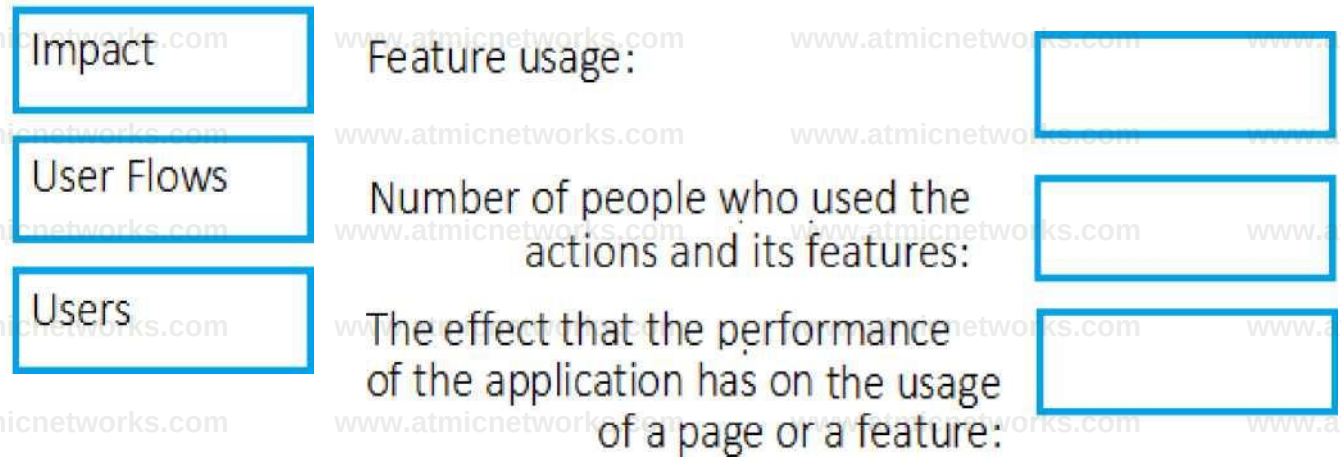

Your company wants to use Azure Application Insights to understand how user behaviors affect an application.
Which Application Insights tool should you use to analyze each behavior? To answer, drag the appropriate tools to the correct behaviors. Each tool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Box 1: User Flows
The User Flows tool visualizes how users navigate between the pages and features of your site. It's great for answering questions like:
How do users navigate away from a page on your site?
What do users click on a page on your site?
Where are the places that users churn most from your site?
Are there places where users repeat the same action over and over?
Box 2: Users
Counting Users: The user behavior analytics tools don't currently support counting users or sessions based on properties other than anonymous user ID, authenticated user ID, or session ID.
Box 3: Impact
Impact analyzes how load times and other properties influence conversion rates for various parts of your app. To put it more precisely, it discovers how any dimension of a page view, custom event, or request affects the usage of a differentpage view or custom event.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-flows
https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-impact https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-troubleshoot

Question: 294

You have an Azure DevOps project named Project1 and an Azure subscription named Sub1. Sub1 contains an Azure virtual machine scale set named VMSS1. VMSS1 hosts a web application named

WebApp1. WebApp1 uses state full sessions.

The WebApp1installation is managed by using the Custom Script extension. The script resides in an Azure Storage account named sa1.

You plan to make a minor change to a UI element of WebApp1 and to gather user feedback about the change.

You need to implement limiteduser testing for the new version of WebApp1 on VMSS1.

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

NOTE: Each correct selection is worth one point.

A. Modify the load balancer settings of VMSS1.
B. Redeploy VMSS1.
C. Upload a custom script file to sa1.
D. Modify the Custom Script extension settings of VMSS1.
E. Update the configuration of a virtual machine in VMSS1.

Answer: B,C,D

Explanation:

Question: 295

You are monitoring the health and performance of an Azure web app by using Azure Application Insights.

You need to ensure that an alert is sent when the web app has a sudden rise in performance issues and failures.

What shouldyou use?

A. custom events
B. Application Insights Profiler
C. usage analysis
D. Smart Detection
E. Continuous export

Answer: D

Explanation:

Smart Detection automatically warns you of potential performance problems and failure anomalies in your web application. It performs proactive analysis of the telemetry that your app sends to Application

Insights. If there is a sudden rise in failure rates, or abnormal patterns in client or server

performance, you get an alert.

Reference:

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

Question: 296

You are integrating Azure Pipelines and Microsoft Teams.
You install the Azure Pipelines app in Microsoft Teams.
You have an Azure DevOps organization named Contoso that contains aproject name Project1.
You subscribe to Project1 in Microsoft Teams.
You need to ensure that you only receive events about failed builds in Microsoft Teams.
What should you do first?

A. From Microsoft Teams, run @azure pipelines subscribehttps://dev.azure.com/Contoso/Project1.
B. From Azure Pipelines, add a Publish Build Artifacts task to Project1.
C. From Microsoft Teams, run @azure pipelines subscriptions.
D. From Azure Pipelines, enable continuous integration for Project1.

Answer: A

Explanation:

To start monitoring all pipelines in a project, use the following command inside a channel: @azure pipelines subscribe [project ur1]
The project URL can be to any page within your project (except URLs to pipelines).
For example:
@azure pipelinessubscribe https://dev.azure.com/myorg/myproject/
Reference:
https://docs.microsoft.com/en-us/azure/devops/pipelines/integrations/microsoft-teams

Question: 297

You have an Azure DevOps organization named Contoso.
You need to receive Microsoft Teams notifications when work items are updated.
What should you do?

A. From Azure DevOps, configure a service hook subscription
B. From Microsoft Teams, configure a connector
C. From the Microsoft Teams admin center, configure external access
D. From Microsoft Teams, add a channel
E. From Azure DevOps, install an extension

Answer: A

Explanation:
Service hooks let you run tasks on other services when events happen in your Azure DevOps projects. For example, create a card in Trello when a work item is created or send apush notification to your team's mobile devices when a build fails. You can also use service hooks in custom apps and services as a more efficient way to drive activities when events happen in your projects.

Note: Service hook publishers define a set of events. Subscriptions listen for the events and define actions to take based on the event. Subscriptions also target consumers, which are external services that can run their own actions, when an event occurs.

Reference:

https://docs.microsoft.com/en-us/azure/devops/service-hooks/overview

Question: 298

You create an alert rule in Azure Monitor as shown in the following exhibit.

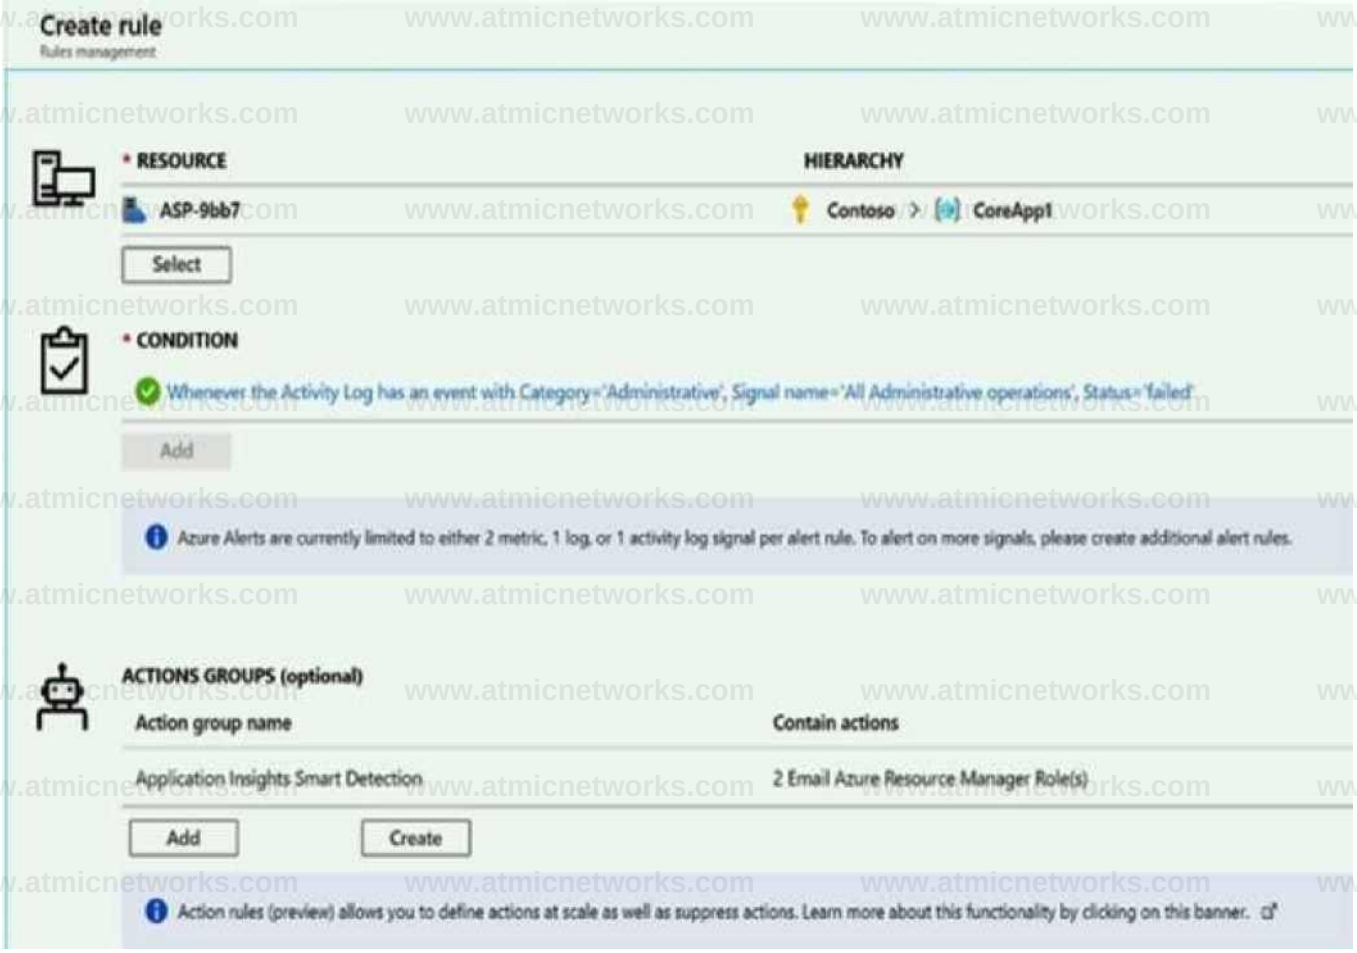

Which action will trigger an alert?

A. a failed attempt to delete the ASP-9bb7 resource
B. a change to a role assignment forthe ASP-9bb7 resource
C. a successful attempt to delete the ASP-9bb7 resource
D. a failed attempt to scale up the ASP-9bb7 resource

Answer: A

Explanation:

Question: 299

You have a web app hosted on Azure App Service. The web app stores data in an Azure SQL database. You need to generate an alert when there are 10,000 simultaneous connections to the database. The solution must minimize development effort.

Which option should you select in the Diagnostics settings of the database?

A. Send to Log Analytics
B. Stream toan event hub
C. Archive to a storage account

Answer: A

Explanation:

ENABLE DIAGNOSTICS TO LOG ANALYTICS
This configuration is done PER DATABASE

1. Click on Diagnostics Settings and then Turn On Diagnostics

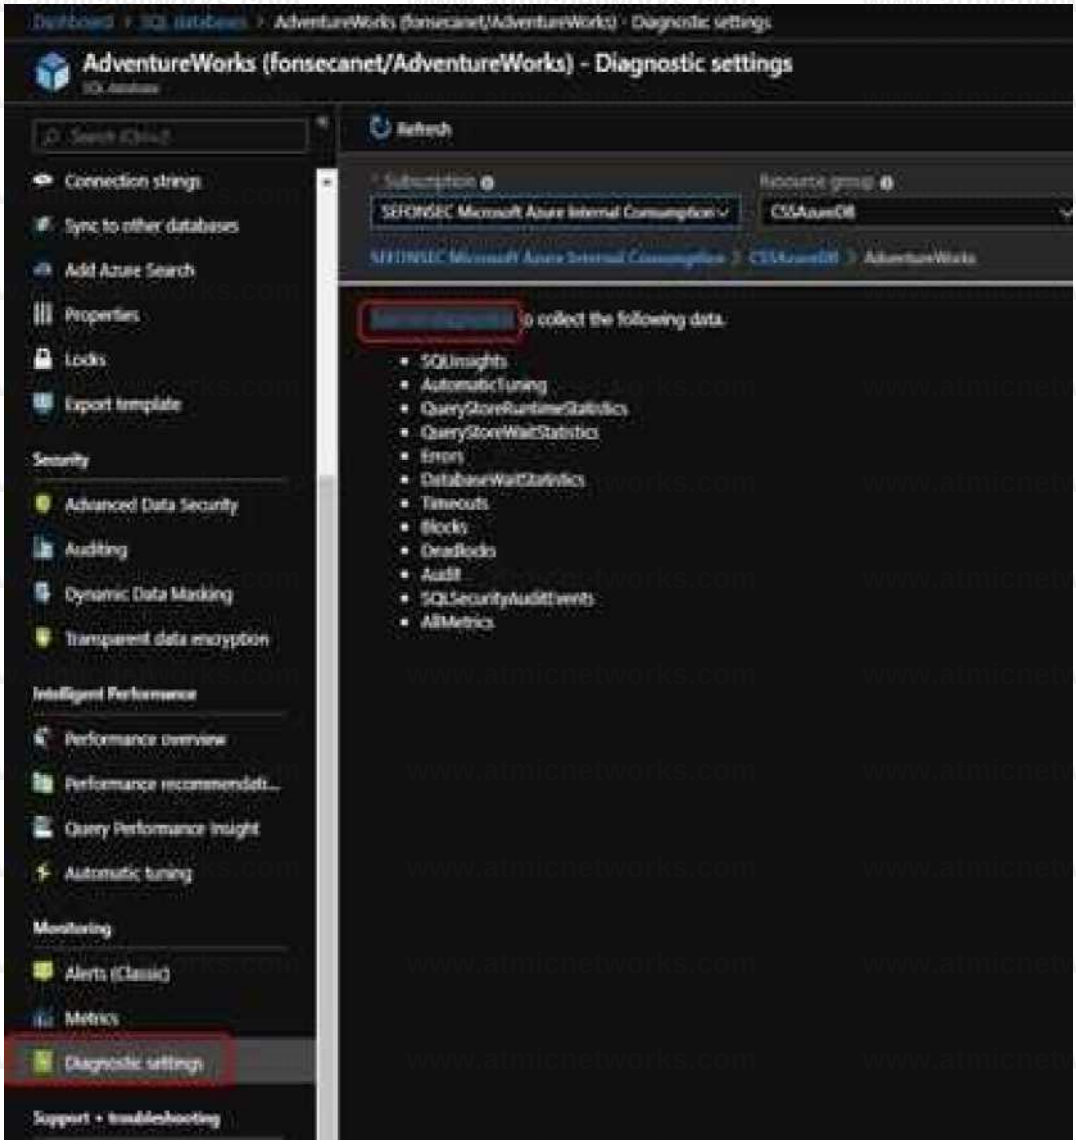

2. Select to Send to Log Analytics and select the Log Analytics workspace. For this sample I will selected only Errors

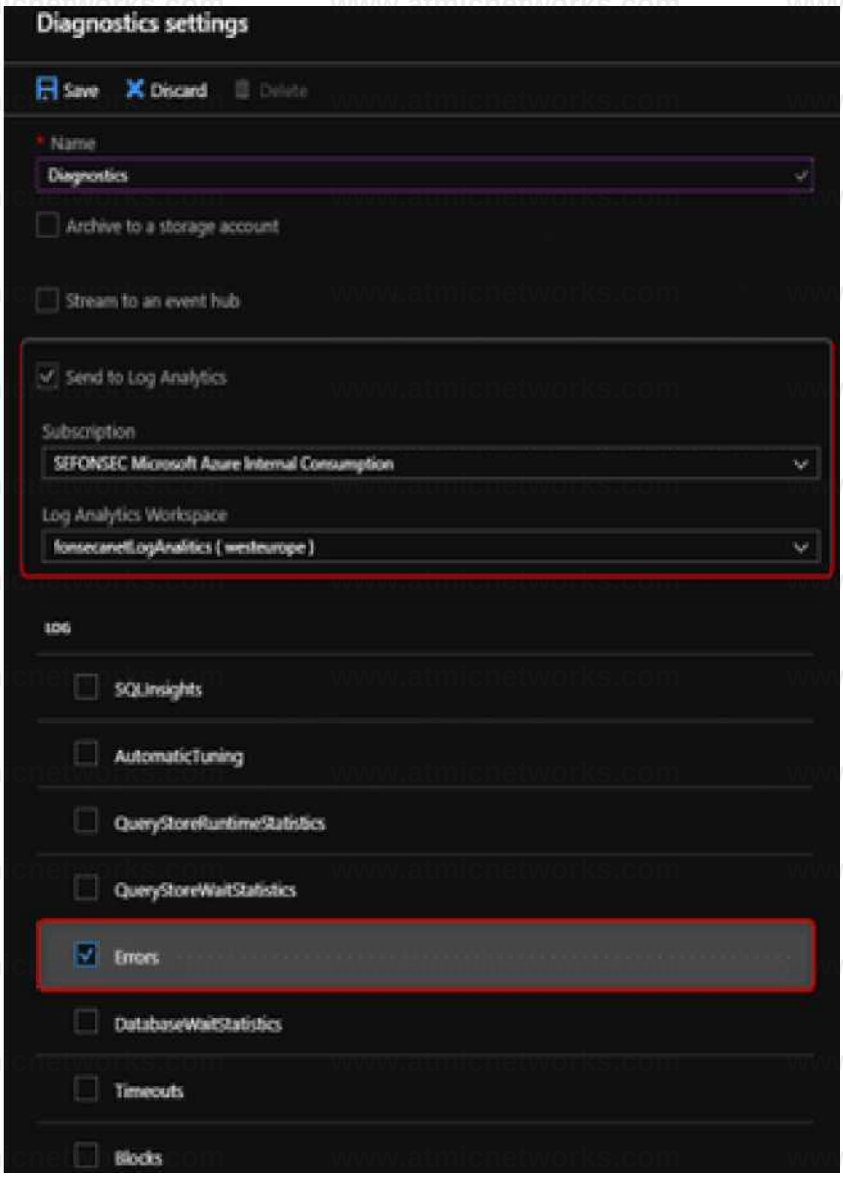

Reference:
https://techcommunity.microsoft.com/t5/azure-database-support-blog/azure-sql-db-and-log- analytics-better-together-part-1/ba-p/794833

Question: 300

You use GitHubfor source control.
A file that contains sensitive data is committed accidentally to the Get repository of a project.
You need to delete the file and its history form the repository.
Which two tools can you use? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. the git filter-branch command
B. BFG Repo-Cleaner
C. the git rebase command
D. GitHub Desktop

Answer: A,B

Explanation:

To entirely remove unwanted files from a repository's history you can use either the git filter-branch command or the BFG Repo-Cleaner open source tool.

Reference:
https://docs.github.com/en/github/authenticating-to-github/keeping-your-account-and-data-secure/removing-sensitive-data-from-a-repository

Question: 301

Your company uses GitHub for source control. The company has a team that performs code reviews.

You need to automate the assignment of the code reviews. The solution must meet the following requirements:

Prioritize the assignment of code reviews to team members who havethe fewest outstanding assignments.

Ensure that each team member performs an equal number of code reviews in any 30-day period.

Prevent the assignment of code reviews to the team leader.

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

NOTE: Each correct selection is worth one point.

A. Clear Never assign certain team members.
B. Select If assigning team members, don't notify the entire team.
C. Select Never assign certain team members.
D. Set Routing algorithm toRound robin.
E. Set Routing algorithm to Load balance.

Answer: A,E

Explanation:

A: To always skip certain members of the team, select Never assign certain team members. Then, select one or more team members you'd like to always skip. In this case select the team leader. E: The load balance algorithm chooses reviewers based on each member's total number of recent review requests and considers the number of outstanding reviews for each member. The load balance algorithm tries to ensure that each team memberreviews an equal number of pull requests in any 30day period.

Incorrect Answers:

D: The round robin algorithm chooses reviewers based on who's received the least recent review request, focusing on alternating between all members of the team regardless of the number of outstanding reviews they currently have.

Reference:
https://docs.github.com/en/organizations/organizing-members-into-teams/managing-code-review-assignment-for-your-team

Question: 302

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.

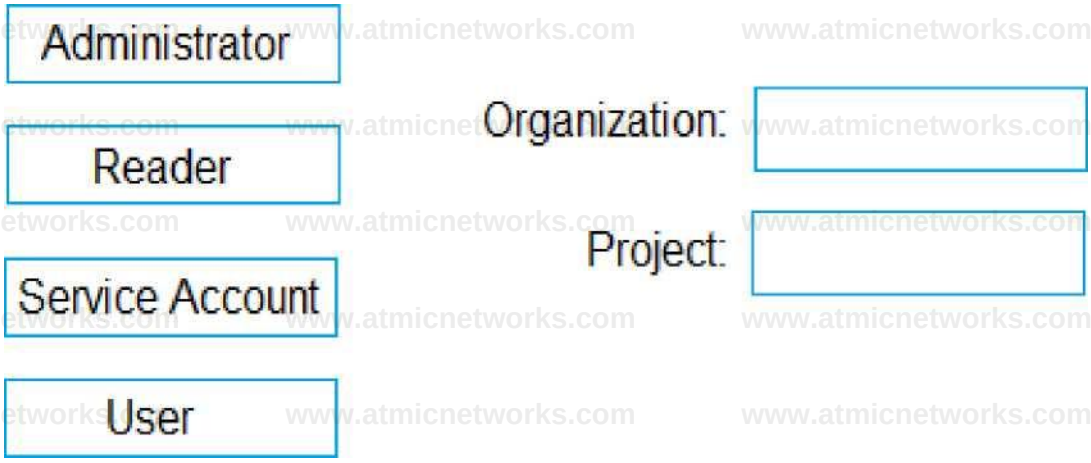

Answer:

Explanation:

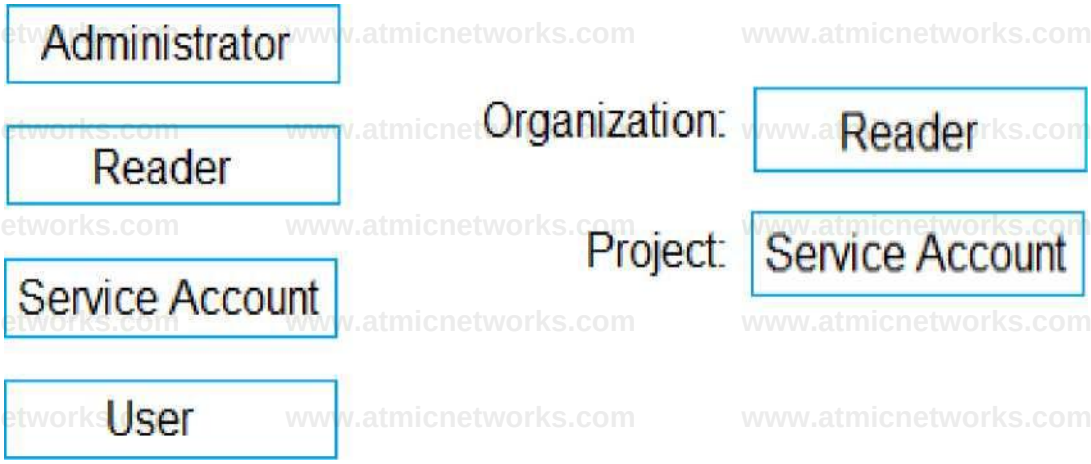

Box 1: Reader

Members of the Reader role can view the organizationagent pool as well as agents. You typically use this to add operators that are responsible for monitoring the agents and their health.

Box 2: Service account

Members of the Service account role can use the organization agent pool to create a project agentpool in a project. If you follow the guidelines above for creating new project agent pools, you typically do not have to add any members here.

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

Question: 303

DRAG DROP

You have a project in Azure DevOps named Project1 that contains two Azure DevOps pipelines named Pipeline1 and Pipeline2.

You need to ensure that Pipeline1 can deploy code successfully to an Azure web app named webapp1. The solution must ensure that Pipeline2 does not have permission to webapp1. 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.

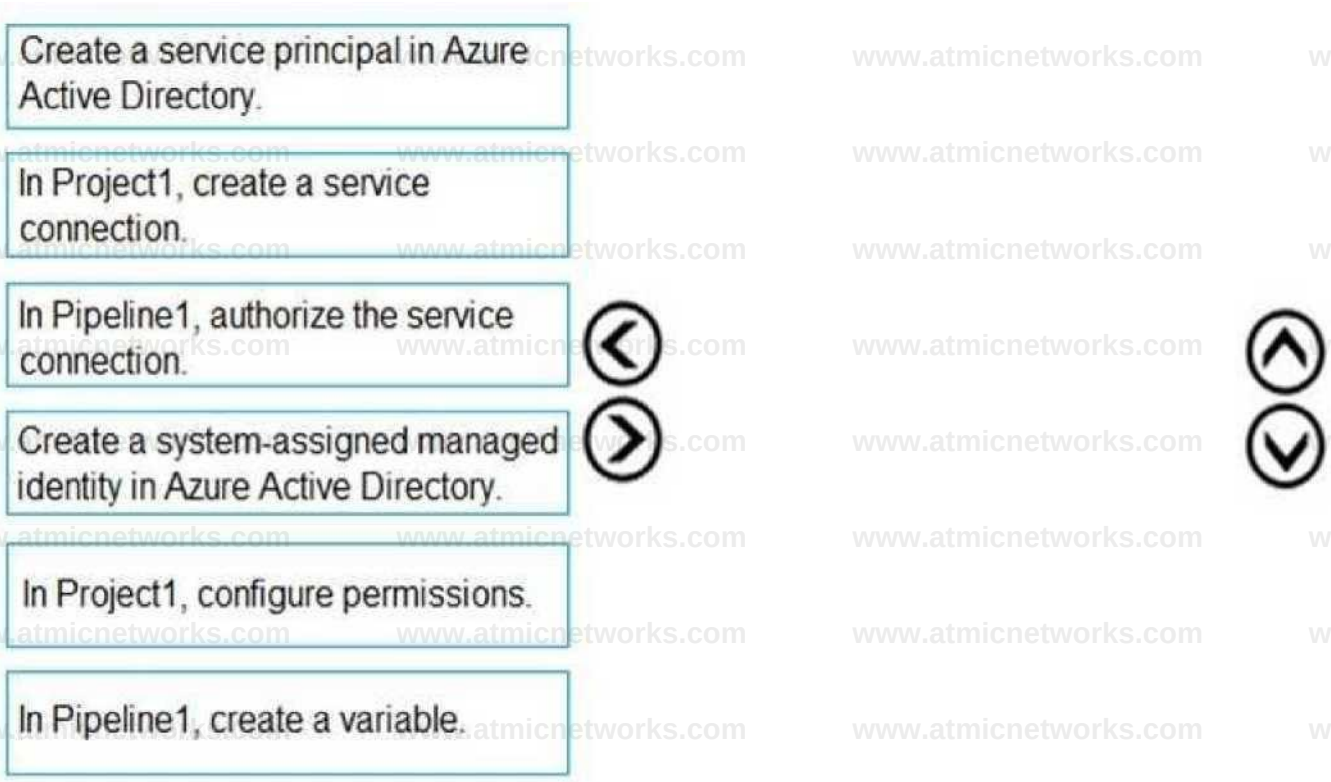

Answer:

Explanation:

Create a service principal in Azure Active Directory,

In Projectl, create a service connection.

In Pipelinel, create a vanable.

In Pipelinel, authorize the service connection.

Create a system-assigned managed identity in Azure Active Directory

In Projectl, configure permissions.

Reference:
https://docs.microsoft.com/en-us/azure/devops/pipelines/library/connect-to-azure?view=azure- devops

Question: 304

DRAG DROP

You use Azure Pipelines to automate Continuous Integration/Continuous Deployment (CI/CD) for an Azureweb 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.

Actions

- 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.

Answer Area

Answer:

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: 305

Note: This question is part ofa 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 anAzure virtual machine scale set named VMSS1 that is configured for auto scaling.

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

Answer: B

Explanation:

Question: 306

DRAG DROP

You have a GitHub organization named org1 and an Azure tenant named Tenant1.

You need to enable single sign-on (SSO) in Azure Active Directory (Azure AD) for the users in org1.

Which URIs should you use for the SAML configuration in Azure AD? To answer, drag the appropriate URIs to the correct settings. Each URImay 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.

URIS

https //github com/orgs/org1

, ,i;
https://gitnub.com/orgs/org1/sso

https://login microsoftonline com/tenantI

https://github.com/orgs/org1/saml/consume

https//login microsoftonline comtenantI com

Answer Area

Identifier (Entity ID): _____

Reply URL (Assertion Consumer Service URL):

Sign on URL

Answer:

Explanation:

Reference:
https://docs.microsoft.com/en-us/azure/active-directory/saas-apps/github-tutorial

Question: 307

You have an AzureDevOps organization named Contoso and an Azure subscription. The subscription contains an Azure virtual machine scale set named VMSS1 that is configured for auto scaling.

You have a project m Azure DevOps named Project 1. Project! is used to build a web app named App1 and deploy App1 to VMSSt.

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

Solution: From Azure Monitor, configure the auto scale settings.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Question: 308

You have an Azure solution that contains a build pipeline in Azure Pipelines. You experience intermittent delays before the build pipeline starts. You need to reduce the time it takes to start the build pipeline. What should you do?

A. Split the build pipeline into multiple stages.
B. Purchase an additional parallel job.
C. Create a new agent pool.
D. Enable self-hosted build agents.

Answer: B

Explanation:

https://docs.microsoft.com/en-us/azure/devops/pipelines/troubleshooting/troubleshooting

Question: 309

HOTSPOT

You plan to use Desired State Configuration (DSC) to maintain the configuration state of virtual machines that run Windows Server.

You needto perform the following:

Install Internet Information Services (IIS) on the virtual machines.

Update the default home page of the IIS web server.

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

NOTE:Each correct selection is worth one point.

Answer Area

```
Configuration WebServerConfig {
Import-DscResource -ModuleName PsDesiredStateConfiguration
Node 'localhost' {

                    ▼ Webserver {
```

Service
WindowsFeature
WindowsOptionalFeature
WindowsProcess

```
        Ensure = "Present"
        Name = "Web-Server"


    ▼ DefaultHomePage {
```

Archive
File
Package
Script

```
        Ensure = 'Present'
        SourcePath = 'Wserverl
\DSCResources\web\index.htm'
        Destinationpath = ’c:\inetpub\wwwroot'
      J
   I
}
```

Answer:

Explanation:

See the answer in image
Answer Area

Configuration WebServerConfif (litport-OstResource -todulHUne PsDesiredStateConfiguration
Mode ' localhost' (

WindowsFwture * Observer (Ensure • "Present' NMt • "ueb-Server"

File * OefaultHowePpge { Ensure • Present' SourcePath • 'Wserverl
VOSCResources\web'.index,ht*'
DestinationPath ■ 'c:\inetpub\wwwroot'

Box 1: Windows Feature

Example:

```
Configuration Website Test {

#   Import the module that contains the resources we're using.
Import-Disc Resource –Module Name Ps Desired State Configuration

#   The Node statement specifies which targets this configuration will beapplied to. Node 'localhost' {

#   The first resource block ensures that the Web-Server (IIS) feature is enabled.
Windows Feature Web Server {
Ensure = "Present"
Name = "Web-Server" }

Box 2: File
Example continued:
#   The second resource block ensures that the website content copied to the website root folder. File Website
Content { Ensure = 'Present'
Source Path = 'c:\test\index.htm'
Destination Path = 'c:\inetpub \wwwroot' }
```

Reference:
https://docs.microsoft.com/en-us/powershell/scripting/dsc/quickstarts/website-quickstart

Question: 310

You are evaluating the use of code review assignments in GitHub.
Which two requirements canbe met by using code review assignments' Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point

A. Automatically choose and assign reviewers based on a list of available personnel
B. Automatically choose and assign reviewers based on who has the most completed review requests.
C. Ensure that each team member reviews an equal number of pull requests during any 30-day period.
D. Automatically choose and assign reviewers based on who received the least recent reviewrequests.

Answer: A,C

Explanation:

Question: 311

You haw an Azure subscription that contains multiple Azure services. You need to send an SMS alert when scheduled maintenance is planned for the Azure services. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Create an Azure Service Health alert.
B. Enable Azure Security Center.
C. Create and configure an action group
D. Create and configure an Azure Monitor alert rule

Answer: A,D

Explanation:

Question: 312

You have a project m Azure DevOps that has a release pipeline.
You need to integrate work item tracking and an Agile project management system to meet the following requirements:

- Ensure that developers can track whether their commits are deployed to production.
- Report the deployment status.
- Minimize integration effort.

Which system should you use?

A. Trello
B. Jira
C. Basecamp
D. Asana

Answer: B

Explanation:

Jira Software is a development tool used by agile teams to plan,track, and manage software releases. Using Azure Pipelines, teams can configure CI/CD pipelines for applications of any language, deploying to any platform or any cloud.

Note: Microsoft and Atlassian have partnered together to build an integration betweenAzure Pipelines and Jira Software.

This integration connects the two products, providing full tracking of how and when the value envisioned with an issue is delivered to end users. This enables teams to setup a tight development cycle from issue creationthrough release. Key development milestones like builds and deployments associated to a Jira issue can then be tracked from within Jira Software.

Reference:
https://devblogs.microsoft.com/devops/azure-pipelines-integration-with-jira-software/

Question: 313

You have several Azure Active Directory (Azure AD) accounts.

You need to ensure that users use multi-factor authentication (MFA) to access Azure apps from untrusted networks.

What should you configure in Azure AD?

A. access reviews
B. managed identities
C. entitlement management
D. conditional access

Answer: D

Explanation:

You can configure a Conditional Access policy that requires MFA for access from untrusted networks.

Reference:
https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/howto-conditional- access-policy-all-users-mfa

Question: 314

You configure Azure Application Insights and the shared service plan tier for a web app.

You enable SmartDetection.

You confirm that standard metrics are visible in the logs, but when you test a failure, you do not receive a Smart Detection notification

What prevents the Smart Detection notification from being sent?

A. You must restart the web app before Smart Detection is enabled.
B. Smart Detection uses the first 24 hours to establish the normal behavior of the web app.
C. You must enable the Snapshot Debugger for the web app.
D. The web app is configured to use the shared service plan tier.

Answer: B

Explanation:

Question: 315

HOTSPOT

You have an Azure subscription that contains the resources shown in the following table.

Name	Type
DF1	Azure Data Factory
SQL1	Azure SQL Database
KV1	Azure Key Vault

You plan to create a linked service in DF1. The linked service will connect to SQL1 by using Microsoft SQL Server authentication. The password for the SQL Server login will be stored in KV1.

You need to configure DF1 to retrieve the password when the data factory connects to SQL1. The solution must use the principle of least privilege.

How should you configure DF1? To answer, select theappropriate options in the answer area.

NOTE:Each correct selection is worth one point.

Answer:

Explanation:

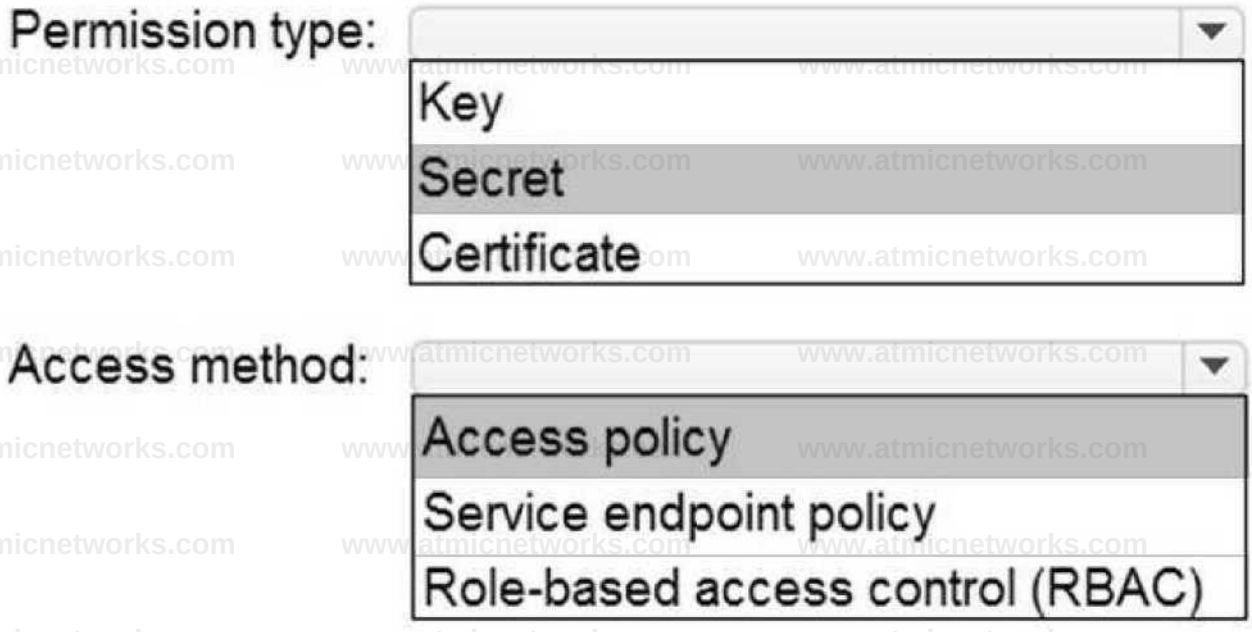

Box 1: Secret

Store credential in Azure Key Vault by reference secret stored in key vault.

To reference a credential stored in Azure Key Vault, you need to:

Retrieve data factory managed identity

Grant the managed identity access to your Azure Key Vault. In your key vault-> Access policies -> Add Access Policy, search this managed identity to grant Get permission in Secret permissions dropdown.

It allows this designated factory to access secret in key vault.
Create a linked service pointing to your Azure Key Vault.
Createdata store linked service, inside which reference the corresponding secret stored in key vault.

Box 2: Access policy

Reference:
https://docs.microsoft.com/en-us/azure/data-factory/store-credentials-in-key-vault

Question: 316

Note: This question is part of a series of questions that present the same scenario. Each question in theseries 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 Git repository in Azure DevOps.

Solution: Configure the build pipeline to use a Microsoft-hosted agent pool running the Windows Server 2019 with Visual Studio 2019 image. Include the Java Tool Installer task in the build pipeline.

Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

To build and deploy Windows, Azure, and other Visual Studio solutions you'll needat least one Windows agent. Windows agents can also build Java and Android apps.

The Azure Pipelines agent pool offers several virtual machine images to choose from, each including a broad range of tools and software. One such image is Windows Server 2019with Visual Studio 2019.

Reference:
https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/v2-windows?view=azure-devops

https://docs.microsoft.com/en-us/azure/devops/pipelines/agents/hosted?view=azure- devops&tabs=yaml

Question: 317

You have a project in Azure DevOps.

You create the following YAML template named Template1.yml.

```
steps:
- script: 9pm install
- script: yarn install
- script: 9pm run compile
```

You create the following pipeline named File1.yml.

```
parameters:
users steps:
- task: MyTask@1
- script: echo Done
```

You need to ensure that Template1.yaml runs before File1.yml.

How should you update File1.yml?

A. parameters: user steps: extends: template: template1.yml- task: MyTask@1 - script: echo Done
B. template: template1.yml parameters: user steps:- task: MyTask@1 - script: echo Done
C. extends: template: templatel.yml parameters: user steps:- task: MyTask@1 - script: echo Done
D. parameters: user steps: - template: templatel.yml- task: MyTask@1 - script: echo Done

Answer: C

Explanation:

Azure Pipelines offers two kinds of templates: includes and extends. Included templates behavelike #include in C++: it's as if you paste the template's code right into the outer file, which references it. To continue the C++ metaphor, extends templates are more like inheritance: the template provides the outer structure of the pipeline and a set of places where the template consumer can make targeted alterations.

Example:

```
extends:
template: template.yml@templates
parameters:
usersteps:

- script: echo This is my first step
- script: echo This is my second step
```

Reference:

https://docs.microsoft.com/en-us/azure/devops/pipelines/security/templates

Question: 318

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 theappropriate options in the answer area. NOTE: Each correct selection is worth one point. Answer Area "•

```
steps;
-    task; AiureResovrce"ana(ertee#lateOeploy«*nt#J
   Inputs;
     deployventScopt: 'Resource Group'
     acureResourceflanagsrConnectlont Fabr , i an Corporate! o41fb3eda2aa 42ffl alm Sfcc4efcc5db Subicrlptlonld:
     'd41de0e<Ja2»a-42fa-»7acdfcc6e*0:5dt action:                    *
rMourceftroupfcnna: Fabrtcw' location: 'West US' teaplateLOCatlon; 'Linked artifact' depioyacntHocie:   *
```

Answer: see the answer below.

Explanation:

Answer is below

Answer Area

```
step*:
-    tacki A;ure»ecovrce"an*|erteaplateCSeploymenx^J
   Inputs:
     deploymentScope; ■Resource Group'
     aiura*e*ourceHana|erConnect!on: 'Fabrlkan torpor ate(ellfMed a«aa 42ro-»7« BfccbafacSdb subscription!#: aaideOec-
     aiao-i^te-afac-sfccoe'Oiodt action:          Create Or Update Resource Goup *
resourct&roupMaM: 'ranrsKao' location: 'West US' teaplatalocatlon; Linked artifact'
     dcpieyaentHo4e: incremental *
```

Question: 319

DRAG DROP

You have an Azure Kubernetes Service (AKSJ 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.

Values

- Always
- initialDelaySeconds
- livenessProbe
- Never
- periodSeconds
- readinessProbe
- successThreshold
- Value

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
                                    : 3
            timeoutSeconds: 1
  osType: Linux
  restartPolicy:
tags: null
type: Microsoft.ContainerInstance/containerGroups
```

Answer

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: 320

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 filesby using commits to the repository.

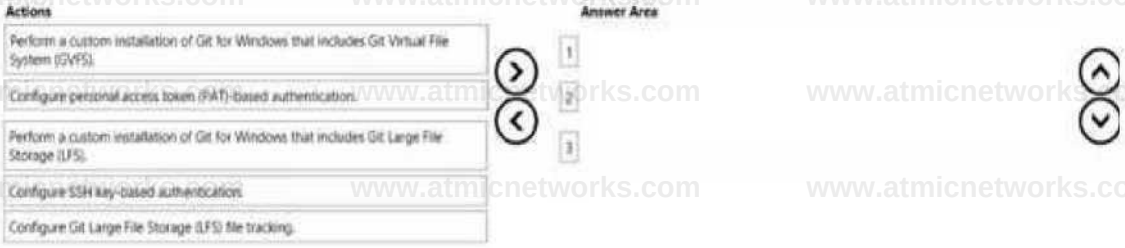

Answer:

Explanation:

Question: 321

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.

NOTE: Each correct selection is worth one point.

Git | helper-script > changes.txt

Answer Area

git * * | helper - script > changes.txt

Answer: see the answer below.

Explanation:

Answer below

get diff v1.4.16 v1.6.12 | helper-script > changes.txt

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: 322

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)

Answer: C

Explanation:

Question: 323

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 therepository.
- 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 beused 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.

Azure DevOps groups

- Build Administrators
- Contributors
- Project Administrators
- Project Collection Administrators
- Project Collection Valid Users

Answer Area

Project manager:

Lead developer:

Developer:

Answer:

Explanation:

Answer Area

Project manager: Project Collection Administrators

Lead developer [Project Adminrstrators

Developer Contributors

Question: 324

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

Answer: B,D

Explanation:

Question: 325

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 notat all. You may need to drag the spirt bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point

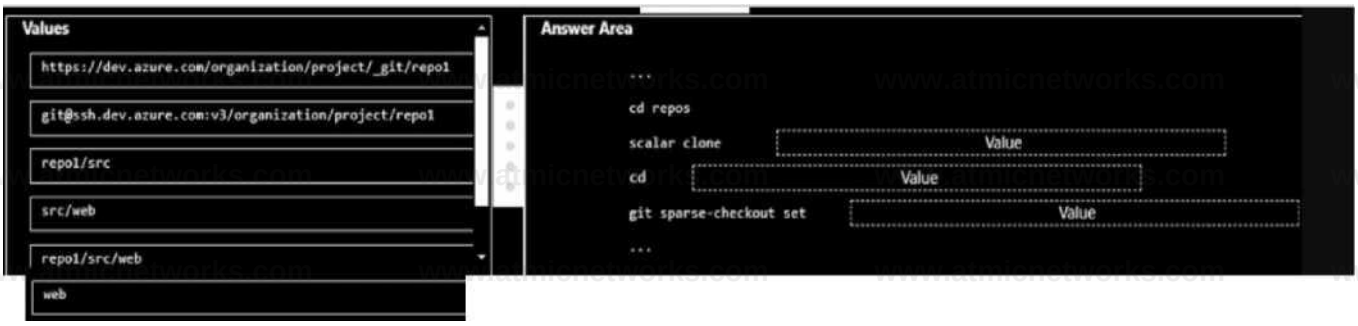

Answer:

Explanation:

git@ssh.dev.azure.com:v3/organization/project/repo1 src/web
web

Question: 326

DRAG DROP

You have an app named App1. Youhave a Log Analytics workspace named Workspace 1 that contains two tables named Events and Logs. App1 manages events in multiple locations and writes togs to

Workspace1.

You need to query Workspace1 for all log entries related to Asia that occurred duringthe last two days.

In which order should you arrange the query statements? To answer, move all statements from the let of statements to the answer area and arrange them in the correct order.

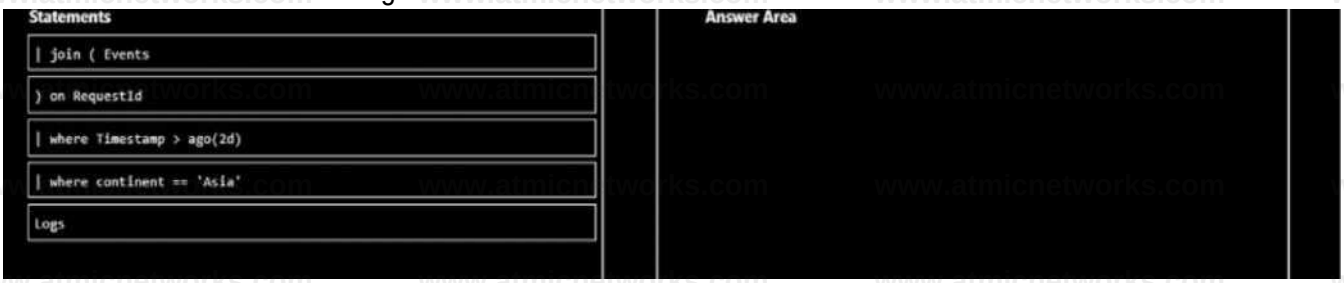

Answer:

Explanation:

Question: 327

DRAG DROP

You have a project in Azure DevOps named Project1 that has a release pipeline in Azure Pipeline named ReleaseP1. you need toensure 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 r 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

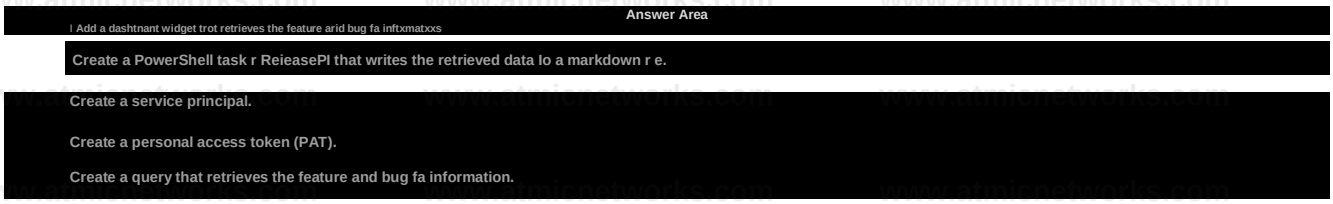

Answer:

Explanation:
Create a service principal.
Create a PowerSheel task in RealeaseP1 that writes the retrieved data to a markdown file.
Add a dashboard widget that retrieves the feature and but fix information.

Question: 328

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 togs 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

Answer: C

Explanation:

An analytics feature in a monitoring solution would allow you to parse logs from multiple sources and analyzethem 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 customdashboards and alerts to monitor your pipelines and quickly identify and troubleshoot any issues.

Question: 329

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.

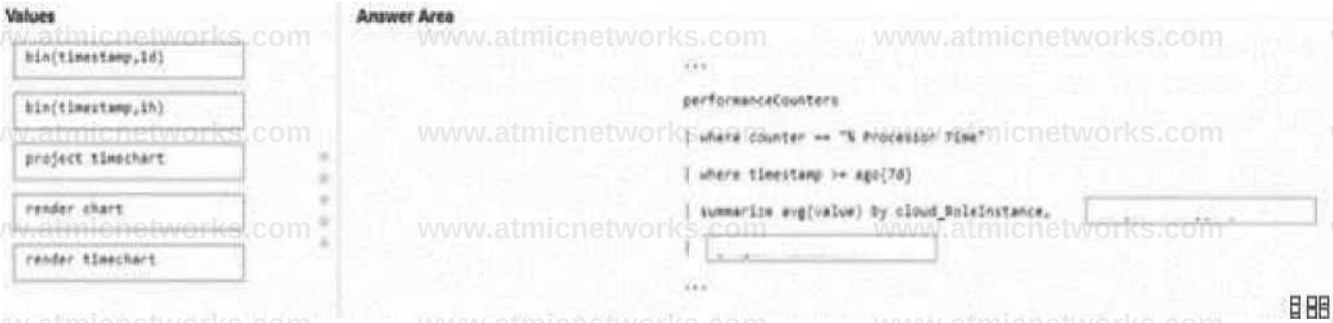

Answer:

Explanation:

https://docs.microsoft.com/en-us/azure/azure-monitor/app/performance-counters

Question: 330

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 solutionmust meet the following requirements:

- Ensure that User 1 can actively push to repo1.
- Ensure that Usef2 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 use?

Answer:

Explanation:

Question: 331

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.

Answer: D

Explanation:

Question: 332

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

Answer: A

Explanation:

Question: 333

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 scan directory.
Exclude scanning of markdown files

Answer:

Explanation:

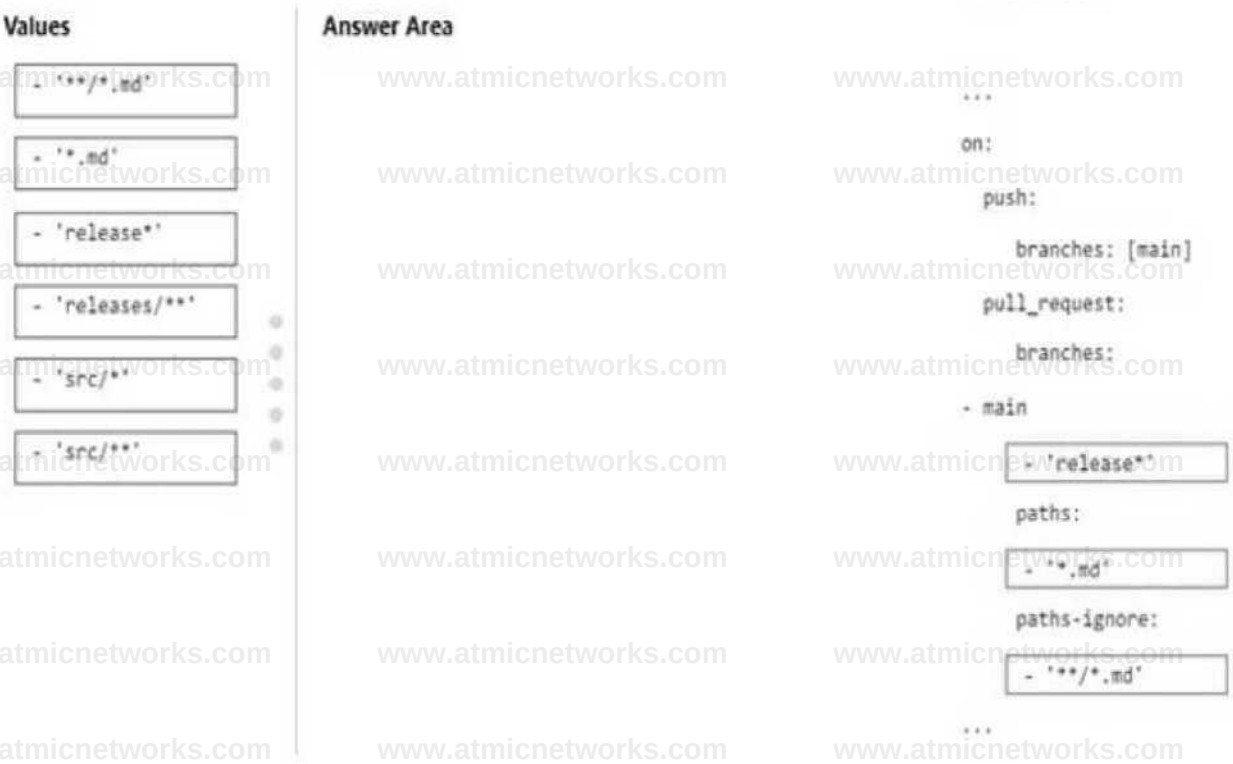

Question: 334

You have project in Azure DevOps.
You create the following template named Template1.yml.

```
steps:
```

- script; r>p* install
- script: yarn install
- script: npr run compile

You create the following pipeline named File1.yml.

You need to ensure that Template1.yaml runs before File1.yml.
How should you update File1.yml?

A. parameters: usersteps: extends: template: template!.y«l - task: Hylaskwl - script: echoDone

3. extends: template: template!.y»l parameters: usersteps: - task: MyTask^l • script: echo Done

C. parameters: usersteps: • template: template!.y«l • task: My ask^i • script: echo Done

D. template: template!.yel parameters: usersteps: - task: MyTask^l - script: echo Done

A. Option A B. Option B C. Option C D. Option D

Answer: B

Explanation:

Question: 335

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 eventsin 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?

Answer:

Explanation:

Statements

| join (Eventi
| where continent = 'Asis'

Answer Area

```
1  ) on RequestId
2  | where Timestamp > ago(2d)
3  Logs
```

Question: 336

DRAG DROP

You have web app named App1 that uses Application Insights in Azure Monitor to Store log dat a. 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, result Code, and duration fields.

Only requests made the last hour.

How shout you complete the query?

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
...
```

Answer:

Explanation:

Values

- extend
- project
- select
- timestamp >= ago(1hr)
- timestamp -gt ago(1hr)

Answer Area

```
requests
| where [ t line steep -qt ago(lhr)
| where resultcode « '404' and (client-City « 'London" or client_City ■» "Paris")
| I select    ti'estarp, wrl, resultcode, duration
```

Question: 337

HOTSPOT

You have an Azure subscription that contains two resource groups named Contoso RG and Contoso Dev, 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.

The **[answer choice]** setting must be changed to prevent the modification of existing databases and web apps in ContosoRG.

- Action
- Template location
- Deployment mode
- Deployment scope

Pipeline! will retrieve the ARM template from the **[answer choice].**

- output of the continuous integration build location specified in the Linked artifact variable default branch of the Grt repository of Contoso Data Factory

Answer

Explanation:

Deployment Mode
Location specified in the Linked artifact variable

Question: 338

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 Pipeline' 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 tor compliance.
C. To Pipeline1, add a step that runs a What If deployment before thedeployment step.
D. Configure a deployment gate for Pipeline' that uses Azure Automation to run a What If deployment

Answer: A

Explanation:

Question: 339

DRAG DROP
You have an Azure subscription that uses Azure Monitor and contains a Log Analytics workspace.
Youhave an encryption key.

You need to configure Azure Monitor to use the key to encrypt log data

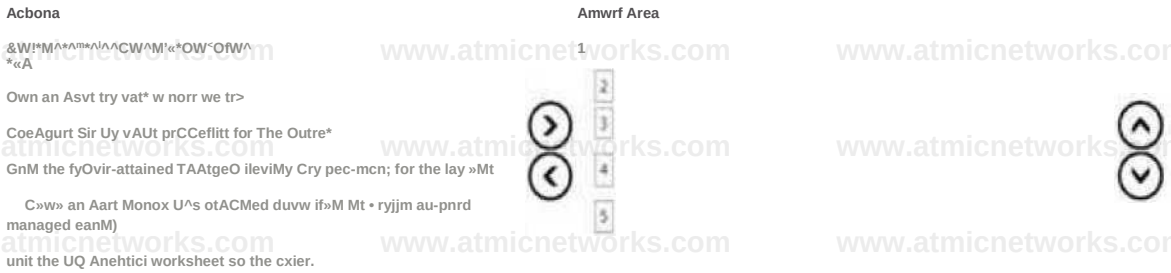

Answer:

Explanation:

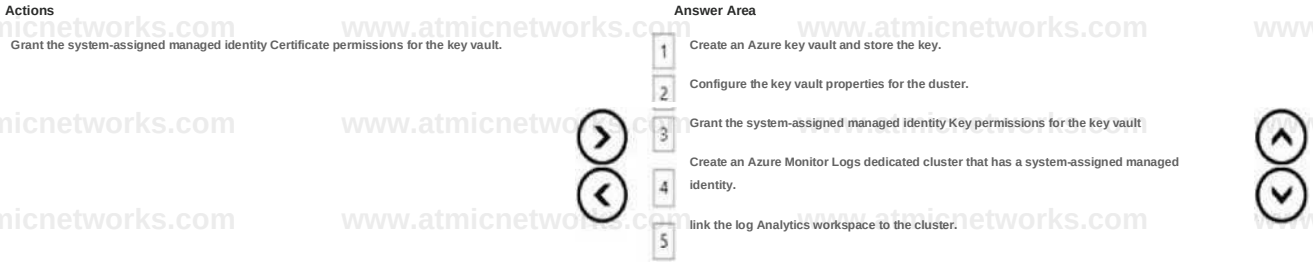

Question: 340

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 beversioned.

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 inAzure Storage and enable blob versions.
B. Install the Get IFS extension and associate the extension to ZIP files.
C. Use G Zip to compress the file before committing the file.
D. Install the get-stash extension and associate the extension to ZIP files.
E. Install the get-fat extension and associate the extension to ZIP files.

Answer: B,D

Explanation:

Question: 341

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 Statuschanges. Which services can be notified by using a web hook?

A. Service1 only
B. Service2 only
C. Service1 and Service2 only

Answer: C

Explanation:

Question: 342

You are designing a YAML template for use with Azure Pipelines. The template Will include theOutputfile 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. $(paraaeters['outputfile'])

B. $ {parameters*outputfile;'

C. S(parameters.outputfile

D $(para«eters[outputfile])

```
E . ${{parameters['output file']}}
```

A. Option A B. Option B C. Option C D. Option D E. Option E

Answer: C,D

Explanation:

Question: 343

You manage projects by using Azure Boards.

You have a current work item name item A that is dependent 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?

A. FromBacklogs, open the context menu, select Add link and then select item3. Set Link type to Related and add the ID of item A

B. From item A, open the Links tab, and then select Add link. Set Link type to Successor and add the ID of item B.

C. 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 item B.

D. From item A, open the Links tab, and then select Add link. Set Link type to Reference and add the ID Of item B.

Answer: B

Explanation:

Question: 344

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 outcomesfor 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.

Answer: B

Explanation:

Question: 345

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 isworth 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)

Answer: C,D

Explanation:

Question: 346

You have an Azure subscription that containsAzure 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 area.

Answer: See the image in explanation for answer.

Explanation:

Answer is as below.

```
inputs:
  key;                    v '"yarn" | "$(Agent.OS)" | yarn.lock'
  path:                   ' $(YARNCACHE_FOLDER)
displayflame: Cache Yarn packages
- script: yarn --frozen-lockfile
```

Question: 347

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 threeactions 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. Addthe package URL to the Get 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.

Answer: A,B,E

Explanation:

Question: 348

You have an Azure subscriptionthat 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

Answer: B

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: Thisblock 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 My Dsc Configuration. 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 tworesource blocks, each of which call the Windows Feature resource."

https://docs.microsoft.com/en-us/powershell/dsc/configurations/configurations?view=dsc- 1.1#configuration-syntax

Question: 349

HOTSPOT

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?

Answer Area

branch
switch
reset

git restore —staged ^ «»in*3 App.exe.config

*-patch

--merge

I -source

Answer:

Explanation:

Answer Area

git checkout ^ main

git restore —Staged ˅ «aln«»3 App.exe.config

Question: 350

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

Answer: B

Explanation:

Question: 351

HOTSPOT

You use Get for source control.

You need to optimize theperformance of a repository. The solution must meet the following requirements:

Permanently remove all items referenced only in the ref log.

Remove history that is NOT in any current branch.

How should you complete the command? To answer, select theappropriate options in the answer area.

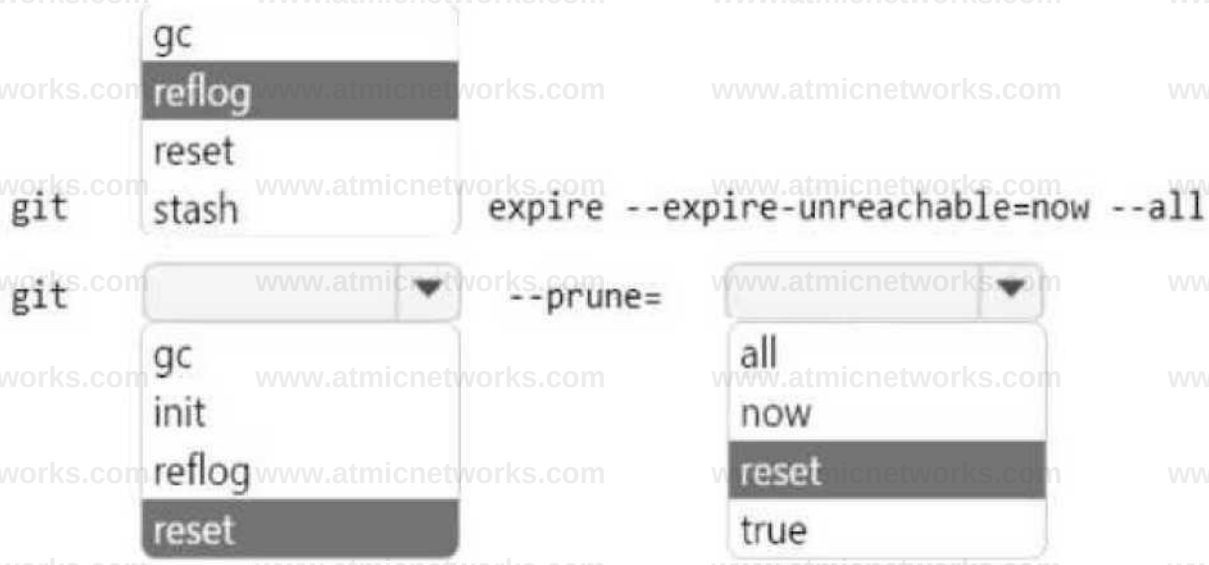

Answer:

Explanation:

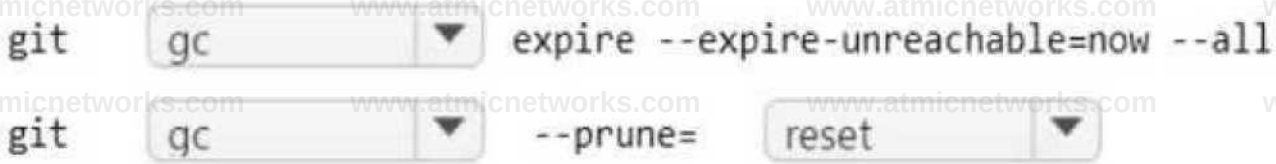

Question: 352

You havean 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.

Answer: B

Explanation:

Question: 353

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 youmodify?

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

Answer: A

Explanation:

To minimize the amountof 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 theMicrosoft documentation.

Question: 354

You have an Azure subscription that contains the resources shown in the following table.

DepPipeline1 and ADFPipeline1 use a single credential that is stored in Vault'.

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

Answer: A

Explanation:

Question: 355

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 thedocuments 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 acomplete solution.

A. Jeky11
B. Medium
C. caret
D. WordPress
E. Doc FX

Answer: A,D

Explanation:

A. Jekyll is a static site generator that can be used to generate a static website from Markdown files stored in a GitHub repository. Jekyll supports the use of Markdown for writing documentation, and it can automatically recompile the website whenever changes are made to the documentation.
E. Doc FX is another tool that can be used to compile a static website from Markdown files stored in a GitHub repository. It isan open-source tool that can be used to generate API documentation, reference documentation, and other types of documentation from source code and other files. Doc FX supports the use of Markdown for writing documentation and it can automatically recompilethe website whenever changes are made to the documentation.

Question: 356

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?Toanswer, 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.

Answer:

Explanation:

git lfs config add "*.psd"

git lfs migrate import ..include-"*.psd: --everything git lfs update

Question: 357

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 thefollowing 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

Answer: B

Explanation:

System-assigned managed identities enable Azure resources to authenticate to cloud services without storing credentials

in code. They also support granting privileges to the identity, making them the ideal choice for this scenario.Source: Microsoft

Question: 358

You use GitHub for source control and Azure Boards for project management. GitHub and Azure Boards are integrated.
You plan to create a pull request in GitHub.
You need to automatically link the request to an existing Azure Boards work item by using the text of AB#<WORKITEMNUMBER>.
To which two elements can you add the text? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. milestone
B. comment
C. title
D. description
E. label

Answer: C,D

Explanation:

You can add the text "AB#<WORKITEMNUMBER>" to the title or description of the pull request in GitHub, which will automatically link the request to an existing Azure Boards work item with that number.

Reference:

Azure DevOps Docs: Link a work item to a pull requesthttps://docs.microsoft.com/en- us/azure/devops/repos/git/pull-requests?view=azure-devops#link-a-work-item-to-a-pull-request GitHub Docs: Creating a pull requesthttps://docs.github.com/en-us/github/collaborating-with-issues- and-pull-requests/creating-a-pull-request

Question: 359

DRAG DROP

You have an Azure DevOps pipeline that is used to deploy a Node.js app.
You need to ensure that the dependencies are cachedbetween builds.
How should you configure the deployment YAML? 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.

inputs:

key: 'np« | *J(A|e«t.O$)' I peckege*lock.json' restoreKeys: |

Values

- always()
- build.sh
- eq(variables.CACHE_RESTORED, 'true')
- integrationtest.sh
- ne(variables.CACHE_RESTORED, 'true')
- npm install

Answer Area

np« I 'S(A#tnt.OS)'

path: S(*»P«jCMfig.CKhe)

cecheHitVar: CACHE RESTOPED

• script:

condition:

Value

Value

Answer:

Explanation:

```
inputs:
  key: 'npm | "$(Agent.OS)" | package-lock.json'
  restoreKeys: |
    npm | "$(Agent.OS)"
  path: $(npm_config_cache)
  cacheHitVar: CACHE_RESTORED

- script:    npm install
  condition: ne(variables.CACHE_RESTORED, 'true')
```

Question: 360

HOTSPOT

You are designing YAML-based Azure pipelines for the apps shown in the following table

You need to configure the YAML strategy value for each app. The solution must minimize appdowntime. Which value should you configure for each app? To answer, select the appropriate options in the answer area.

Appl: canary rolling rvnonce

App?

canary rolling

ru nance

Answer:

Explanation:

App1Canary and App2 rolling

App1 Canary would minimize app downtime for the first app, as it would only deploy new code when the canary has confirmed that it is functional - and if there are any issues, it would roll back to the previous version of the code.
App2 rolling would be the second option, as it would allow for frequent deployments of new code, while still giving the developers enough time to fix any issues that may have been introduced during new code deployments.

Question: 361

DRAG DROP

You useGet for source control.
You delete a file, commit the changes, and continue to work.
You need to recover the deleted file.
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.

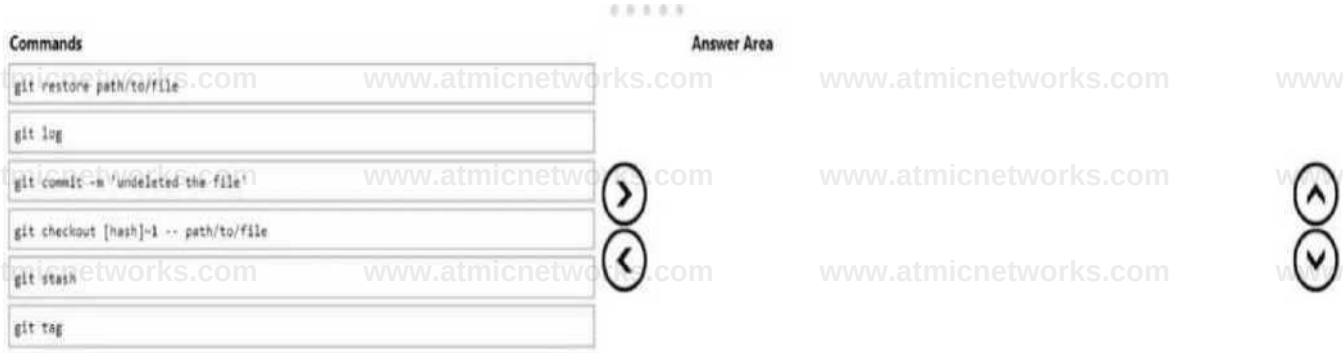

Answer:

Explanation:

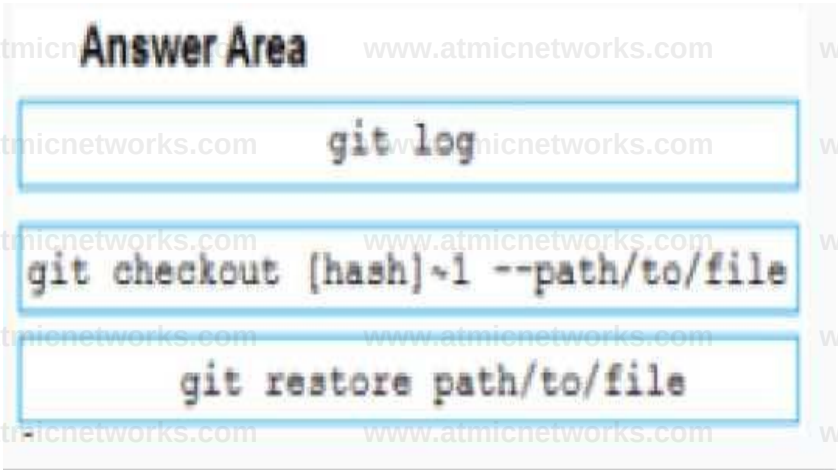

Question: 362

DRAG DROP

You have an app named App1. You have a Log Analytics workspace named Workspace 1 that contains two tables named Events andLogs. 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, moveall statements from the list of statements to the answer area and arrange them in the correct order.

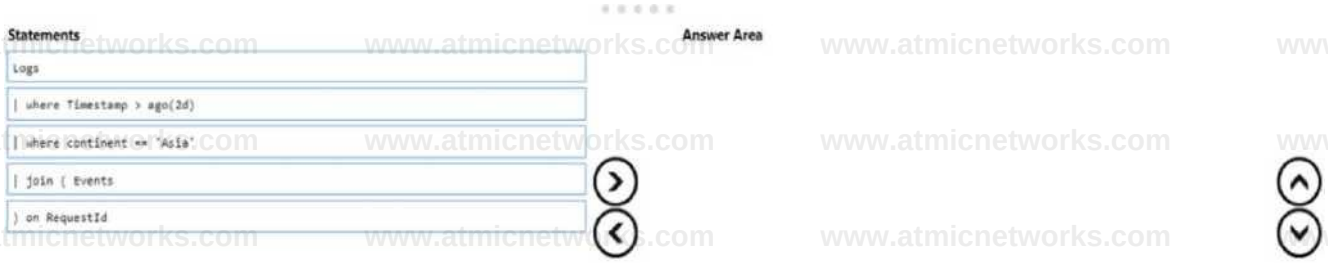

Answer:

Explanation:

Answer Area

| join (Events

logs

| where Tlmstaap > ego (2d)

I where continent *• 'Asia'

) on RequestId

Question: 363

You have an Azure subscription that contains the resources shown in the following table.

Project1 produces 9pm packages that are published to Feed1. 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 tags the packages as release.

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 tags the packages as release.

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 promotes a package to the @release view.

D. Create a feed view named @default. After the 9pm packages test successfully, configure a release pipeline that promotes a package to the @default view.

Answer: C

Explanation:

By creating a feed view named "release" and setting it as the default view, packages that are published to the feed will not be immediately available for consumption. After the 9pm packages are tested successfully,you can configure a release pipeline that promotes a package to the @release view. This ensures that only tested packages are available for consumption and minimizes development effort as it doesn't require any additional steps to be taken by the consumerof the feed. Reference:

Azure DevOps Docs: Create a feed and viewshttps://docs.microsoft.com/en-us/azure/devops/artifacts/feeds/create-feed?view=azure-devops

Azure DevOps Docs: Promote a packagehttps://docs.microsoft.com/en-us/azure/devops/artifacts/feeds/promote-package?view=azure-devops

Question: 364

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.

Answer:

Explanation:

Question: 365

DRAG DROP

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

You need to create a Docker 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: Eachcorrect selection is worth one point.

Answer:

Explanation:

mcr.microsoft.com/dontnet/aspnet:5.0 dotnet publish -c Release -o out dotnet restore

Question: 366

You have a GitHub repository that contains workflows. The workflows contain steps that execute predefined actions. Each action has one or more versions.

You need to request the specific version of an action to execute.

Which three attributes can you use to identify the version? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. the SHA-based hashes
B. the tag
C. the runner
D. the branch
E. the serial

Answer: A,B,D

Explanation:

https://docs.github.com/en/actions/using-workflows/workflow-syntax-for-github-actions
"We strongly recommend that you include the version of the action you are using by specifying a Git ref, SHA, or Docker tag. If you don't specify a version, it could break your workflows or cause unexpected behavior when the action owner publishes an update. Using the commit SHA of a released action version is the safest for stability and security. Ifthe action publishes major version tags, you should expect to receive critical fixes and security patches while still retaining compatibility. Note that this behavior is at the discretion of the action's author. Using the default branch of an action may beconvenient, but if someone releases a new major version with a breaking change, your workflow could break."

Question: 367

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains aunique 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 9pm 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 using pipeline artifacts.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Pipeline artifacts are a way to persist build outputs,test results, and other files generated during a pipeline run. They allow you to share data between stages, jobs, and pipelines, and to persist data for longer than the lifetime of a pipeline run. While artifacts can be useful for sharing data between pipeline runs and reducing the time required to download dependencies, they are not a solution for reducing the time required to install JavaScript packages from 9pm during a pipeline run.

The solution of reducing the pipeline execution time could be achievedby using package caching, which allows you to store and reuse 9pm packages from previous pipeline runs.There are several package caching options available for Azure Pipelines, including the 9pm task, the 9pm cache task, and the 9pm ci task. All of theseoptions allow you to configure caching for your 9pm packages, which can significantly reduce the time required to install packages during subsequent pipeline runs. Another solution could be using a dedicated agent that has those packages already installed,this way the pipeline doesn't have to install them again.

You can find more information on package caching by following this linkhttps://docs.microsoft.com/en-us/azure/devops/pipelines/tasks/package/npm- cache?view=azure-devops

Question: 368

You have an on-premises app named App1 that accesses Azure resources by using credentials stored in a configuration file.

You plan to upgrade App1 touse an Azure service principal.

What is required for App1 to programmatically sign in to Azure Active Directory (Azure AD)?

A. the application ID, a client secret, and the object ID
B. a client secret, the object ID, and the tenant ID
C. the application ID, a client secret, and the tenant ID
D. the application ID, a client secret, and the subscription ID

Answer: C

Explanation:

https://docs.microsoft.com/en-us/azure/active-directory/develop/app-objects-and-service- principals "When you've completed the app registration,you've a globally unique instance of the app (the application object) which lives within your home tenant or directory. You also have a globally unique ID for your app (the app or client ID). In the portal, you can then add secrets or certificates and scopes to make your app work, customize the branding of your app in the sign-in dialog, and more."

Question: 369

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'

OependsOn • *[Kindow$Feature]Web-Server'

Hindis Feature Neb-Server {

Ensure • 'Present*

Name • 'web-Server*

You have the following Local Configuration Manager (LCM) configuration.

localCo^igurat 1 onHan age r {

ConfigurationMode ■ "ApplyAn^onitor"

AefreshFrequencyMins ■ 30

ConfigurMionPodeFrtquencyt'ins • 69 **RtfreshHode ■**

'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.

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		
If manual changes are made to the configuration of the virtual machine, the configuration @ will reapply automatically	O	
If the Web-Server Windows feature is uninstalled from the virtual machine, the discrepancy will be reported in a log entry within 60 minutes	#	

Answer:

Explanation:

Yes " Dependson are defined"

Yes "Mode is Apply and Monitor" No"Configuration Mode FreqMin : 60 "

Configuration Mode Freq Min:How often, in minutes, the current configuration is checked and applied. This property is ignored if the Configuration Mode property is set to Apply Only.

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=dsc- 1.1

Question: 370

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.

Values

- Category
- DscReportStatus
- Message
- OperationName
- Resource
- ResultType

Answer Area

```
AzureDiagnostics
| where [          ] == "DscNodeStatus"
| where [          ] contains ""
| where [          ] != "Compliant"
```

Answer:

Explanation:

```
AzureDiagnostics
| where [ Category ] == "DscNodeStatus"
| where [ OperationName ] contains ""
| where [ ResultType ] != "Compliant"
```

https://learn.microsoft.com/fr-fr/azure/automation/automation-dsc-diagnostics

Question: 371

You have a project in Azure DevOps named Project1.

You implement a Continuous Integration/Continuous Deployment (CI/CD) pipeline that uses PowerShell Desired State Configuration (DSC) to configure theapplication infrastructure.

You need to perform a unit test and an integration test of the configuration before Project1 is

deployed.

What should you use?

A. the PS Script Analyzer tool
B. the Pester test framework
C. the PS Code Health module
D. the Test-Ds Configuration cmdlet

Answer: B

Explanation:

You should use the Pester test framework to perform a unit test and an integration test of the configuration before Project1 is deployed. The Pester test framework is a PowerShell testing framework that can be used tovalidate PowerShell DSC configurations.

Question: 372

You use Azure DevOps processes to build and deploy code.

You need to compare how much time is spent troubleshooting issues found during development and how much time is spent troubleshooting issuesfound in released code.

Which KPI should you use?

A. defect escape rate
B. unplanned work rate
C. defect rate
D. rework rate

Answer: A

Explanation:

The defect escape rate is a metric that assesses the collective quality of software releases by evaluating how oftenerrors are discovered and rectified in the pre-production process versus during production.

The defect escape rate is a KPI (Key Performance Indicator) that measures how many defects are found in released code versus how many are found during development.This KPI can help you to compare how much time is spent troubleshooting issues found during development versus how much time is spent troubleshooting issues found in released code. The higher the defect escape rate, the more defects are found in releasedcode, and thus more time is spent troubleshooting issues in released code.

Question: 373

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 webapp 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.

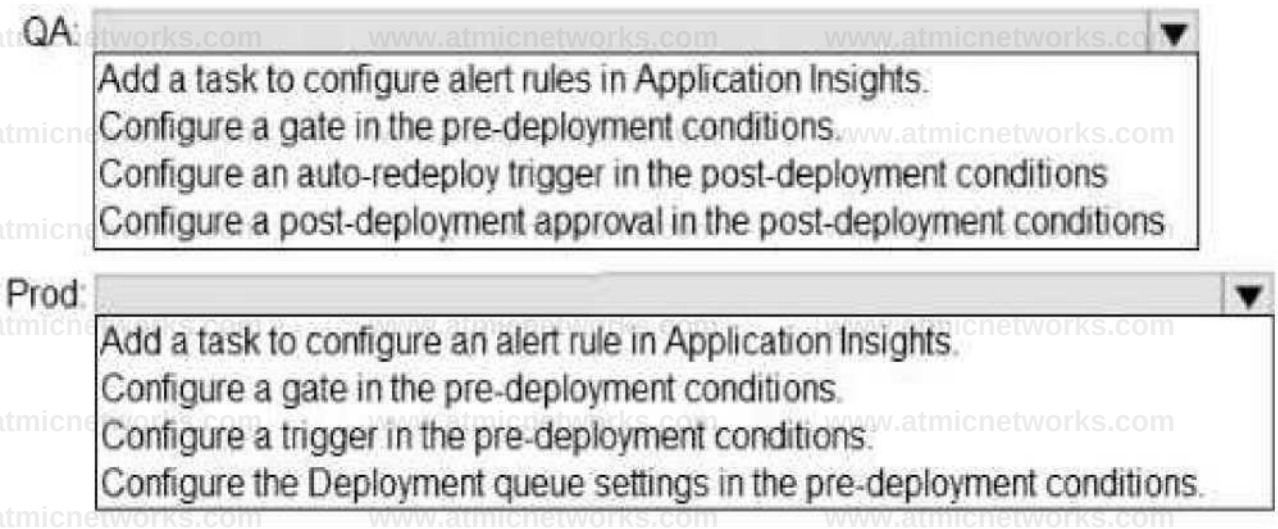

Answer:

Explanation:

QA: Add a task to configure alert rules in Application Insights

Pro: Configure a gate in the pre-deployment conditions

Question: 374

You have a project in Azure DevOps named Project1.

You need to ensure that all new pipelines in Project1 execute three specific tasks during pipeline execution.

What should you create?

A. a task group

B. a JSONtemplate

C. a YAML template

D. a PowerShell task

Answer: C

Explanation:

A task group in Azure DevOps is a collection of tasks that can be reused across multiple pipelines. You can create a task group that contains the three specific tasks that you need to execute during pipeline execution, and then reference that task group in all new pipelines in Project1. This way, you can ensure that the three specific tasks are executed in all new pipelines without having to manually add them to each pipeline.

Question: 375

DRAG DROP

You use Extra beam Fusion SIEM and the Azure cloud platform.

You need to integrate Extra beam and Azure. The solution must use OAuth authentication.

Which three actions should you perform in sequence? To answer, move the appropriate actions fromthe list of actions to the answer area and arrange them in the correct order.

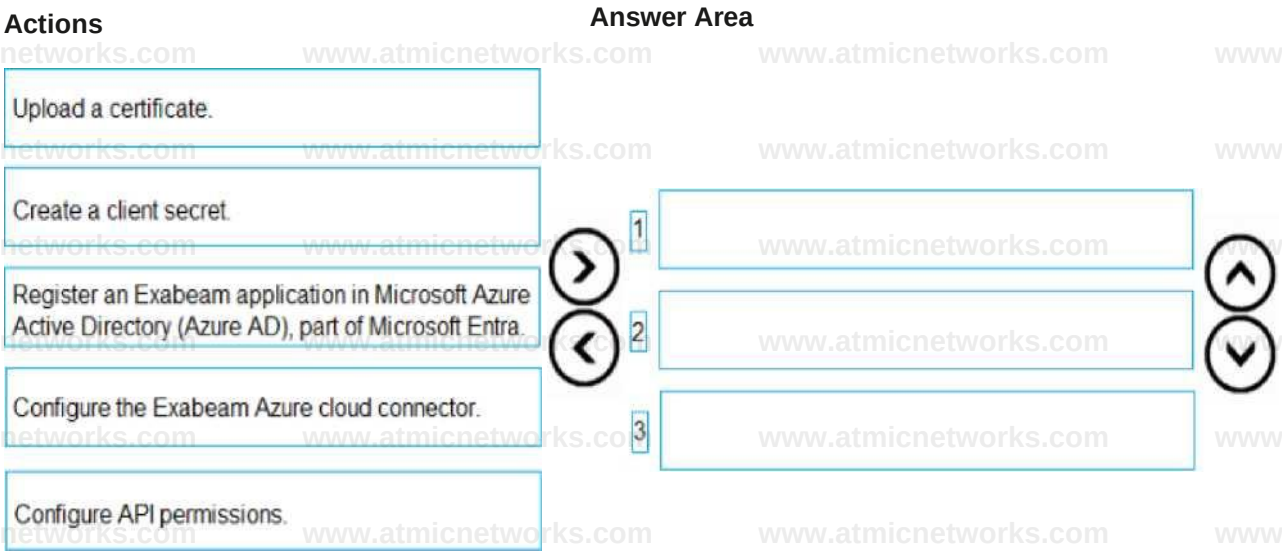

Answer:

Explanation:

Create a client secret
Register an Exabeam application in Microsoft Azure...

Configure the Exabeam Azure Cloud Connecter

Question: 376

You plan to create a GitHub workflow that will use GitHub Actions. The actions will require a 256-KB secret.

You need to recommend a solution to storeand encrypt the secret. The secret value must be accessible only to the workflow. The solution must minimize administrative effort

What should you recommend?

A. Store the secret in the organization-level GitHub secrets.
B. Store the secret in the repository-level GitHub secrets.
C. Encrypt the secret value and store the value in the repository. Store the decryption key in the repository-level GitHub secrets.
D. Encrypt the secret value and store the value in the repository. Store the decryption key in the organization-level GitHub secrets.

Answer: C

Explanation:

https://docs.github.com/en/actions/security-guides/encrypted-secrets "To use secrets that are larger than 48 KB, you can use a workaround to store encrypted secrets in your repository and save the decryptionpassphrase as a secret on GitHub." Because it requires less administrative privilege it's at repository level

Question: 377

You have a pipeline named Pipeline1 in Azure Pipelines.
You need to create a service connection to enable Pipeline1 to downloada public container image.
Which type of service connection should you create?

A. a Docker host
B. Azure Service Fabric
C. Azure Kubernetes Service (AKS)
D. a Docker registry

Answer: D

Explanation:

Question: 378

You have an Azure Automation account that contains a webbook. The webbook is used to configure the application infrastructure of an Azure subscription.
You have a project in Azure DevOps named Project). Project1 contains a repository that stores code for the webbook.
You need to ensure that every committedchange to the code will update automatically and publish the webbook 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 theAutomation account
D. the Service connections settings for Project1

Answer: C

Explanation:

Question: 379

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 webbook
D. a gist

Answer: C

Explanation:

Question: 380

You have a GitHub Enterprise account.
You need to enable push protection for secret scanning of the account repositories.
What should you dofirst?

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.

Answer: C

Explanation:

Question: 381

DRAG DROP

You have an Azure Repos repository namedrepo1.

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.

Answer

Explanation:

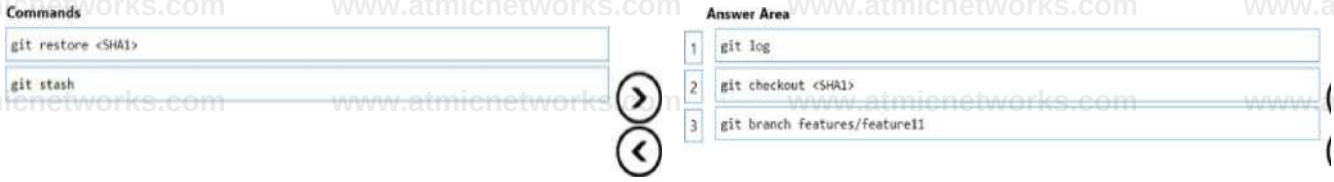

Question: 382

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.

Mow should you configure the variable? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

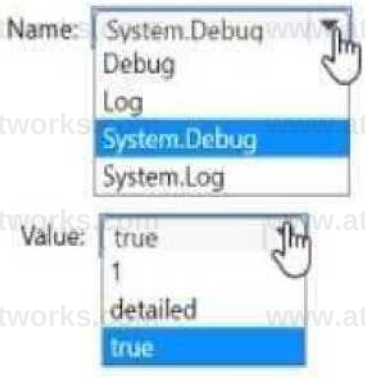

Answer:

Explanation:

Answer Area

Name: System.De bug

Value: true

Question: 383

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.

Answer Area

Store the password as a; | Secret' J

Certificate
Key
Secret

Grant Pipeline: access to VaultI by modifying the:] Access policies

Access control (IAM| settings
Access policies
Security settings

Answer:

Explanation:

Answer Area

Store the password as a: Secret

Grant Pipelinel access to Vault! by modifying the: Access policies

Question: 384

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

Answer: D

Explanation:

Question: 385

You have a GitHub repository that contains multiple versions of an Azure Pipelines template. You plan to deploy multiple pipelinesthat 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

Answer: B

Explanation:

Question: 386

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
10.3	Manually pushed io Feedl
1.4.0	Manually pushed to Feedl
2.0.0	Available from an upstream source
2.3.1	Satea from an upstream source

You need to perform a buildof Project1. Which version of Package1 will be used?

A. 1.0.3
B. 1.4.0
C. 2.0.0
D. 2.3.1

Answer: D

Explanation:

Question: 387

HOTSPOT

You have an Azure web app named webapp1 that uses the .NET Core runtime slack. You have an Azure Application insights resourcenamed Applnsights1 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 Applnsights1 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.

Answer Area

From the code repository of webappt;

Disable adaptive sampling
P Enable fued-rate sampling

itatMnlnsighlKtK

From Apptnsights 1: Modify the Usage and estimated costs settings
Configure Continuous export
Configure the Smart Detection settings
Modify the Usage and estimated costs settings

Answer:

Explanation:

Answer Area

From the code repository of webappl:

: I Modify Applicationinsghtsxcnfn ^[k

Disable adaptive sampling

|. Enable fnedrate sampling _____ .

Modil hcationii

From Apptnaghtsl: Modify the Usage ana esumarea costs settings

Configure Continuous export

Configure the Smart Detection settings

Modify the Usage and estimated costs settings.

Question: 388

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

Answer: D

Explanation:

Question: 389

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
E. Option A
F. Option B
G. Option C
H. Option D

Answer: C

Explanation:

Question: 390

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 appropriateoptions in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
Cara* storageAccount strlog

 var jtorijfAccjufltHMeToUse • *S[storageAccount}${uMqoe$trlng(resource6roup().id;}' resource Invoices tor
age 'Microsoft .Storages storageAccountig2O22-t9*tl * • {

    naoe: storageAccounViameTouse location: 'vastus' slmt (
        naae: *Standardj6RS*

  kind          ' ssor»c<v2'
[kind
  param: properties: type bar          I properties * (
Wn(,:          rfflcOnly: true
  param
properties-
type- vat
```

Answer:

Explanation:

Answer Area

```
paraa storageAccount string van steeageAccountNaneToUse *
■i(storageAccount)${unipueString(resourceGroup(). W!) resource invoiceStorage
'Microsoft.Storage/storageAccountsS2322-35-01' • {
   name: storageAccountHaneTollse location: eastus'
sku: { naan: ' StandardCPS'

 kind          ■'     ^ tor agev?-

 properties      " {
   SUpportsHttpsTrafficOnly: true }
```

Question: 391

Note: This question is partof 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 m 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

Answer: A

Explanation:

Question: 392

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 actionsto the answer area and arrange them in the correct order.

Answer:

Explanation:

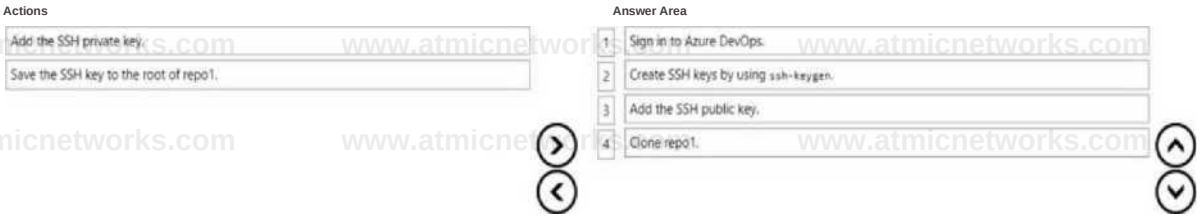

Question: 393

DRAG DROP

You are implementing a newproject in Azure DevOps.

You need to assess the performance of the protect. The solution must identity 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 spirt bar between panes or scroll to view content.

NOTEEach correct selection is worthone point.

Answer:

Explanation:

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- | Cycle time

The percentage of defects found m production j Bug report rates

Question: 394

DRAG DROP

You create a Git repository named Repo1 in Azure Repos.
You need to configure Repo1 to meet the following requirements:

- Work items must be linked to a pull request.
- Pull requests must have a minimum of two reviewers.
- Pull requests must complete a code review by using a thirty-party tool.

The solution must minimize administrative effort.
Which type of policy should you usefor 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.

Policies

- Branch
- Build
- Check in
- Status

Answer Area

Work items most be linked to a pull request

Puli requests must complete a code review Dy 4 two-party tool

Pull requests must have a minimum ot wo

Answer:

Explanation:

Policies

- Build
- [OISS ----------] •
- Status

Answer Area

Work items must be linked to a pull request Status

Pull requests must complete a code review Dy grands a two party toor

Puli requests must have a minimum of wo Branch reviewers

Question: 395

You have a project in Azure DevOps.

You need to push notificationsaboutpullrequests to a Microsoft Teams channel. The solution must minimize development effort.
What shouldyou do?

A. Install the Azure Repos app for teams and configure a subscription to receive notifications in the

channel.

B. Use AzureAutomation to connect to the Azure DevOps REST API and send messages to Teams.

C. Install the Azure Pipelines app for Teams and configure a subscription to receive notifications in the channel.

D. Use an Azure function to connect to the Azure DevOps RESTAPI and send messages to Teams.

Answer: C

Explanation:

Question: 396

Youhave a GitHub repository that is integrated with Azure Boards Azure Boards has a work item that has the number 715.

You need to ensure that when you commit source code in GitHub, the work item is updated automatically.

What should you include in the commit comments?

A. @714

B. =715

C. the URL of the work item

D. AB#715

Answer: B

Explanation:

Question: 397

You manage source code control and versioning by using GitHub.

A large file it committed to arepository accidentally.

You need to reduce the sizeof the repository. The solution must remove the file from the repository.

What should you use?

A. bfg

B. 1fs

C. gvfs

D. init

Answer: A

Explanation:

Question: 398

You manage code by using GitHub.
You need to ensure that repository owners are notified if a new vulnerable dependency or malware is found in their repository.
What should you do?

A. Configure branch protection rules for each repository.
B. Configure Dependabot alerts.
C. Configure CodeQL scanning actions.
D. Subscribe all the repository owners to the GitHub Advisory Database.

Answer: B

Explanation:

Question: 399

You have multiple teams that work on multiple projects in Azure DevOps.
You need to plan and manage the consumers and producers for each project. Thesolution must provide an overview of all the projects.
What should you do?

A. Install the Dependency Tracker extension and create dependencies for each project.
B. Add a Parent or Child link to the feature or user story for the items of each project.
C. Add a Predecessor or Successor link to the feature or user story for the items of each project.
D. Create a custom query to show the consumers and producers and add a widget to a dashboard.

Answer: D

Explanation:

Question: 400

You have a project in Azure DevOps namedProject that contains a Kanban board named Board1.
You create a Microsoft Teams channel and add the Azure Boards app to the channel.
You need to ensure that users can create work items in Board1 from Microsoft Teams.
Which command should you run?

A. @azure boards sign in
B. @azure boards link

C. @azure boards create

D. @azure boards subscriptions

Answer: D

Explanation:

Question: 401

DRAG DROP

You have a project in Azure DevOps.

You need to configure a dashboard. The solution must include thefollowing metrics:

- Bottlenecks in the software development process
- A burndown chart for the work in a single iteration
- How long it takes to close a work item after the item was started

Which type of widget should you use for each metric? To answer, drag the appropriate widget types to the correct metrics. Each widget type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

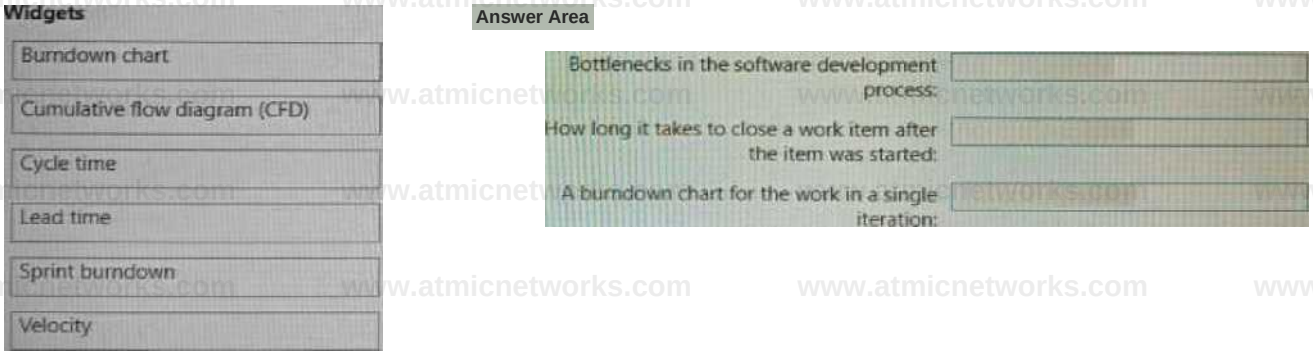

Answer:

Explanation:

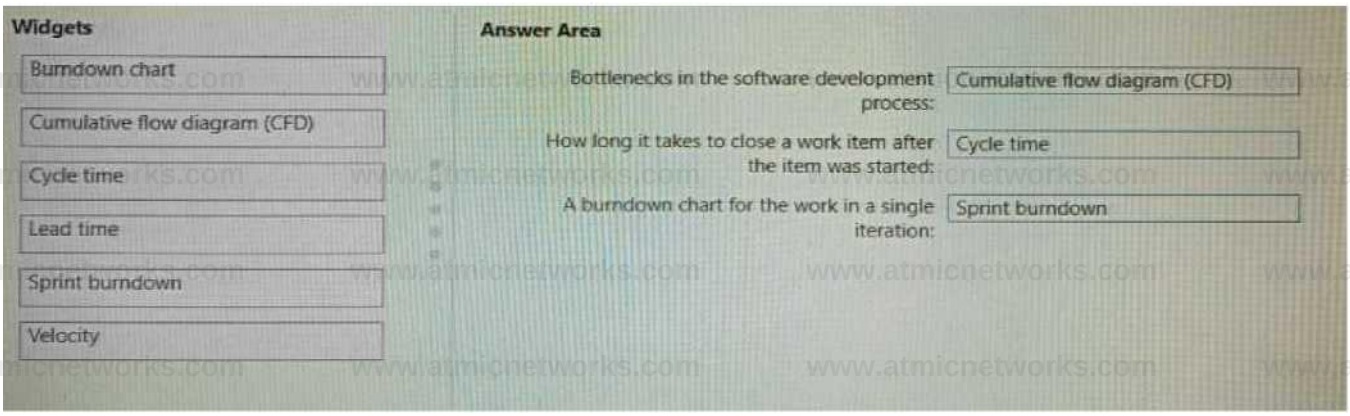

Question: 402

You have an Azure key vault named KV1 and three web servers.

You plan to deploy an app named App1 to the web server. You need to ensure that App1 can retrieve a secret fromKV1. The solution must meet the following requirements:

- Minimize the number of permission grants required
- Follow the principle of least privilege.

What should you include in the solution?

A. role-based access control (RBAQ permissions
B. a system-assigned managed identity
C. a user-assigned managed identity
D. a service principal

Answer: B

Explanation:

Question: 403

DRAG DROP

You use Semantic Versioning (Sem Ver) as a dependency versioning strategy.

You perform changes to code as shown in the followingtable.

Name	New functionality	Change incompatibility
Change1	Yes	Minor
Change2	Yes	Significant
Change3	No	Minor

Which part of the version should you increment for each change? To answer, drag the appropriate parts to the correct changes. Each part may be used once, more than once, or not at all. You may need to

NOTE: Each correct selection is worth one point.

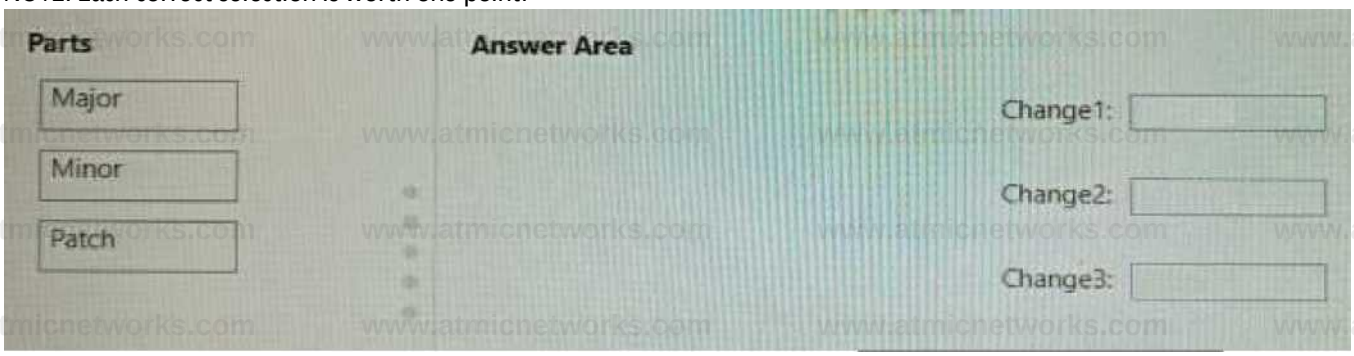

Explanation:

drag the split bar between panes or scroll toview content.

Answer:

Minor

Question: 404

DRAG DROP

You need to deploy Internet Information Services (IIS) to an Azure virtual machine that runs Windows Server 2022

How should you complete the Desired State Configuration (DSC) configuration script? To answer, drag the appropriate values to the correct locations. Eachvalue 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.

Answer:

Explanation:

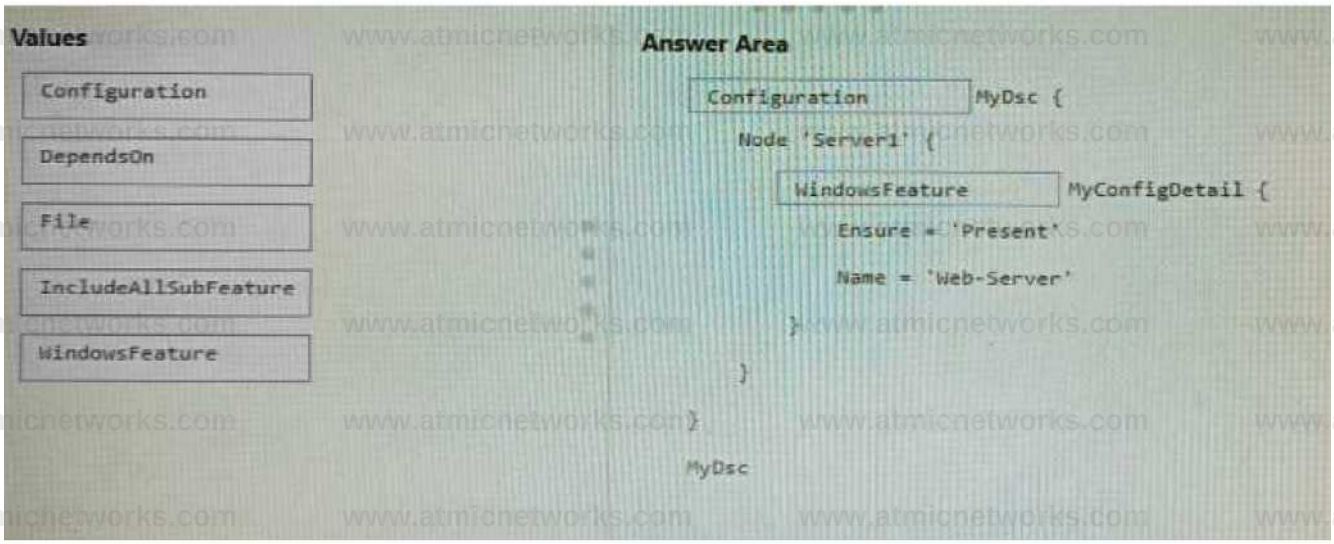

Question: 405

You have an Azure subscription linked to a tenant in Microsoft Azure Active Directory (Azure AD), part of Entrap. The tenant is licensed for Azure AD Premium Plan 1.

A securityreview indicates that too many users have privileged access to resources.

You need to deploy a privileged access management solution that meets the following requirements:

- Enforces time limits on the use of privileged access
- Requires approval to activate privileged access
- Minimizes costs

What should you do first?

A. Configure alerts for the activation of privileged roles.
B. Enforce Azure Multi-Factor Authentication (MFA) for role activation.
C. Configure notifications when privilegedroles are activated.
D. Upgrade the license of the Azure AD tenant.

Answer: D

Explanation:

Question: 406

You have an Azure web app named webapp1 that uses the .NET Core runtime stack. You have an Azure Application Insights resource named AppInsight1. Webapp1 sendstelemetry data to AppInsights1.

You need to ensure that webapp1 sends the telemetry data at a fixed sampling rate.

What should you do?

A. From the code repository of webapp1, modify the ApplicationInsights.config file.
B. From the code repository of webapp1, modify the Startup.cs file.
C. From AppInsights1. modify the Usage and estimated costs settings.
D. From AppInsights1, configure the Continuous export settings.

Answer: B

Explanation:

Question: 407

DRAG DROP

You have a large repository named Repo1 that contains a directory named directory 1.

You plan to modify files in directory1.

You need to create a clone of Repo1. The solution must minimize the amount of transferred data. 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.

Answer:

Explanation:

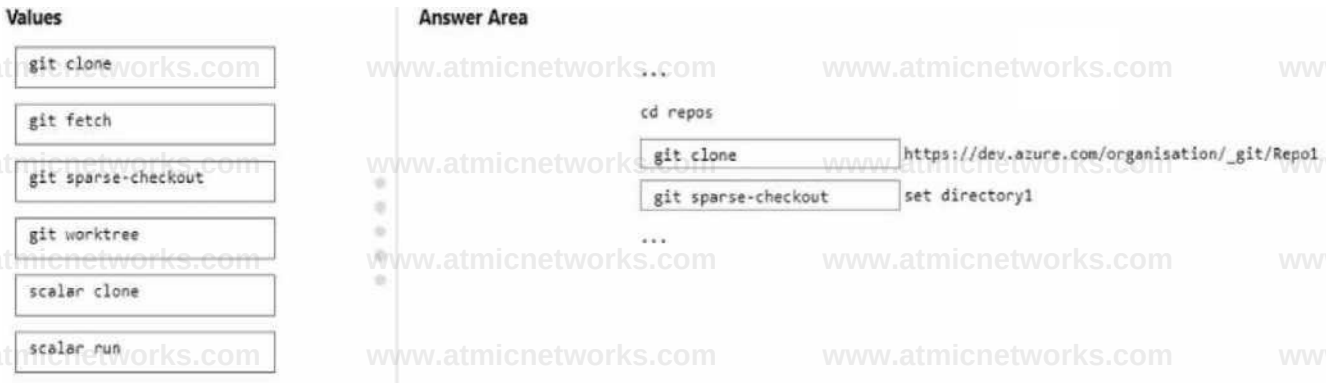

Question: 408

Note: This question is part of a series of questions thatpresent 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 pipeline that is used to deploy a web app. The pipeline includes a test suite named TestSuite1. TestSuite1 is used tovalidate the operations of the web app.

TestSuite1 fails intermittently.

You identify that the failures are unrelated to changes in the source code and execution environment.
You need to minimize troubleshooting effort for the TestSuite1 failures.
Solution: You enable flaky test management.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

Question: 409

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 pipeline that is used to deploy a web app. The pipeline includes a test suite named TestSuite1. TestSuite1 is used to validate the operations of the web app.
TestSuite1 fails intermittently.
You identify that the failures are unrelated to changes in the source code and execution environment.
You need to minimize troubleshooting effort for the TestSuite1 failures.
Solution: You implement the Test Results Trend widget.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Question: 410

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 pipeline that is used to deploy a web app.The pipeline includes a test suite named TestSuite1. TestSuite1 is used to validate the operations of the web app.
TestSuite1 fails intermittently.

You identify that the failures are unrelated to changes in the source code and execution environment.
You need to minimize troubleshooting effort for the TestSuite1 failures.
Solution: You enable Test Impact Analysis (TIA).
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Question: 411

You have an Azure subscription that contains a Log Analytics workspace named WS1 and a virtual machine named VM1.
You need to install the Microsoft Enterprise Cloud Monitoring extension on VM1.
Which two values are required to configure the extension? Each correct answer presents part of the solution.
NOTE: Each correct answeris worth one point.

A. the secret key of WS1
B. the ID of the subscription
C. the system-assigned managed identity of VM1
D. thelDofWS1
E. the resource ID of VM1

Answer: D,E

Explanation:

Question: 412

You have an app named App1 that uses Application Insights to monitor application performance.
You need to analyze how often a page in App1 is accessed.
Which pane in Application Insights should you use?

A. Events
B. Sessions
C. Impact
D. Users

Answer: D

Explanation:

Question: 413

You have an app named App1 that you releaseby using Azure Pipelines. App1 has the versions shown in the following table.

You complete a code change to fix a bug that was introduced in version 3.4.3.

Which version number should you assign to the release?

A. 3.4.4
B. 3.4.8
C. 3.5.0
D. 4.0.1

Answer: D

Explanation:

Question: 414

You plan to publish build artifacts by using an Azure pipeline.

You need to create an. artifactignore file that meets the following requirements:

- Includes all files in the build output folder and all subfolders
- Excludesfiles that have the .dll extension

What should you include in the file?

A.
```
./**
!*.dll
```

B.
```
**/*
!*.dll
```

C.
```
*/**
*.dll
```

D.
```
**/*
#*.dll
```

A. Option A B. Option B C. Option C D. Option D

Answer: B

Explanation:

Question: 415

DRAG DROP

You have an Azure Key Vault that contains an encryption key named key1.

You plan to create aLog Analytics workspace that will store logging data.

You need to encrypt the workspace by using key1.

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 themin the correct order.

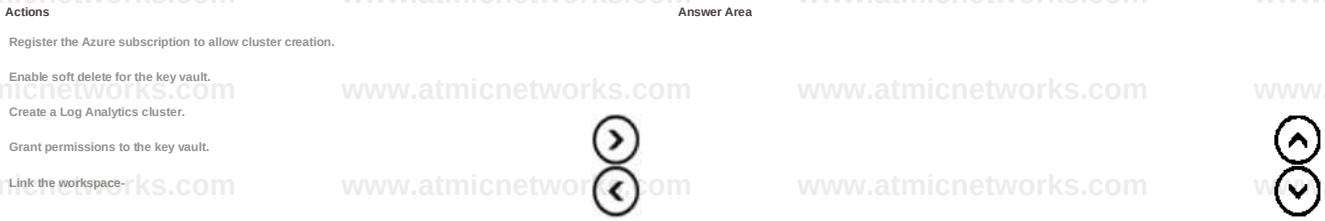

Answer:

Explanation:

Question: 416

HOTSPOT

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

You plan to use Azure Monitor to manage logs.

You need to ensure that the users can perform the actions shown in following the table.

The solution must follow the principle of least privilege.

Which role should you assign to each user? To answer select the appropriate options in the answerarea.

NOTE: Each correct selection is worth one point.

Answer Area

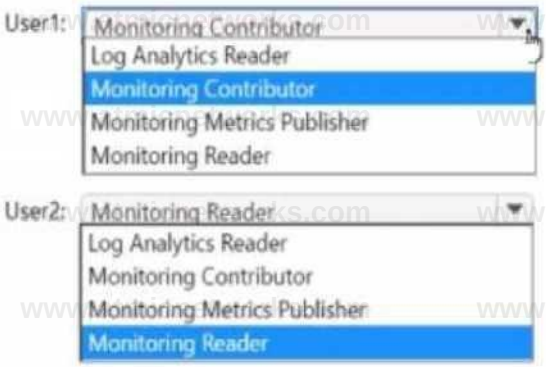

User	Action
Use ri	• Create private monitoring dashboards. • Search usage data for an Azure Monitor workspace.
User 2	• View autoscale settings. • View alert activities and settings.

Answer:

Explanation:

Answer Area

Used: Monitoring Contributor

User2: Monitoring Reader

Question: 417

You plan to use Azure DevOps to build and deploy anapp that will be hosted in a Kubernetes cluster.
You need to scan the app image for vulnerabilities before the image is deployed to the cluster.
What should you include in the solution?

A. Microsoft Defender for DevOps
B. Microsoft Defender for Storage
C. Microsoft Defender for Containers
D. Microsoft Defender for App Service

Answer: C

Explanation:

Question: 418

You use GitHub for source control
You are evaluating whether to use proxying to add a private upstream MyGet package feed to your MyGet feed.
What are twopossible advantages of this approach? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. minimizes latency when accessing the package
B. provides automatic authentication
C. minimizes the impact on your storage quota
D. minimizes the impact of upstream source availability issues

Answer: B,C

Explanation:

Question: 419

DRAG DROP
You have an Azure Pipeline.

You need to store configuration values as variables.

At which four scopes can the variables be defined, andwhat is the precedence of the variables from the highest precedence to lowest precedence? To answer, move the appropriate scope from the list of scopes to the answer area and arrange them in the correct order.

Answer:

Explanation:

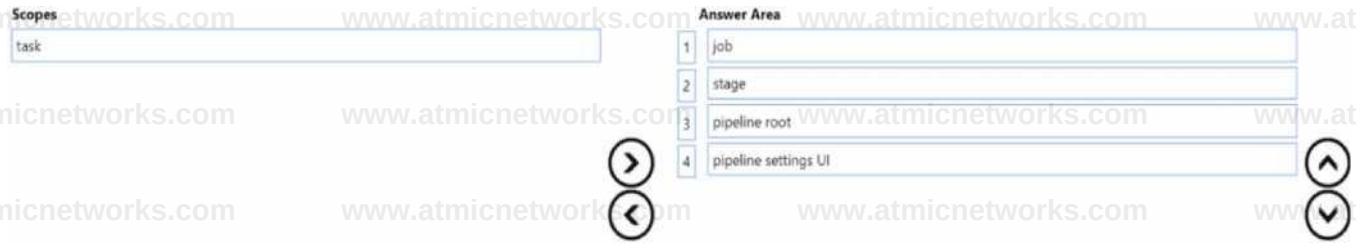

Question: 420

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 2022. You need to recommend adeployment 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 shouldyou 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

Answer: C

Explanation:

Question: 421

You use Azure Pipelines to build and test code projects.

You notice an increase in cycle times.

You need to identify whether agent pool exhaustionis causing the issue.

What are two possible ways to achieve this goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. View the Pipeline duration report.
B. Query the PipelineRun/PipelineRuns endpoint.
C. View the pool consumption report at the organization level.
D. Query the TaskAgentPoolSizeSnapshots endpoint

Answer: C,D

Explanation:

Question: 422

You have a project in Azure DevOps named Project1 that contains two environments named environment1 andenvkonment2.
When a new version of Project1 is released, the latest version is deployed to environment2, and the previous version is redeployed to environments
You need to distribute users across the environments. The solution must meet the following requirements:

- New releases must be available to only a subset of the users.
- You must gradually increase the number of users that can access environment2.

What should you use?

A. web app deployment slots
B. Azure Traffic Manager
C. VIP swapping
D. Azure Load Balancer

Answer: A

Explanation:

Question: 423

You use Git for source control.
You enable GitHub code scanning.
You raise a pull request from a non-default branch. In the code scanning output you receive the following error message: "Analysis not found."
You need to ensure that the code scanning completes successfully for the pull request.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Add a new workflow for code scanning.
B. Add the name of the non-default branch to the on: push specification in the code scanning workflow.
C. Update the code in the pull request.
D. Add the name of the default branch to the on: push specification in the code scanning workflow.
E. Delete the pull request, and then raise the request again from the default branch.

Answer: C,D

Explanation:

Question: 424

DRAG DROP

You have a tendril in Microsoft Azure Active Directory (Azure AD), part of Microsoft Entr a. The tenant contains three groups namedGroup1. Group2, and Group3.

You create a new project in Azure DevOps named Project1.

You need to secure the service connections for Project1. The solution must meet the following requirements:

- The members of Group1 must be able to share and unshare aservice connection with other projects.
- The members of Group2 must be able to rename a service connection and update the description.
- The members of Group3 must be able to use the service connection within build or release pipelines.
- The principle of least privilege must be followed.

Which permission should you grant to each group? To answer, drag the appropriate permissions to the correct groups. Each permission may be used once, more than once, or not at all. You may need to dragthe split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Permisssions

Contributor

Creator

Organization-level Administrator

Project-level Administrator

User

Answer Area

Groupl:

Group2:

Group3:

Answer:

Explanation:

Question: 425

You have an Azure subscription that contains an Azure container registry. The container registry contains an ACR Tasks task namedTask1. Task1 is configured to run once every five days.

You need to trigger Task1 to run immediately.

Which command should you run?

A. az acr build
B. az acr task run
C. az acr run
D. az acr taskrun

Answer: B

Explanation:

Question: 426

DRAG DROP

You have therepositories shown in the following table.

TyPe	URL
Azure Repo^	https:// dev.azu remarryconrow/projectl /_gi V project 1 git
G tHub	https^/g ithuhuconVconto so/p reject 1 .g it

You need to migrate the contents of the GitHub repository to the Azure Repos repository. The solution must ensure that the Azure Repos repository only contains branches and history from the GitHub repository.

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.

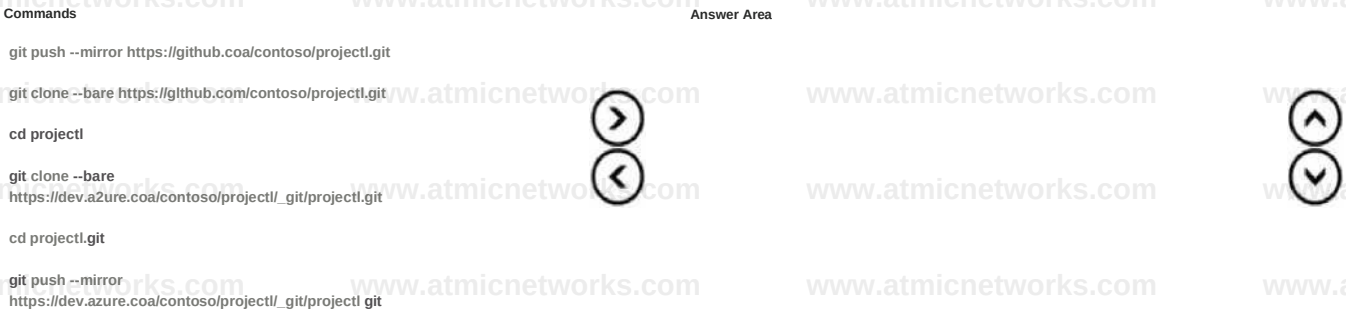

Answer:

Explanation:

Question: 427

You manage code by using GitHub.

You plan to use Dependabot to scan for code dependencies.

You need to identify when scanning will be triggered automatically.

Which two actions will trigger a scan? Each correct answer presents a complete solution.

NOTE: Each correct solution is worth one point

A. The dependency graph of a repository changes.
B. Any commit is pushed.
C. A branch is forked.
D. A pull request is created.
E. A new advisory is added.

Answer: B,D

Explanation:

Question: 428

HOTSPOT

You plan to use Desired StateConfiguration (DSC) to maintain the configuration of a server that runs Windows Server 2022. The server must have the following features installed:

- A web server
- An email server

How should you complete the DSC configuration file? To answer,select the appropriate options in the

answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
configuration RequiredFeatures
{
    I«tport-D-sc Resource -ModuleNaae PSBesiredStateConfiguration
    Node localhost
    {
        hindowsFeatureSet RequiredNindowsFeatures
        {
            Hase
                - @('M?il Server" "IIS")
                MsrSKHP-Sf™^^                    5")
                <- ri ' Mail-Server", "Web-Server") -
                @^ SMTP-Server", 'Web-Server')
            = 'Present '
                Install
                Ensure
                Enforce
                Required
            InclqdeAllSubFeature = itruc
        }
    }
}
```

Answer:

Explanation:

Answer Area

```
configuration RequiredFeatures
{
    libport-Ose Resource -KoduleHaae PSUesiredStateConfiguration
    Mode localhost
    TindoKsFeatureSet RequiredWindowsFeatures
    Name = (g}f SMTP'Server" "Web Server") ^
    Ensure = 'Present'
    IncludeAllSubFeat
}
```

Question: 429

DRAG DROP

You have a project in Azure DevOps.

You need to configure a dashboard that will provide information on the following metrics.

- How long it takes to close a work item
- The number of completed backlog items
- How long it takes to restore failed services

Which type of widget should you use for each metric? To answer, drag the appropriate widget types to the correct metrics. Each widget 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.

Widgets	Answer Area
Bumdown	How long it takes to close a work item.
Cycle time	The number of completed backlog items:
Lead time	How long it takes to restore failed services:
Velocity	

Answer:

Explanation:

To configure a dashboard that will provide information on the metrics you mentioned, you can use the following types of widgets:

How long it takes to close a work item: Cycle time. This widget shows the average time it takes for work items to go from a specified start state to a specified end state1.

The number of completed backlog items: Velocity. This widget shows the amount of work delivered by a team within a

sprint or iteration2.

How long it takes to restore failed services: Lead time. This widget shows the average time it takes for work items to go from creation to completion3.

Question: 430

You have a public GitHub repository named Public1.

A commit is made to Public1. The commit contains a pattern that matches a regular expression.

Who is notified first when the commit is made?

A. the owner of Public1
B. the secret scanning partner
C. the committer
D. the administrator of the GitHub organization

Answer: B

Explanation:

Question: 431

You have an Azure subscription that contains an Azure Pipelines pipeline named Pipeline1 and an app namedApp1.

Pipeline1 is used to automate the building of App1.

You have a Slack channel named App1chat that includes an incoming webhook.

You need to ensure that when a successful build of App1 is created, a notification is sent to App1chat by using the webhook.

What should you use?

A. an alert rule
B. a notification
C. an action group
D. a subscription

Answer: D

Explanation:

Question: 432

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 pipeline that is used to deploy a web app. The pipeline includes a test suite named TestSuite1.

TestSuite1 is used to validate the operations of the web app.

TestSuite1 fails intermittently.

You identify that the failures are unrelated to changes in the source code and execution environment.
You need to minimize troubleshooting effort for the TestSuite1 failures.
Solution: You increase code coverage.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Question: 433

You use Azure Pipelines to build and deploy an app named Appl. You plan to monitor App1 by using Application Insights.
You create an Application Insights instance named All. You need to configure App1 to use All. Which file should you modify?

A. appsettings.son
B. launchSettings.json
C. startup.cs
D. project.son

Answer: A

Explanation:

Question: 434

HOTSPOT

You have an Azure subscription that contains a user named User1.
You have an Azure Resource Manager (ARM) template named Template 1.
You plan to perform the following actions:

- Deploy an Azure key vault named KV1.
- Deploy Azure resources by using Template1 to retrieve secrets from KV1

You need to ensure that User1 can deploy Template1 The solution must follow the principle of least privilege.
Which permission should you grant to User1, and which parameter should be specified when you create KV1? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Permission. Mu rosoft. KevVault/v.iults/deploy/ac turn

Microsoft KeyVault/vaiilts/keys/read

M ic rosoft Resou rces/subsc opt ions/resou rt e<jr oops/read

Parameter enabled-for-template deployment

• enable-rbacautltonzation

enabled for deployment

- -enabled • for-template-deploymerit

Answer:

Explanation:

Answer Area

Permission: Microsoft KeyVault/vaolts/deployMtion

Parameter -enabled-for-template deployment

Question: 435

DRAG DROP

You have an Azure subscription that contains a project in Azure DevOps named Project1. In Microsoft Azure Active Directory (Azure AD), part of MicrosoftEntr

a. you have three users that require access to Project! as shown in the following table.

Name	Title	Requirement
Used	Project Manager	View repertories
Useri	Development lead	Create repositories and manage permissions
UserS	Developer	Create branches and tags

You need to ensure that the users have the appropriate permissions. The solution must use the principle of least privilege.

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

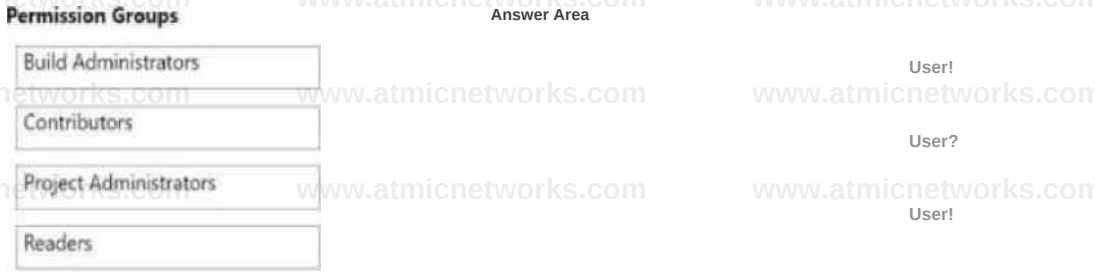

Answer:

Explanation:

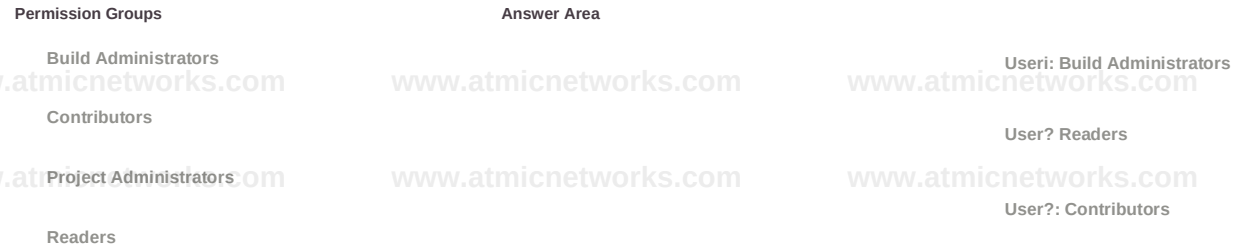

Question: 436

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

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

- Enable users to test new features by using aspecific URL.
- Minimize the effort required to promote a test version to production.
- Minimize the effort required to revert production code to the previous version.

Which strategy should you recommend?

A. A/B
B. feature toggle
C. progressive exposure
D. blue/green

Answer: A

Explanation:

Question: 437

You have an app that is deployed to two environments named Production-A and Production-B by using Azure Pipelines.

You need to configure a release pipeline that will mark the app as complete and ready forrelease into the Production-B environment. The solution must meet the following requirements:

- Ensure that there are no active Azure Monitor alerts in the Production-A environment before the app is marked as complete.
- Minimize administrative effort.

What should you do?

A. To the Production-B environment stage, add a pre-deployment gate that will query Azure Monitor. B. To the Production-A environment stage, add a post-deployment gate that will query Azure Monitor.

C. To the Production-A environment stage, add a post-deployment approval.

D. To the Production-A environment stage, add a pre-deployment gate that will query Azure Monitor.

Answer: B

Explanation:

Question: 438

You have an Azure subscription.

You use Bicep templates to deploy websites and AzureSQL infrastructure.

You need to automate the deployments by using Azure Pipelines and a self-hosted agent that runs on two virtual machines.

The solution must minimize administrative effort.

What should you do first?

A. Create an Azure Automation account.

B. On each virtual machine, enable a system-assigned managed identity.

C. Create a user-assigned managed identity.

D. Create a service principal.

Answer: D

Explanation:

Question: 439

HOTSPOT

You have a .NET app named App1.

You need to upload App1 to GitHubPackages.

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

NOTE: Each correct selection is worth one point.

Answer Area

Answer:

Explanation:

Answer Area

make " nuget push ^ "bin/Release/app1.1.0.0.nupkg" --api-key PAT_TOKEN --source "github"

Question: 440

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? Toanswer, move the appropriate tasks from the list of tasks to the answer area and arrange them in the correct order

Tasks Answer Area

:: Microsoft Purview Scan

:: Advanced Security AutoBuild

;: Advanced Security Initialize CodeQL

:: Advanced Security Perform Coded. Analysis

;; Advanced Security Dependency Scanning

Answer:

Explanation:

Advanced Security Intialize CodeQL

Advanced Security AutoBuild

Advanced Security Perform CodeQL Analysis

Question: 441

DRAG DROP

You have a GitHub repository named repo1.

You migrate repo1 to an Azure Repos repository named repo2.

After themigration, 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 dragthe split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

- repo1
- repo1 clone URL
- repo1 repo URL
- repo2
- repo2 clone URL
- repo2 repo URL

Answer Area

git cion* --bar*

git r*not* add --nirrorsfetch upstream

git f*tch upstream --tags

git push origin --*11

Answer

E:

Values

- repo1
- repo1 clone URL
- repo1 repo URL
- repo2
- repo2 clone URL
- repo2 repo URL

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
...
```

Question: 442

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

Answer: A

Explanation:

Question: 443

DRAG DROP

You have an Azure subscription.

You have the standards shown in the following table.

Name	Contents
Stand a rd 1	Security requirements
Standard^	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: Eachcorrect selection is worth one point.

Answer:

Explanation:

Question: 444

HOTSPOT

You have anAzure Pipelines pipeline named Pipeline1 that has the following YAML definition.

```
pool:
  name: Azure Pipelines.
  demands:
  - msbuild
  - visualstudio

staps:
- task: VSB id Id-1
  displayNama: 'Build solution "Wiln'
  inputs:
    solution: 'HParameter#.Solution}'
    platform: ffSulldPlitfora)[1]
    configuration: 'S(BuildCon-iguration
* task: CopyFile^i
  displaying: *COpy iles to: 5(Suiid.ArtifactStag;ngDirectory l' inputs:
  Source?older; '$(Sy$tem.Defaultwork ingDirectory)*
  Contents ■ '*■ bin\f (EuildConfigurjtianJV[1]
  TargetFolder: '$(Build.ArtifactStagingPirectory)*
condition: succetdedOrFailed )
```

- task; PublishBuildArtifacts#!

```
display-Name; 'Publish Artifact: drop'
inputs:
  PathtoPublish '$(Build,ArtifactStagingDirectory ' condition: succeededorfailed-)
```

For each of the foltawing statements, select Vies if True. Otherwise select No

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Pipelinel consists of three stages	O	0
$(8uildPlatform) is a predefined variable	0	O
$(System.De-FaultWorkingOirectory) is a predefined variable		

Answer:

Explanation:

Answer Area

Statements	Yes	No
Pipelinel consists of three stages	$	0
$(8uildPl«tform) is a predefined variable.		$
{(System.DefaultWorkingDirectory) is a predefined variable		n

Question: 445

HOTSPOT

You are using an Azure Pipelines pipeline to build and deploy a web app.

You need to implement a testing strategy that meets the following requirements:

- Validates thescalability of the app
- Tests a code module without testing the module's dependencies
- Validates the interactions and dependencies between the app and its dependent services

Which type of test should you implement for each requirement? To answer,select the appropriate options in the answer area.

Answer Area

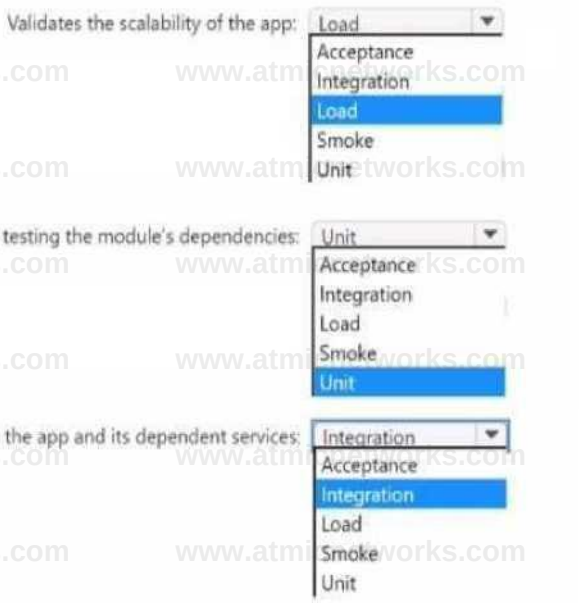

Answer:

Explanation:

Answer Area

Validates the scalability of the app: Load ’

Tests a code module without testing the module's dependencies Unit ’

Validates the interactions and dependencies between the app and its dependent services Integration ’

Question: 446

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 thesplit bar between panes or scroll to view content.

Answer:

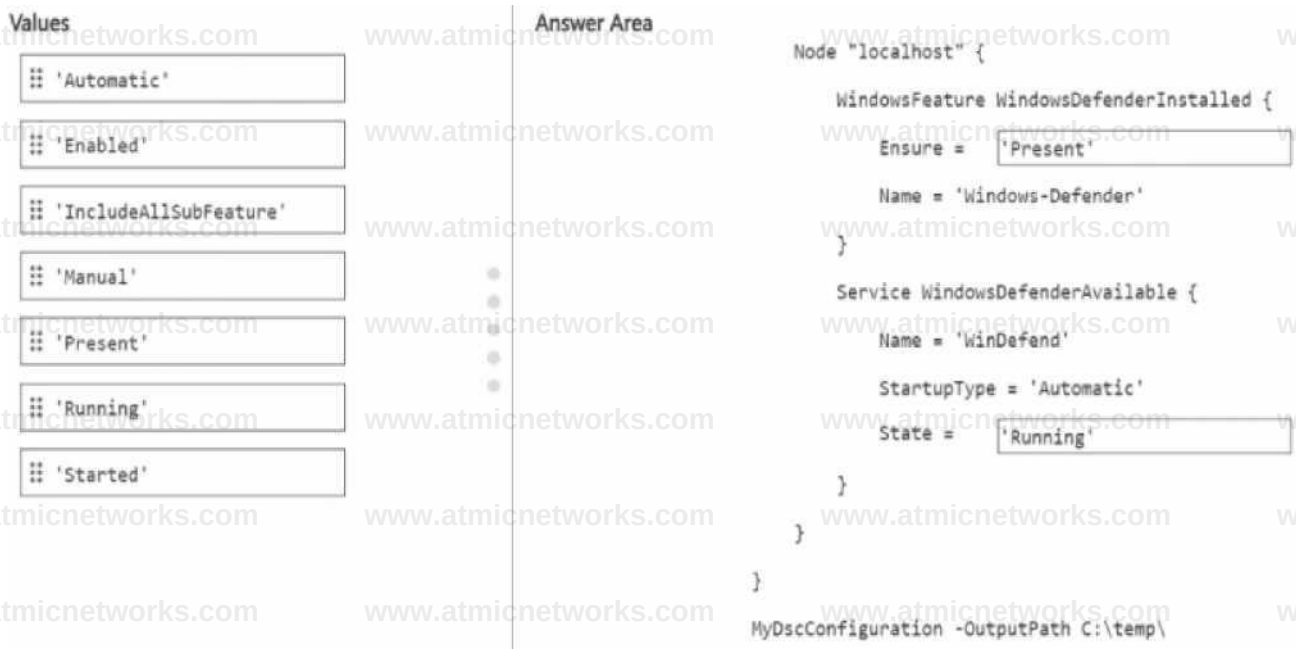

Question: 447

NOTE: Each correct selection is worth one point.

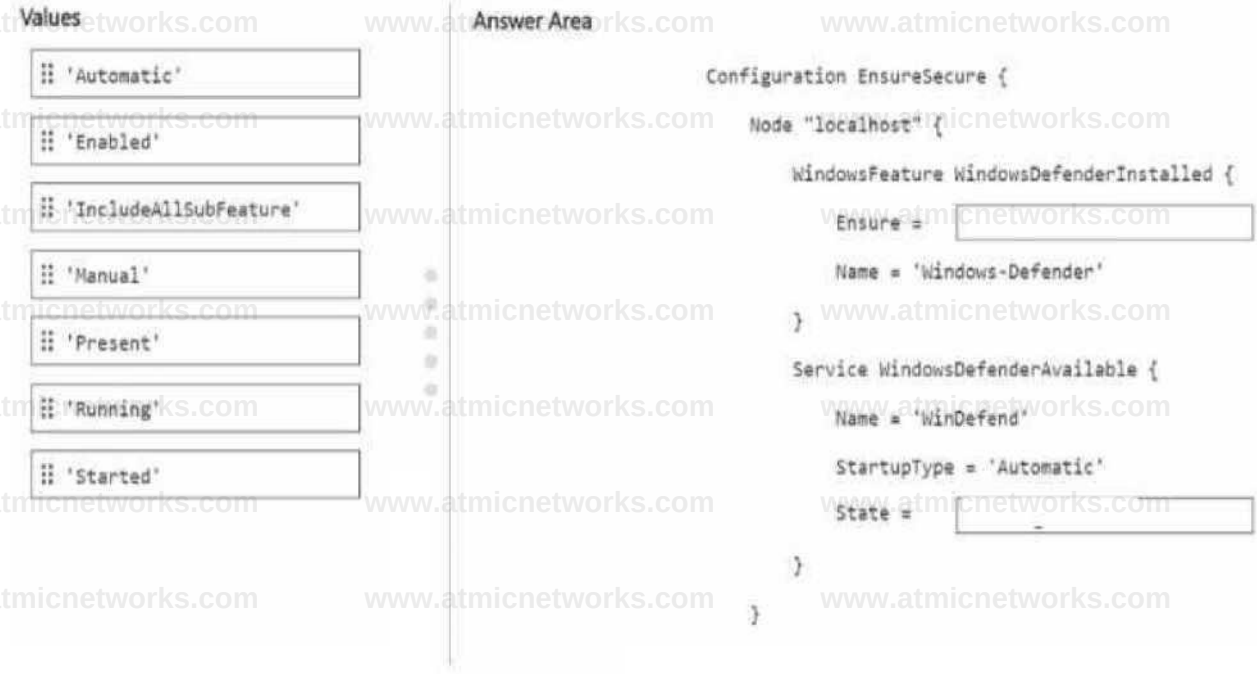

Explanation:

You manage projects by using Azure Boards. You manage project code by using GitHub. You have a work itemthat 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 Addlink, select Existing item and then enter the URL of the commit.

Answer: A,B

Explanation:

Question: 448

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.

- Thetest 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.

StkfM*

- stage: HdldjApp

jobs:

- stage; Tt&t.App

oependsOn: rBuild_Appj

jobs:

- stags Teit_0at3ba&e dtp end : On [Bu ild _App) jobs:

- stag? Pud ish_T^st Jesuits

jobs:

Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Question: 449

HOTSPOT

You have an Azure Pipelines pipeline that has the following definition.

```
- jab: A steps:
  4 jab ft steps here

- jabt £ dependsOn: A steps:
  ■*? jib E Step! !■-: 7

- job: c
  dependsOn 1 A steps;
  - jab C steps here
  - job: D depenctsDr.:
    - C stepst
    5 jab D steps here
  - job: E dependsOn:
    - 8
    steps;
    - job E steps here
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Job E can start before job C completes.	0	®
Job B can run concurrently with job C.	0	
Job D can start before job B completes.		0

Answer Area

Statements	Yes	No
Job E can start before job C completes.		
Job B can run concurrently with job C.		
Job D can start before job B completes.	n	0

Answer:

Explanation:

Question: 450

You have a project in Azure DevOps that uses an Azure Boards board and stores code in a GitHub repository. The repository contains a file named README.md.

You need to ensure thatREADME.md includes the status of the work items on the board. The solution must minimize administrative effort.

What should you do first?

A. Enable GitHub annotations for the board.
B. Install the Azure Boards app for GitHub.
C. Create a GitHub personalaccess token (PAT).
D. Select Allow anonymous users to access the status badge.

Answer: D

Explanation:

Question: 451

You are developing an Azure Pipelines pipeline.

You need to configure a check in the pipeline that will query Azure Boards to ensure that there are no active work item issues before the pipeline deploys a build to production.

Which type of check should you implement?

A. post-deployment approvals
B. manual validations
C. pre-deployment gates
D. pre-deployment approvals

Answer: C

Explanation:

Question: 452

You have an Azure pipeline that is used to deploy an app named App1.

You need to ensure that new versions of App1 are released only if they exceed performance baselines. The solution must minimize administrative effort.

What should you configure?

A. an Azure Repos branch policy
B. an Azure Pipelines release artifact
C. an Azure Pipelines deployment gate
D. an Azure Monitor alert

Answer: C

Explanation:

Question: 453

DRAG DROP

You have a GitHub repository.

You need to configure Dependabot dependency scanning. The solutionmust meet the following requirements:

- Automatically open a pull request to resolve an alert.
- Automatically open a pull request when a dependency is updated.

What should you enable for each requirement? To answer, drag the appropriate featuresto the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Question: 454

You are interrogating logs by using KQL.
You execute the query shown in the following exhibit.

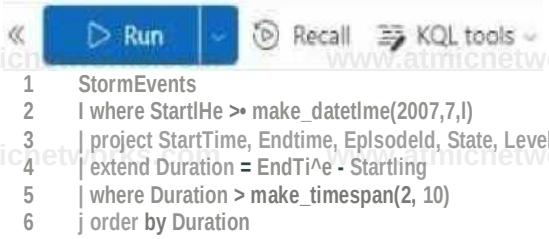

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Answer Area

(Answer choice) columns will be returned in the query output. **Five** _____ |
One

Five

Six Seven

The data type of the Duration column is (answer choice) timespan **datetime integer string**

timespan

Answer:

Explanation:

Answer Area

[Answer choice] columns will be returned in the query output. Five

The data type of the Duration column is [answer choice), timespan

Question: 455

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 valuesof the secrets from being logged.

What should you do?

A. Apply a prefix of secret to the name of the variables.
B. Pass the secrets on the command line instead of in the pipeline variables.
C. Echo the values of the secrets to the command line.
D. Storethe secrets in the environment variables instead of the pipeline variables.

Answer: D

Explanation:

Question: 456

You have an Azure Resource Manager (ARM) template that contains the following expression.

[if [pari^t£^Ci^^ 'is', is')]

You need to migrate the template to Bicep.

Which expression should you run?

Which expression should you run?

A.

iif(isCoaplete, 'la", '2a')

B.

if(isComplete, "la', '2a')

C.

if('isCo»plete' '1B else 'la

D.

isComplete ? 'la' : '2a'

Answer: D

Explanation:

Question: 457

DRAG DROP

You use an Azure pipeline to build a .NET app that has NuGet dependencies.

You need to ensure that the pipeline caches required NuGet packages.

How should you configure thepipeline? 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.

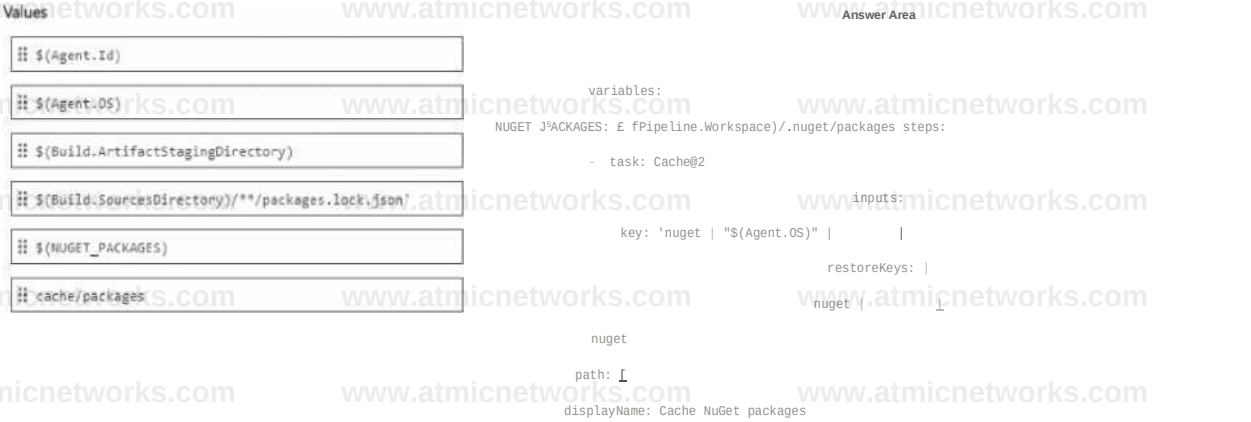

Answer:

Explanation:

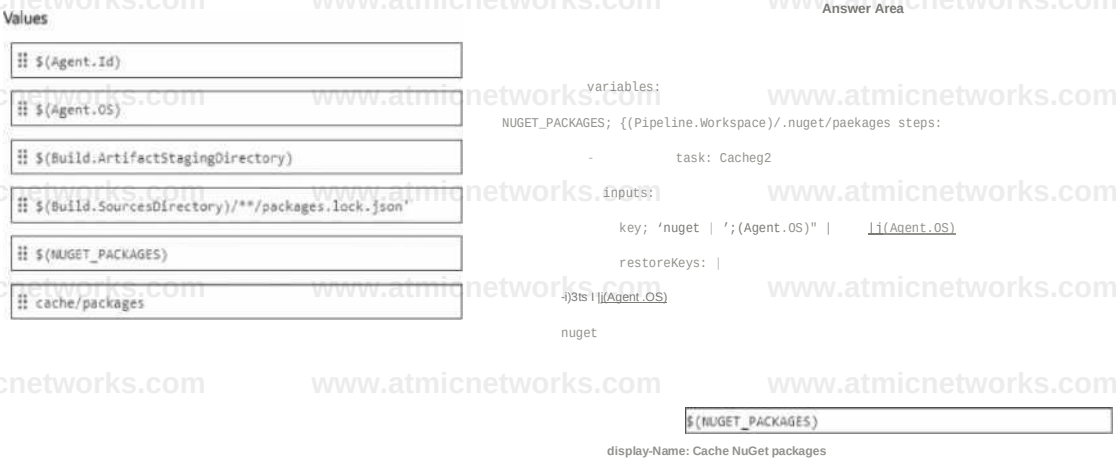

Question: 458

You have an Azure subscription that contains an Azure Pipelines pipeline named Pipeline1 and a user named User1.

Pipeline1 is used to build and test an app named Appl. User1 isassigned the Contributors role for Pipeline1.

You plan to test App1 by using an Azure Deployment Environments environment.

You need to ensure that User1 can provision the environment. The solution must follow the principle of least privilege.

Which role should you assign to User1?

A. DevCenter Project Admin
B. Deployment Environments User
C. Build Administrators
D. Contributors

Answer: B

Explanation:

Question: 459

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unit 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 if the review screen.

You use an Azure Pipelines pipeline to build and release web apps.

You need to configure the pipeline to meet the following requirements:

- Only run when there is a change in the/webapp folder.
- Only run when a pr is created.

Solution: You configure the pipeline definition by using the following elements.

```
trigger:
  paths;
    include; /wefeapp
  brand es ■
    include: nefs/hwd/pr
```

* * *

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Question: 460

Note: This question is part of a series of questionsthat present the same scenario. Each question in the series contains a unit 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 questionin this section, you will NOT be able to return to it. As a result, these questions will not appear if the review screen.

You use an Azure Pipelines pipeline to build and release web apps.

You need to configure the pipeline to meet the following requirements:

- Only run when there is a change in the /webapp folder.
- Only run when a pr is created.

Solution: You configure the pipeline definition by using the following elements.

```
trigger:
  paths.
    include: /NetApp
  braneKes:
    include: pr
```

Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

Question: 461

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unit 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 if the review screen.

You use an Azure Pipelines pipeline to build and release web apps.

You need to configure the pipeline to meet the following requirements:

- Only run when there is a change in the /webapp folder.
- Only run when a pr is created.

```
pr:
  naths:
    include: /pr
  hranches:
    inciu-ie: r zTj/i'iod irebapp
```

.

Solution: You configure the pipeline definition by using the following elements.

Does this meetthe goal?

A. Yes
B. No

Answer: B

Explanation:

Question: 462

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 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 mustbe run after successful completion of the build stage.
- The PubHsh_Test_ftesuHs 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: SuilMpp
jobs;
```

- stage: Test_App

dependsOat [Build App]

jobs:

- stage: Testjlfttobase

■.^LcrdsOn: [Build AppJ

jobs:

- stage: ³ublisr._Test^esults

jobs:

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Question: 463

You use an Azure Pipelines pipeline to build and test an app named App1.
Your company's development department works in the feature branches.
You need to ensure that a pull request will merge into the main branch only when testing covers more than 90 percent of the code, What should you do?

A. Create a Publish Test Results task.
B. Configure a branch policy for the main branch.
C. Create a code coverage configuration YAML file.
D. Configure a branch policy for thefeature branches.

Answer: C

Explanation:

Question: 464

You use GitHub Enterprise for source control repositories. The repositories store C# code. You need to enable CodeQL scanning for the repositories. What should you do?

A. Push a GitHub Actions workflow to all the repositories.
B. Enable Dependabot security updates.
C. Enable Dependabot alerts.

D. Configure a required GitHub Actions workflow for all the repositories.

Answer: D

Explanation:

Question: 465

You have a GitHub repository that uses GitHub Actions and stores accesskeys by using GitHub encrypted secrets.
You plan to update the secrets by using the GitHub REST API.
You need to wrap the secrets before adding them to a REST-based call.
Which encryption library should you are that you use?

A. CryptoNet
B. libsodium
C. hashlib
D. BouncyCastle

Answer: B

Explanation:

Question: 466

HOTSPOT
You have a GitHub repository.
You plan to use the repository to create a container.
You need to configure GitHub Actions to deploy the container to GitHub Container Registry.
How should you completethe script? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
steps:

* name: Step 1

  run:    echo JGITHUB TOKEN 1 docker login ghcr.io -u USERNAME -password-stdin
          docker build -file Dockerfile -tag ${{ github.repository )):latest -label "runnumber=J(( github.run_number}}'
    ' docker login ghcr.io -u USERNAME -password-stdin
  run:    docker pull ghcr.io/J(( github repository owner |)/J(( github.repository |):latest
          docker push ghcr io/J({ github.repository owner ))/$(( github repository )}.latest
-  name:
          echo JGITHUB.TOKEN | docker login ghcr io -u USERNAME -password-stdin
S^HMMnMMMMMVpwMB^BMB^BHBMB^^MVWWWVW^^H^^^MMmi^^H^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^M
  run:    uses azure/docker-login@v1

- name: Step 2

  run:    docker build file Dockerfile -tag SU github.repository iHatest - label “runnumber^H github.run number IT
          docker build —file Dockerfile -tag J[{ github.repository |) latest label ~runnumber=J|( github.run number ||*
          docker login ghcr.io -u USERNAME -password-stdin
  run:    docker pull ghcr.io/J(( github.repository.owner }}/$(( github repository )):latest
          docker push ghcr io/J(( github.repository.owner ))/${{ github.repository )) latest echo JGITHUB.TOKEN | docker login ghcr.io -u
          USERNAME -password-stdin uses: azure/docker-login@v1 _____

-  name: Step 3

  run:    docker push ghcr io/J{| qithub repository owner ))/$(( github repository ||:latest
          docker build file Dockerfile - tag ${{ github.repository )):latest -label "runnumber=J{{ github.run.number)}"
          docker login ghcr.io -u USERNAME -password-stdin
          docker pull ghcr.io/$|( github.repository.owner }[/${{ github.repository ||:latest _____
          docker push ghcr io/$(( github repository owner ))/$(( github repository )] latest
          echo JGITHUB.TOKEN | docker login ghcr.io -u USERNAME -password-stdin
          uses: azure/docker-login@v1
```


Answer:

Explanation:

Answer Area

```
steps:

-  name; Step 1

  run:    echo JGITHUB.TOKEN | docker login ghcr.io u USERNAME -password-stdin
-  name: .Step 2

  run:    docker build . -file Dockerfile -tag S([ github.repository ]}:latest -label"runnumber = S{{ github.run.number ))*'
-  name: Step 3

  run:    docker push ghcr.io/S{{ github.repository.owner }}/${{ github repository }}:latest
```

Question: 467

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

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

- Enable testing and monitoring of changes during a gradual rollout.
- Control the number of users that will receive new code releases.

Which strategy should you recommend?

A. A/B

B. progressive exposure

C. blue/green
D. feature toggle

Answer: A

Explanation:

Question: 468

2You have a GitHub repository.

You need to ensure that all changes to code are validated by your company's security department before the main branch is deployed.

Which two actions can you perform? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Create a branch protection rule for the feature branches.
B. Create a branch protection rule for the main branch.
C. Create a LICENSE file.
D. Require signed commits.
E. Create a CODEOWNERSfile.

Answer: B,E

Explanation:

Question: 469

You have an Azure DevOps subscription that contains the projects shown in the following table

Name	Number of users	Repository	Visibility
Projectl	100	Project! public repository	Public
Project?	5	Projectl public repository	Public
Project3	2	Private GitHub repository	Private
Project4	1,000	Public GitHub repository	Public
Projects	ISO	Public GitHub repository	Private

You build apps for the projects by using Azure Pipelines.

Which two projects meet the criteria for granting free parallelJobs? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Project1
B. Project4
C. Project2

D. Project3
E. Project5

Answer: C,D

Explanation:

Question: 470

HOTSPOT

You have an Azure subscription that contains an Azure load testing instance named Test1 and an Azure App Service web app named App1.

You implement a CI/CD pipeline named Pipeline1 that uses Test1. Pipeline1 uses the following alt- config-pipeline1.yam1 file.

version: v0.1 testId: 65e74cf8-e02e-4e71-a656-3f8649425a74 displayName: myTest testPlan: url-test.json testType: URL description: engineinstances: 2
subnetId: publicIPPisabled : 'false failurecriteria:

- MyRequest: p90(response_time_ms) >100 - MyRequest: max(response_time_ms) >250

autoStop:
errorpercentage: 9© timewindow: 60 splitAllCSVs: false

keyVaultReferenceIdentityType: SystemAssigned

ed

Pipeline! uses the followina url-test. json file

{ “version”: "1.0", "scenarios": { ”requestGroupl”; { “requests”: [{

"requestName": "MyRequest", "responsevariables": [], "queryparameters": [], "requestType": "URL", "endpoint": "https:7/www.contoso.com", "headers": {}, "body": null, "method": “GET”, "requestBodyFormat": null

1, “csvDataSetConfigList" []

“testSetup": [{ “virtualUsersPerEngine": 10, "durationInSeconds": 100, "loadType": "Linear", “scenario”: “requestGroupl”, “rampUpTimeInSeconds”: 60 }] }

For each of the following statements, select Yes if thestatement is true, otherwise select No.

Answer Area

Statements	Yes	No
Testi will terminate after 180 seconds.	O	O
Testi will fail if a request exceeds 250 ms.	O	O
Testi will run at peak load for 60 seconds.	O	O

Answer:

Explanation:

Yes
Yes No

Question: 471

You have an Azure DevOps project that produces Node Package Manager (npm) packages. Multiple projects consume the packages. You need to configure Azure Artifacts to ensure that both thelatest and pre-release versions of the packages are available for consumption. What should you do?

A. Create two feed views named @prerelease and @default. Configure a release pipeline that promotes a package to the @default view after successful testing.
B. Create two feed views named @prerelease and @release. Set @release as the default view. Configure a release pipeline that tags the packages as release after successful testing.
C. Create a feed view named @prerelease. Configure a release pipeline that tags the packages as release after successful testing.
D. Create two feed views named @prerelease and @release. Set @release as the default view. Configure a release pipeline that promotes a package to the @release view after successful testing.

Answer: C

Explanation:

Question: 472

You use Azure Pipelines to build and test code.
You need to analyze the agent pool usage.
What are two ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Reviewthe historical graph for the agent pools.
B. Review the Pipeline duration report.
C. Query the TaskAgentPoolSizeSnapshot/TaskAgentPoolSizeSnapshots endpoint.
D. Query the PipelineRun/PipelineRuns endpoint.

Answer: A,D

Explanation:

Question: 473

HOTSPOT

You have an Azure subscription that contains multiple web apps.

You need to enable Change Analysis for the web apps.

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

NOTE: Each correct selection is worth one point.

Answer Area

▼ "Microsoft.Ch»"jeAnilyjl$"

NewAzRole Assignment -ObjectType
Register-AzProviderFeature ProviderNamespace
Register-AzResourceProvider -ProviderNamespace

Swabapp_list * Gat-AstftbApp | where-Object (S_.kind -eq 'app'}

forvacn (Swebapp in Sw«b»pp_list)

Stags » Swebapp.Tags

Stags ["hidden-related;diagnostics/changeAnalysis$canEnabled"]*Strue

▼ I Swebapp.Id -Tag Stags -Force

Invoke-AzResourceAction -ResourceId
Set-AzPolicyAssignment -Identityld
Set AzResource ResourceId

Answer:

Explanation:

Register-AzResourceProvider -ProviderNamespace
Invoke-AzResourceAction -ResourceId

Question: 474

HOTSPOT

You deploy multiple apps by using Azure Pipelines.

You need to recommend a solution to manage the configuration of the apps. The solution must meet the following requirements:

- Support sharing a set of feature flags across multipleapps.
- Support deploying configurations to different deployment slots.

What should you use for each requirement? To answer, select the appropriate options in the answer are a. NOTE: Each correct selection is worth one point.

Answer Area

Support shanng a set of feature flags across multiple apps |

Azure App Configuration Azure App Service settings
Azure DevOps task groups
Azure DevOps variable groups Azure Key Vault

Support deploying configurations to different deployment slots [

Azure App Configuration Azure App Service settings Azure DevOps task groups Azure DevOps variable groups
Azure Key Vault

Answer:

Explanation:

Azure App Configuration Azure App Service settings

Question: 475

You build a Node.js app named App1 by using GitHub.
You need to test App1. The solution must ensure that App1 meets the following requirements.

- Works correctly with versions 10,12, and 14 of Node.js
- Works correctly with Ubuntu versions 22.04and 20-04
- Minimizes the runtime of the test stage

What should you use?

A. jobs
B. matrices
C. environments
D. concurrency

Answer: B

Explanation:

Question: 476

You have an Azure subscription that includes an app named App1.
You have an Azure DevOps project that contains two environments named Staging and Production.
You use Azure Pipelines to deploy App1.
You need to validate the performance of App1 in the Staging environment before it is deployed to Production. The solution must minimize administrative effort.
What should you do in the Azure DevOps project?

A. In the production branch policy, add a status check to query Azure Monitor Alerts for active alerts.
B. In the Staging environment, add a check to query Azure Monitor Alerts for active alerts.
C. In the Production environment stage, add a post-deployment approval for the Azure Monitor Alerts group
D. In the Production environment, add a check to query Azure Monitor Alerts for active alerts.

Answer: C

Explanation:

Question: 477

You have an Azure Pipelines pipeline named Pipeline1 and a user named User1. Pipeline1 contains a temporary final stage named final1. You need to ensure that User1 can delete final1 when testing is complete. The solution must follow the principle of least privilege. At which level should you grant permissions to User1?

A. organization
B. pipeline
C. project
D. stage

Answer: D

Explanation:

Question: 478

DRAG DROP

You have a GitHub repository named repo1 and an Azure key vault named kv1.

In repo1, you plan to create a workflow named Workflow 1 that will deploy adatabase server by using credentials stored in kv1.

You need to ensure that Workflow1 can retrieve the credentials from kv1.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answerarea and arrange them in the correct order.

Actions

Grant secret permissions to kv!

Create a personal access token (PAT) in GitHub.

Create a service principal in Microsoft Entra ID,

Grant key permissions to kv1.

Answer Area

Answer

Explanation:

Create a service principal in Microsoft Entra ID.
→ Needed for authentication between GitHub Actions and Azure.
Grant secret permissions to kv1.
→ Allows the service principal to read secrets from Key Vault.
Create a personal access token (PAT) in GitHub.
→ Required for GitHub to authenticate securely and use the service principal in the workflow.

Question: 479

DRAG DROP

You plan to use Azure DevOps in development processes.
You identify the following issues with the existing development processes:

- Frequent discovery of bugs in production
- Increasing detection of bugs by customers
- Slow resolution of bugs and slow resumption of production

You need to implement Key Performance Indicators (KPIs) to track each issue.
Which KPI should you implement for each issue? To answer, drag the appropriate KPIs to the correct issues. Each KPI may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Answer

Explanation:

Question: 480

HOTSPOT

You have a GitHub Enterprise account.

You have three projects named Project 1, Project2, and Project3. Project! must meet the following requirements:

- Perform concurrent testing across 700 hosts.
- Support workflow job durations of at least five hours.

Project2 must meet the following requirements:

- Perform concurrent load testing across 35 hosts that run macOS.
- Support workflow job durations of at least one hour.

Project3 must meet the following requirements:

- Perform concurrent load testing across 200 hosts run that Windows.
- Support workflow job durations of at least 24 hours.

You need to recommend an architecture for the GitHub runners of the projects. The solution must minimize costs. What should you recommend for each project? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

1[1›0^1] Self hosted GitHub runners
GitHub-hosted larger runners
Self-hosted GitHub runners
I Standard GitHub-hosted runners

Project? Self-hosted GitHub runners GitHub hosted larger runners
■If hosted GitHub runners
Standard GitHub-hosted runners

Projects: Self hosted GitHub runners GitHub-hosted larger runners Self-hosted GitHub runners
Standard GitHub-hosted runners

Answer:

Explanation:

Answer Area

Project!:	Self hosted GitHub runners
Projects	Self-hosted GitHub runners
Project3:	Self-hosted GitHub runners

Question: 481

DRAG DROP

You have a GitHub repository named Repo1.

You plan to implement continuous integration and continuous delivery (CI/CD) by using GitHub Actions.

You need to create a workflow.

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.

Answer:

Explanation:

Actions

;; Create a project

:: Create a directory named workflows

Answer Area

1 ■; Create a directory named .github/workflows.

2 :: Create a YAML file

3 :: Define the workflow in the YAMl file

4 :: Commit the YAML file lo Repo'

Question: 482

HOTSPOT

You have an Azure Pipelines pipeline that has the following definition.

```
Debug: true steps:

- task: CopyFilesf2 displayName: Copy Certificate inputs:
    Contents: certificate.pfx
    [redacted] targetFolder: '$(0uild.ArtifactStagingDirectory)'

    task: AzureKeyVaultS?
  inputs: atureSubscription: 'Vour-Azure-Subscription* KeyVaultNaoe: 'kvl'
  SecretsFilter; '•* RunAsPrelob: true

  task: AtureAppServiceSettingsai inputs:
  ature$ubscription: 'Your-Azure-Subscription' appMame: appservicel
  resourceGrouDName: rd connectionstrings: S(SecretConnectionString)

- task: AzureFunctionApp^2 inputs;
  connectedServiceHaneARH: 'Your-Azure-Subscription' appType:
  'functionAppLinux'.
  appName: 'appname' package: 'S(System.DefaultWorkingDirectory)/•"/“.zip'
  appSettings: $:appSettings)
```

For each of the following statements, select Yes if the statement is true, otherwise select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The secrets in kvl are available to CopyFiles@2.	O	0
The connection string is recorded in the log file.	O	
The secrets in kvl is available to AzureFunctionApp@2.	0	O

Answer:

Explanation:

Answer Area

Statements	Yes	No
The secrets in kv1 are available to CopyFiles@2.		
The connection string is recorded in the log file.		
The secrets in kv1 is available to AzureFunctionApp@2.		

Question: 483

You have a GitHub repository.
You need to create a tag named v3.0.5 and ensure that the tag is available in the remote repository. Which two commands should you run? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. git coamit -m 'tag v3.6.5'
B. git push origin v3.0.5
C. git tag V3.9.5
D. git push -force
E. git add 'tag v3.6.5'

Answer: B, C

Explanation:

Question: 484

You have an Azure subscription that uses Azure Boards and Azure Repos.
You have GitHub repositories.
You need to ensure that all activity in the repositories is traced by using Azure Boards. The solution must minimize development effort
What should you use?

A. work item linking
B. a project backlog
C. Application Insights in Azure Monitor
D. Delivery Plans

Answer: A

Explanation:

Question: 485

You have an infrastructure as code (IaC) solution that is managed by using Bicep files and Azure Pipelines.

You need to ensure that pipeline actions that significantly affect the existing infrastructure are reviewed before they are run.

What should you configure?

A. a release pipeline that performs a repository build validation
B. a branch policy that performs a repository check for comment resolution
C. a build pipeline with a Bicep file that performs a what-if operation
D. a build pipeline with a Bicep file that performs a deploy operation

Answer: C

Explanation:

Question: 486

Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals. More than one solution in the set might solve the problem. It is also possible that none of the solutions in the set solve the problem.

After you answer a question in this section, you will NOT be able to return. As a result, these questions do not appear on the Review Screen.

You have an on-premises web server named Server1.

You need to ensure that you can build .NET apps and deploy the apps to Server1 by using Azure Pipelines.

Solution; You deploy Managed DevOps Pools and select the Azure Pipelines - Windows Server 2022 image.

Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

Question: 487

Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals. More than one solution in the set might solve the problem. It is also possible that none of the solutions in the set solve the problem.

After you answer a question in this section, you will NOT be able to return. As a result, these questions do not appear on the Review Screen.

You have an on-premises web server named Server1.

You need to ensure that you can build .NET apps and deploy the apps to Server1 by using Azure Pipelines.

Solution: You deploy Managed DevOps Pools and select the Ubuntu Server 22.04 LTS - x64 Gen 2 image.

Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Question: 488

Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals. More than one solution in the set might solve the problem. It is also possible that none of the solutions in the set solve the problem.

After you answer a question in this section, you will NOT be able to return. As a result, these questions do not appear on the Review Screen.

You have an on-premises web server named Server1.

You need to ensure that you can build .NET apps and deploy the apps to Server1 by using Azure Pipelines.

Solution: You deploy a virtual machine scale set and select the [small disk] Windows Server 2019 Datacenter - x64 Gen 1 image.

Does this meet the goal?

A. Yes

B. No

Answer: A

Explanation:

Topic 5, Labs & Tasks

Question: 489

SIMULATION

Task 1

You need to ensure that an Azure Web App named az400-38443478-main can retrieve secrets from an Azure key vault named az400-3844J478-kv1 byusing a system managed identity The solution must use the principle of least privilege.

Answer: See the solution below in explanation.

Explanation:

To ensure that your Azure Web App namedaz400-38443478-maincan retrieve secrets from an Azure Key Vault namedaz400-3844J478-kv1using a system managed identity with the principle of least privilege, follow these detailed steps:

Enable a System Managed Identity for the Azure Web App:

Navigate to the Azure Portal.

Go to the Azure Web Appaz400-38443478-main.

SelectIdentityunder theSettingssection.

In theSystem assignedtab, switch theStatustoOn.

ClickSaveto apply the changes.

Grant the Web App Access to the Key Vault:

Go to the Azure Key Vaultaz400-3844J478-kv1.

SelectAccess policiesunder theSettingssection.

Click onAdd Access Policy.

ChooseSecret permissionsand selectGetandList. This grants the app the ability to read secrets, adhering to the principle of least privilege.

Click onSelect principal, search for your Web App nameaz400-38443478-main, and select it.

ClickAddto add the policy.

Don't forget to clickSaveto save the access policy changes.

Retrieve Secrets in the Web App Code:

In your Web App's code, use the Azure SDK to retrieve the secrets.

For example, in a .NET application,you can use theAzure.IdentityandAzure.Security.KeyVault.Secretsnamespaces.

Utilize theDefaultAzureCredentialclass which will automatically use the system managed identity when running on Azure.

```
using Azure.Identity;
usingAzure.Security.KeyVault.Secrets;

var client = new SecretClient(new Uri("https://az400-3844J478-kv1.vault.azure.net/"), new
DefaultAzureCredential());

KeyVaultSecret secret = await client.GetSecretAsync("my-secret-name");
string secretValue =secret.Value;
```

Replace"my-secret-name"with the actual name of the secret you want to retrieve.

By following these steps, your Azure Web App will be able to securely retrieve secrets from the Azure Key Vault using a system managed identity, without needingto store credentials in the code, and adhering to the

principle of least privilege. Remember to replace the placeholder names with the actual names of your Web App and Key Vault.

Question: 490

SIMULATION

Task 2

You need to create an instance of AzureApplication Insights named az400-38443478-main and configure the instance to receive telemetry data from an Azure web app named az400-38443478- main
You must create a Log Analytics workspace before this task.

Answer: See the solution below in explanation.

Explanation:

To create an instance of Azure Application Insights namedaz400-38443478-mainand configure it to receive telemetry data from an Azure web app with the same name, you'll need to follow these steps:
Create a Log Analytics Workspace:
Go to theAzure Portal.
Search for Log Analytics Workspaces and select Add.
Select a Subscription and either use an existing Resource Group or create a new one.
Provide a unique name for your Log Analytics workspace.
Choose the Region that is appropriate for you.
Review the settings and then selectCreate1.
Create an Azure Application Insights Instance:
In the Azure Portal, navigate to Application Insights.
Click on + Create.
Fill in the instance details, ensuring the name isaz400-38443478-main.
Link the instance to the Log Analytics workspace you created in the previous step.
Review and create the Application Insights instance2.
Configure the Azure Web App to Send Telemetry Data:
Go to the Azure Web Appaz400-38443478-main.
Under Monitoring, selectApplication Insights.
Choose to use an existing resource and select the Application Insights instance you created.
Follow the prompts to set up the connection, which may involve adding the appropriate SDK to your web app and configuring the connection string or instrumentation key.
Verify Telemetry Data Reception:
After setting up, send some test traffic to your web app.
Then, go to the Application Insights instance and check the Overview or Performance sections to see if

telemetry data is being received.

Remember to replace placeholder names with the actual names of your resources where necessary.

These steps will help you set up Azure Application Insights to monitor your web app effectively.

Question: 491

SIMULATION

Task 3

You need to ensure that an Azure Web App named az400-38443478-matn supports rolling upgrades The solution must ensure that only 10 percent of users who connect to az400-38443478 main use updated versions of the app. The solution must minimize administrative effort.

Answer: See the solution below in explanation.

Explanation:

To ensure that your Azure Web App namedaz400-38443478-mainsupports rolling upgrades and only 10 percent of users connect to the updated version of the app, you can use deployment slots with the followingsteps:

Create a Deployment Slot:

Navigate to the Azure Portal.

Go to your Web Appaz400-38443478-main.

SelectDeployment slotsin the menu.

Click onAdd Slot.

Name the slot (e.g.,staging) and if needed, clone settings from the production slot.

Configure the Traffic Percentage:

In the Deployment Slots menu, you will see a column forTraffic %.

Set the traffic percentage to10%for thestagingslot1.

This will route only 10% of the traffic to the updated version of the app in thestagingslot.

Deploy the Updated App to the Staging Slot:

Deploy yourupdated application to thestagingslot.

Test the application in thestagingslot to ensure it's working as expected.

Complete the Rolling Upgrade:

Once you're satisfied with the performance and stability of the app in thestagingslot, you can graduallyincrease the percentage of traffic until you're ready to swap with the production slot. To swap slots, go to theDeployment slotsmenu and click onSwapwith the production slot.

By using deployment slots, you can achieve rolling upgrades with minimal administrative effort, as it allows you to test the new version on a subset of users before fully releasing it. Remember to adjust the traffic percentage and monitor the application's performance throughout the process.

Question: 492

SIMULATION

Task 4

Youneed to configure a virtual machine template in a DevTest Labs environment named az400- 38443478-dtl1. The operating system must be based on Windows Server 2016 Datacenter. Virtual machines created from the DevTest Lab must include the Selenium tool and theGoogle Chrome browser.

Answer: See the solution below in explanation.

Explanation:

To configure a virtual machine template in your DevTest Labs environment namedaz400-38443478- dtl1with Windows Server 2016 Datacenter that includes the Selenium tooland the Google Chrome browser, follow these steps:

Create a Custom Image with Windows Server 2016 Datacenter:

In the Azure Portal, go to your DevTest Labaz400-38443478-dtl1.

Navigate toConfiguration and policies>Custom images.

Use an existing VM or create a new one with Windows Server 2016 Datacenter.

After setting up the VM, capture it to create a custom image1.

Install Selenium and Google Chrome on the VM:

Connect to the VM via RDP.

Download and install the Selenium WebDriver for your preferred programming language from the official Selenium website2.

For Google Chrome, download the ofline installer from the official website and install it on the VM3.

Generalize the VM:

Run thesysprepcommand to generalize the VM, which prepares it to be used as a template.

Shutdown the VM aftersysprepcompletes.

Capture the Generalized VM to Create a Template:

In the Azure Portal, navigate to the VM and selectCapture.

Provide the required details and create the image.

Add Selenium and Google Chrome Artifacts to the Template:

Go back to the DevTest Labaz400-38443478-dtl1.

SelectArtifactsand add Selenium and Google Chrome artifacts to the template.

Ensure these artifacts are configured to install during the VM creation process.

Create VMs from the Template:

Now, when youcreate a new VM in the DevTest Lab, select the custom image you created.

The VM will be provisioned with Windows Server 2016 Datacenter, and the Selenium tool and Google Chrome browser will be installed automatically.

By following these steps, you can ensure that all virtual machines created from this template in your DevTest Lab will have the required operating system, tools, and browser installed. Remember to replace placeholder

names with the actual names of your resources where necessary.

Question: 493

SIMULATION

Task 5

You plan to stoic signed images in an Azure Container Registry instance named az40038443478act1.

You need to modify the SKU for az40038443478aa1 to support the planned images. The solution must minimize costs.

Answer: See the solution below in explanation.

Explanation:

To store signed images in an Azure Container Registry (ACR) instance and support your planned images while minimizing costs, you need to modify the SKU of your ACR instance to one that supports content trust and image signing. Here's how you can do it: Determine the Appropriate SKU:

Content trust and image signing are features of thePremiumservice tier of Azure ContainerRegistry1. If cost minimization is a priority, ensure that the Premium tier is necessary for your use case. If you require content trust, the Premium tier is the appropriate choice.

Modify the SKU of the ACR Instance:

Navigate to the Azure Portal.

Go to your ACR instanceaz40038443478act1.

SelectUpdatefrom the overview pane.

Choose thePremiumSKU from the SKU drop-down menu2.

Review the changes and pricing, then save the configuration.

By upgrading to the Premium SKU, you'll be able to store signed images in your ACR instance.

Remember to monitor your usage and costs to ensure they align with your budget and requirements.

Question: 494

SIMULATION

Task 6

You have an Azure function hosted in an App Service plan named az400-38443478-funct

You need to configure az400 38443478-funcl lo update the functions automatically whenever new code is

committedto the main branch of https://githubcom/Azure-Samples/functions-quickstart. NOTE: Access to GitHub is blocked in the exam environment. Access to GitHub is NOT required to complete the task successfully.

Answer: See the solution below in explanation.

Explanation:

To configure your Azure Functionaz400-38443478-functto automatically update whenever new code is committed to the main branch of the specified GitHub repository, you can use GitHub Actions for continuous deployment. Here's how to set it up:

Create a GitHub Actions Workflow:

In your GitHub repository, navigate to the.github/workflows/directory.

Create a new file for your workflow (e.g.,azure-function-cd.yml).

Define the Workflow:

In the workflow file, define the steps for the build and deployment process.

Use theAzure/functions-actionto deploy to your Azure Function App.

Set up triggers for themainbranch to initiate the workflow on every commit.

Generate Deployment Credentials:

In the Azure Portal, navigate to your Function Appaz400-38443478-funct.

Download the publish profile from theOverviewsection by clicking onGet publish profile.

Store the Publish Profile as a GitHub Secret:

In your GitHub repository, go toSettings>Secrets and variables>Actions.

Create a new secret (e.g.,AZURE_FUNCTIONAPP_PUBLISH_PROFILE) and paste the content of the publish profile.

Configure the Workflow to Use the Secret:

In the workflow file, reference the secret to authenticate the deployment to Azure.

Here's a sample GitHub Actions workflow snippet:

```
name: Deploy Azure Function

on:
push:
branches:
- main

jobs:
build-and-deploy:
runs-on: ubuntu-latest
steps:
- uses: actions/checkout@v2
- name: Set up Python version
uses: actions/setup-python@v2
```

with:
python-version: '3.x'

- name: Install dependencies
run: |
pip install -r requirements.txt

- name: Deploy to Azure Functions
uses: Azure/functions-action@v1
with:
app-name: az400-38443478-funct
publish-profile: ${{ secrets.AZURE_FUNCTIONAPP_PUBLISH_PROFILE }}
package: .
Replace theapp-namewith the name of your Azure Function App and ensure the Python version and dependencies match your application's requirements.
By following these steps, yourAzure Function will automatically update whenever new code is pushed to the main branch of the GitHub repository. This setup minimizes manual effort and ensures that your function app is always running the latest code.

Question: 495

SIMULATION

Task 7
You need to prepare a network security group (NSG) named az400 38443478 nsgl 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

Answer: See the solution below in explanation.

Explanation:

To prepare a Network Security Group (NSG) namedaz400-38443478-nsg1for hosting an Azure DevOps pipeline agent, while allowing only the required outbound port for Azure DevOps and denying all other inbound and outbound access to the Internet, follow these steps:
Create the NSG:
Navigate to the Azure Portal.
Go toNetwork Security Groupsand click on+ Create.
Fill in the details, including the nameaz400-38443478-nsg1, and create the NSG.
ConfigureOutbound Security Rules:

Once the NSG is created, go to its settings.

Navigate toOutbound security rules.

Click on+ Addto create a new rule.

Set theDestination port rangesto443, which is the required port for Azure DevOps12.

Set theProtocoltoTCP.

Set theActiontoAllow.

Assign aPrioritynumber (e.g.,100) that does not conflict with existing rules.

Provide a meaningfulNamefor the rule (e.g.,AllowAzureDevOps).

Configure Default Rules to Deny All Other Traffic:

In the sameOutbound security rulessection, edit the default rule to deny alltraffic.

Change theActiontoDenyfor the rule with the lowest priority (highest number).

Ensure that this rule applies to all protocols, source and destination IP ranges, and port ranges.

Associate the NSG with the Appropriate Resource:

Associate the NSGwith the subnet or network interface of the virtual machine or resource where the Azure DevOps pipeline agent will be hosted.

By following these steps, you will ensure that the Azure DevOps pipeline agent can communicate with Azure DevOps services over the required port while blocking all other inbound and outbound Internet access, adhering to the principle of least privilege and security best practices.

Question: 496

SIMULATION

Task 8

You need to ensure that the https://contoso.com/statushook webhookis called every time a repository named az40038443478acr 1 receives a new version of an ima ge named dotnetapp

Answer: See the solution below in explanation.

Explanation:

To ensure that the webhook athttps://contoso.com/statushookis called every time the repository namedaz40038443478acr1receives a new version of an image nameddotnetapp, you can follow these steps to configure a webhook in Azure Container Registry:

Navigate to the Azure Container Registry:

Go to the Azure Portal.

Find and select yourAzure Container Registry instanceaz40038443478acr1.

Create a New Webhook:

UnderServices, selectWebhooks.

Click on+ Addto create a new webhook.

Fill in the form with the following information:

Webhook name: Enter a unique name for your webhook.

ServiceURI: Enterhttps://contoso.com/statushook.

Custom headers: (Optional) Add any headers you want to pass along with the POST request.

Trigger actions: SelectPushto trigger the webhook on image push events.

Scope: Specify the scope asaz40038443478acr1:dotnetappto target the specific image.

Status: Set toEnabled.

Save the Webhook Configuration:

Review the information and clickCreateto save the webhook.

Once configured, the webhook will send a POST request tohttps://contoso.com/statushookwhenever a new version of thedotnetappimage is pushed to theaz40038443478acr1repository in your Azure Container Registry1.

This setup will automate the notification process, ensuring that the specified webhook is called with each new image version, thus fulfilling the task requirements.

Question: 497

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 userthat has an email address of Usfrr2-42147S09@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 asUserl-42147509©ExamUsers.com.

Answer: See the solution below in explanation.

Explanation:

Step 1: Sign Up for Azure DevOps

Navigate to Azure DevOps.

Click on Start 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 on Continue to create the organization.
Step 3: Create a Private Project
In your new organization, click on New 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 Readers group.
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 Users list.
Confirm that the user has the Stakeholder access level and is part of the Readers group.
By following thesesteps, you should be able to complete the task successfully. If you encounter any issues, feel free to ask for further assistance!

Question: 498

SIMULATION
Task 2
For the RG1 lod42147S09 resource group, you need to create an action group named DevOpsAGthat 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.

Answer: See the solution below in explanation.

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 ActionGroup:
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: Select Email/SMS message/Push/Voice.
Action Name: Enter a name for the action (e.g., NotifyOwners).
Email: Select Notify all owners.
Click OK.
Step 3: Enable Common Alert Schema
Common Alert Schema:
In the Advanced tab, enable the Common alert schema option1.
Step 4: Review and Create
Review:
Review all the settings you have configured.
Create:
Click on Review + create and then Create.
By following these steps, you will have successfully created an action group named DevOpsAG that emails the specified users and notifies resource owners using the common alert schema

Question: 499

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

Answer: See the solution below in explanation.

Explanation:

Step1: Create a New Team Dashboard
Navigate to Azure DevOps:
Go to Azure DevOps and sign in with your credentials.
Select Your Project:
Choose Project1 from your list of projects.
Access Dashboards:

In the left-hand menu, select Dashboards.

Create a New Dashboard:

Click on New Dashboard.

Enter the name Dashboard1.

Ensure the dashboard type is set to Team Dashboard.

Click Create.

Step 2: Add the Team Members Widget

Open the Widget Catalog:

After creating the dashboard, thewidget catalog will open automatically. If it doesn't, click on Add Widget.

Search for Team Members Widget:

In the widget catalog, search for Team Members.

Add the Widget:

Click on the Team Members widget and then click Add to 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 on Save to save your changes.

Share the Dashboard:

You can share the dashboardURL 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

Question: 500

SIMULATION

Task 4

Initialize the default main branch, if it does not exist already.

In the UserI-42147509 Azure DevOps project, you need to implement an approval process for the default main branch.

The process must use the four-eyes principle and have at least one approval on every iteration.

Answer: See the solution below in explanation.

Explanation:

Step 1: Initialize the Default Main Branch

Navigate to Azure DevOps:

Go to Azure DevOps and sign in with your credentials.

Select Your Project:

Choose Project1 from your list of projects.

Initialize the Main Branch:

Go to Repos > Files.

If themain branch does not exist, you will see an option to initialize it. Click on Initialize and follow the prompts to create the main branch.

Step 2: Implement an Approval Process for the Main Branch

Navigate to Branch Policies:

Go to Repos > Branches.

Findthe main branch and click on the ... (ellipsis) next to it.

Select Branch policies.

Enable Required Reviewers:

Under Policies, enable Minimum number of reviewers.

Set the minimum number of reviewers to 2 to enforce the four-eyes principle.

Add RequiredReviewers:

Add the users who should review the changes. Ensure that at least one approval is required on every iteration.

Enable Reset Code Reviewer Votes:

Enable the Reset code reviewer votes when there are new changes option to ensure that any new changes require re-approval.

Save Changes:

Click on Save changes to apply the policies.

Step 3: Verify the Approval Process

Create a Pull Request:

Make a change in a branch and create a pull request to merge it into the main branch.

Review and Approve:

Ensurethat the pull request requires at least two reviewers to approve it before it can be merged. By following these steps, you will have successfully initialized themain branch and implemented an approval process that adheres to the four-eyes principle, ensuring that every iteration has at least one approval

Question: 501

SIMULATION

Task 5

For Project1, you need to create a project wiki named Wiki1 that uses the Mermaid syntax to render a diagram A sample of the desired output is stored in C:\Resources\TCPHandshake.png.

Answer: See the solution below in explanation.

Explanation:

Step 1: Create a Project Wiki

Navigate to Azure DevOps:

Go to Azure DevOps and sign inwith your credentials.

Select Your Project:

Choose Project1 from your list of projects.

Create a Wiki:

In the left-hand menu, select Wiki.

Click on Create project wiki.

Enter the name Wiki1 and click Create.

Step 2: Add Mermaid Syntax to Render a Diagram

Open the Wiki Page:

Navigate to the newly created Wiki1.

Edit the Wiki Page:

Click on Edit to start editing the wiki page.

Insert Mermaid Diagram:

Use the following Mermaid syntax to render a diagram. For example, to render a simple flowchart, you canuse:

```mermaid
graph TD;
A-->B;
A-->C;
B-->D;
C-->D;
```

Save the Page:

Click on Save to save your changes.

Step 3: Render the TCP Handshake Diagram

Convert TCPHandshake.png to Mermaid Syntax:

Since you have a sample diagram in C:\Resources\TCPHandshake.png, you need to convert this diagram into Mermaid syntax. Here's an example of how a TCP handshake might look in Mermaid syntax: ```mermaid

```
sequenceDiagram
participant Client
participant Server
Client->>Server: SYN
Server-->>Client: SYN-ACK
Client->>Server: ACK
```

Add the Diagram to the Wiki:

Replace the sample Mermaid syntax with the TCP handshake diagram syntax in the wiki page.

Save the Page:

Click on Save to save your changes.

By following these steps, you will have created a project wiki named Wiki1 and used Mermaid syntax to render a diagram

Question: 502

SIMULATION

Task 6

Initialize the default main branch, if it does not exist already.

For Project 1, you need to implement branch merging restrictions to enable squash merge for all changes merged into the main branch.

Answer: See the solution below in explanation.

Explanation:

Step 1: Initialize the Default Main Branch

Navigate to Azure DevOps:

Go to Azure DevOps and sign in with your credentials.

Select Your Project:

Choose Project1 from your list of projects.

Initialize the Main Branch:

Goto Repos > Files.

If the main branch does not exist, you will see an option to initialize it. Click on Initialize and follow the prompts to create the main branch1.

Step 2: Enable Squash Merge for the Main Branch

Navigate to Branch Policies:

Go to Repos > Branches.

Find the main branch and click on the ... (ellipsis) next to it.

Select Branch policies.

Enable Squash Merge:

Under Policies, scroll down to the Merge strategy section.

Select Squash merge as the required mergestrategy2.

Save Changes:

Click on Save changes to apply the policies.

Step 3: Verify the Squash Merge Policy

Create a Pull Request:

Make achange in a branch and create a pull request to merge it into the main branch.

Complete the Pull Request:

Ensure that the pull request uses thesquash merge strategy by selecting Squash commit under the Merge type in the Complete pull request dialog

Question: 503

SIMULATION

Task 7

Initialize the default main branch, if it does not exist already

In the User 1-42147509 organization, you need to install the Microsoft Security DevOps extension.

Next, create a new starter pipeline named starter1 that will use the following starter code.

```
trigger:
    - Smarter
```

pool:

■ ubuntu-latest allO supported.

vtnlmage: 'i^indp^^-latest'

steps:

Ensure that starter! includes a task that executes the extension and uses the following input*:

- Command: run
- Policy aruredevops
- Publish: true

Save the pipeline to anew branch named starter

Answer: See the solution below in explanation.

Explanation:

Step 1: Initialize the Default Main Branch

Navigate to Azure DevOps:

Go to Azure DevOps and sign in with your credentials.

SelectYour Project:

Choose Project1 from your list of projects.

Initialize the Main Branch:

Go to Repos > Files.

If the main branch does not exist, you will see an option to initialize it. Click on Initialize and follow the prompts to create the main branch1.

Step 2: Install the Microsoft Security DevOps Extension

Navigate to Extensions:

In Azure DevOps, click on the Shopping Bag icon in the top right corner and select Browse Marketplace.

Search for the Extension:

Search for Microsoft Security DevOps.

Install the Extension:

Click on Get it free.

Select your organization (User1-42147509) and click Install.

Follow the prompts to complete the installation2.

Step 3: Create a New Starter Pipeline

Navigate to Pipelines:

Go to Pipelines > New pipeline.

Select the Repository:

Choose Azure Repos Git and select the relevant repository.

Configure the Pipeline:

Select Starter pipeline and replace the default YAML with the following starter code:

```
trigger:
- starter

pool:
vmImage: 'windows-latest'

steps:
- task: MicrosoftSecurityDevOps@1
inputs:
```

command: 'run'
policy: 'azuredevops'
publish: true

Save the Pipeline:
Click on Save and enter starter as the branch name.
Click on Save and run to save the pipeline to the new branch named starter3.
By following these steps,you will have successfully initialized the main branch, installed the Microsoft Security DevOps extension, and created a new starter pipeline named starter1 that includes the specified task

Question: 504

SIMULATION

Task 8
Initialize the default main branch,if it does not exist already.
In Project 1, you need to create a new Azure Pipelines YAML pipeline by using the ASP.NET template. The pipeline must use Azure Repos as the hosting platform and must be created in a new branch named azure-pipelines.

Answer: See the solution below in explanation.

Explanation:

Step 1: Initialize the Default Main Branch
Navigate to Azure DevOps:
Go to Azure DevOps and sign in with your credentials.
Select Your Project:
Choose Project1 from your list of projects.
Initialize theMain Branch:
Go to Repos > Files.
If the main branch does not exist, you will see an option to initialize it. Click on Initialize and follow the prompts to create the main branch1.
Step 2: Create a New Branch for the Pipeline
Navigate to Branches:
Go to Repos > Branches.
Click on New branch.
Create the Branch:
Enter azure-pipelines as the branch name.
Select main as the base branch.
Click Create2.
Step 3: Create a New Azure Pipelines YAML Pipeline
Navigate to Pipelines:
Go to Pipelines > New pipeline.
Select the Repository:
Choose Azure Repos Git and select the relevant repository.

Configure the Pipeline:
Select Starter pipeline.
Replace the default YAML with the ASP.NET template. You can find the ASP.NET template in the Azure Pipelines documentation oruse the following example:

```
trigger:
- main

pool:
vmImage: 'windows-latest'

variables:
buildConfiguration: 'Release'

steps:
- task: UseDotNet@2
inputs:
packageType: 'sdk'
version: '5.x'
installationPath: $(Agent.ToolsDirectory)/dotnet

- script: |
dotnet build --configuration $(buildConfiguration)
displayName: 'Build project'

- script: |
dotnet test --no-build --configuration $(buildConfiguration)
displayName: 'Run tests'
```

Save the Pipeline:
Click on Save and enter azure-pipelines as the branch name.
Click on Save and run to save the pipeline to the new branch named azure-pipelines3.
By following these steps, you will have successfully initialized the main branch, created a new branch named azure-pipelines, and set up a new Azure Pipelines YAML pipeline using the ASP.NET template

Question: 505

SIMULATION

Task 9

In Project1, you need to create a variable group named varGroup1 that will contain the following variables:
serverNane: server1 dbName: db1

Answer: Seethe solution below in explanation.

Explanation:

Step 1: Navigate to the Library Navigate to Azure DevOps:

Go to Azure DevOps and sign in with your credentials.

Select Your Project:

Choose Project1 from your list of projects.

Access the Library:

In the left-hand menu, select Pipelines > Library.

Step 2: Create a Variable Group

Create a New Variable Group:

On the Library page, click on + Variable group.

Configure the Variable Group:

Name: Enter varGroup1.

Description: Optionally, add a description for thevariable group.

Add Variables:

Click on + Add to add a new variable.

Variable Name: Enter serverName.

Value: Enter server1.

Click OK.

Click on + Add again to add another variable.

Variable Name: Enter dbName.

Value: Enter db1.

Click OK.

Save the Variable Group:

Click on Save to save the variable group.

By following these steps, you will have successfully created a variable groupnamed varGroup1 containing the specified variables

Question: 506

SIMULATION

Task 10

For Project1, you need to ensure that artifacts, symbols, and attachments are retained for 60 days.

Answer: See the solution below in explanation.

Explanation:

Step 1: Navigate to Project Settings

Navigate to Azure DevOps:

Go to Azure DevOps and sign in with your credentials.

Select Your Project:

Choose Project1 from your list of projects.

Access Project Settings:

In the left-hand menu, scroll down and select Project settings.

Step 2: Configure Retention Policies

Navigate to Pipelines Settings:

Under Pipelines, select Settings.

Set Retention Policies:

In the Retention section, set the number of days to keep artifacts, symbols, and attachments to 60 days.

Ensurethat the retention policy is applied to all relevant pipelines and branches.

Save Changes:

Click on Save to apply the retention policy1.

By following these steps, you will have successfully configured theretention policy to retain artifacts, symbols, and attachments for 60 days in Project1

Question: 507

SIMULATION

Task 11

You need to write a KQL query that will count the number of inbound requests for each source IP address, for any connection made during the last three months of 2021.

Open Azure Data Explorer by using the following link:

https//dataexploter azure.com/clusters/help/databases/Securiitylogs

The requests are contained in a table named InboundBrowsing in the Securitylogs connection.

The query must return two columns named NumberOfRequests and SourceIP.

Export the query result to C:\Samples

Answer: See the solution below in explanation.

Explanation:

Step 1: Write the KQL Query

Open Azure Data Explorer:

Navigate to Azure Data Explorer and sign in with your credentials.

Access the Securitylogs Database:

Open the Securitylogs database.

Write the Query:

Use the following KQL query tocount the number of inbound requests for each source IP address:

```
InboundBrowsing
| where Timestamp between (datetime(2021-10-01) .. datetime(2021-12-31))
| summarize NumberOfRequests = count() by SourceIP
| project SourceIP, NumberOfRequests
```

Step 2: Export the Query Results

Run the Query:

Execute the query in Azure Data Explorer.

Export the Results:

Once the query results are displayed, click on the Export button.

Choose the export format (e.g., CSV) and specify the export path as C:\Samples.

By following these steps, you will have successfully written a KQL query to count the number of inbound requests for each source IP address during the last three months of 2021 and exported the results to C:\Samples

Question: 508

SIMULATION

Task 12

You need to create a personal access token (PAT) named Token! that has only the following capabilities

- Read write, andmanage code
- Read and execute builds
- Read releases

Token1 must expire in 60 days.

Answer: See the solution below in explanation.

Explanation:

Step 1: Navigate to Personal Access Tokens

Sign in to Azure DevOps:

Go to Azure DevOps and sign in with your credentials.

Access User Settings:

Click on your profile picture in the top right corner.

Select User settings.

Open Personal Access Tokens:

In the user settings menu, select Personalaccess tokens.

Step 2: Create a New Personal Access Token

Create a New Token:

Click on + New Token.

Configure the Token:

Name: Enter Token1.

Organization: Select the organization where you want to use the token.

Expiration: Set the expiration to 60 days.

Set Scopes:

Code: Select Read, Write, & Manage.

Build: Select Read & Execute.

Release: Select Read.

Create the Token:

Click on Create.

Step 3: Save the Token

Copy the Token:

Once the token is created, copy it immediately as it will not be displayed again.

Store the token in a secure location.

By following these steps, you will have successfully created a personal access token named Token1 with the specified capabilities and a 60-day expiration

Question: 509

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.

Thesolution must use the principle of least privilege.

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

Answer: See solution below.

Explanation:

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

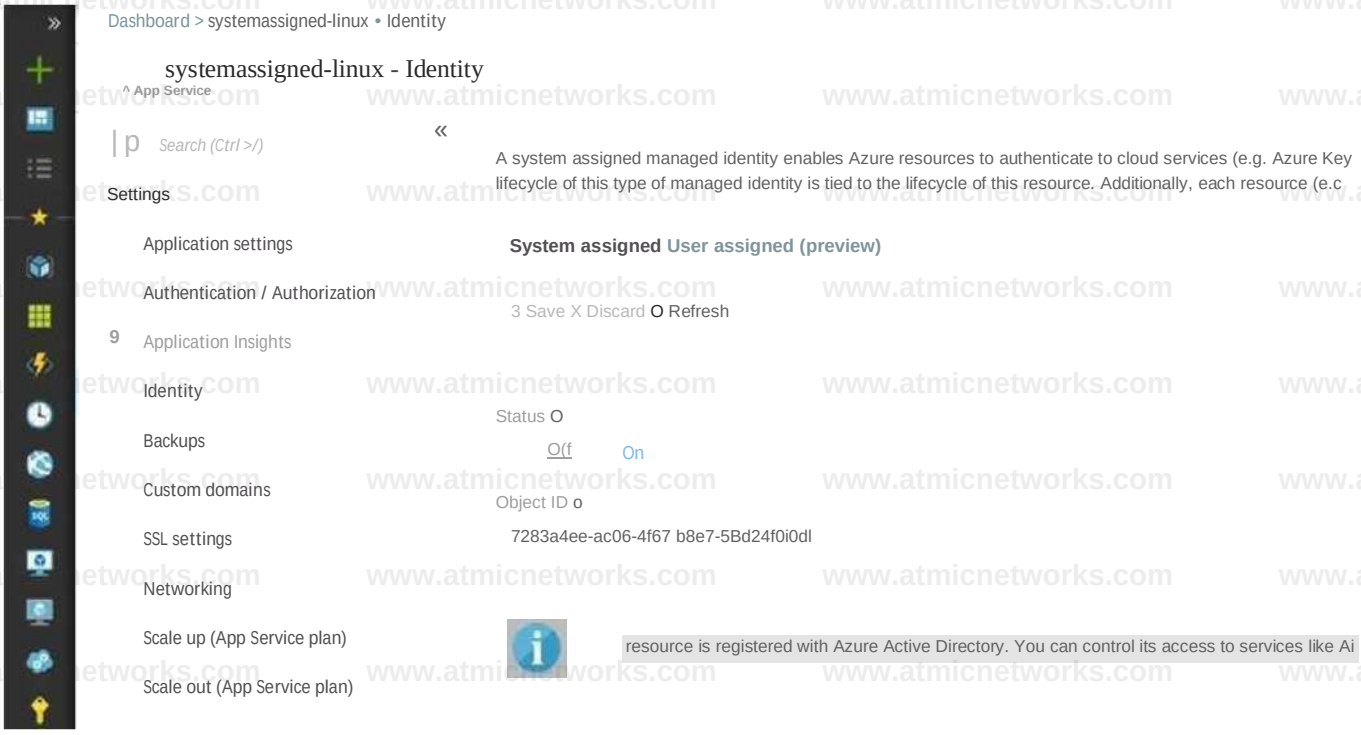

Reference:
https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity

Question: 510

SIMULATION

You have an Azure functionhosted 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.

Answer: See solution below.

Explanation:

1. Open Microsoft Azure Portal

2. Log into your Azure account, select App Services in the Azure portal left navigation, and then select configureaz400-9940427-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 Deployto 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: 511

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.

Answer: See solution below.

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, selectDeployment slots > Add Slot.

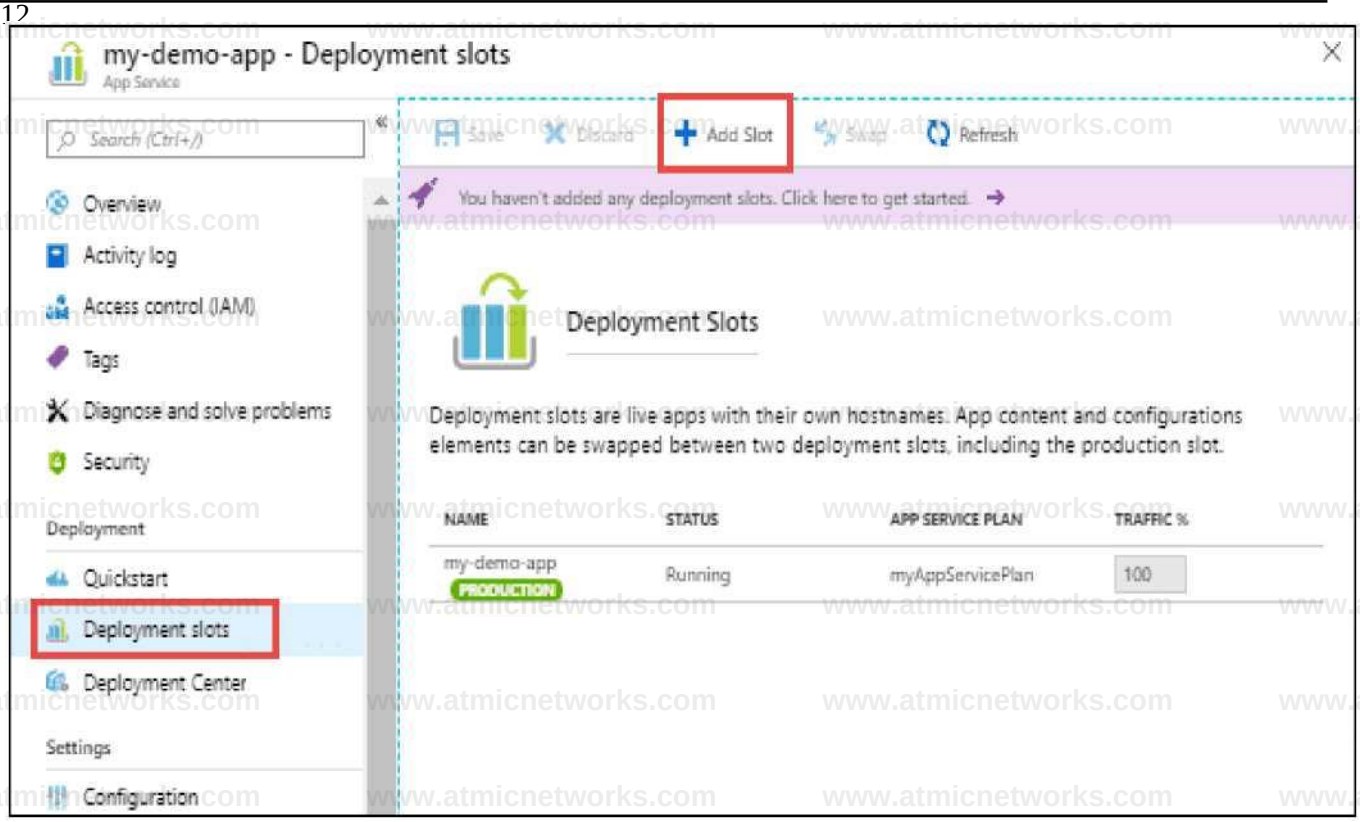

3. In the Add a slot dialog box, give the slot a name, and select whether to clone an app configuration from another deployment slot. Select Add to continue.

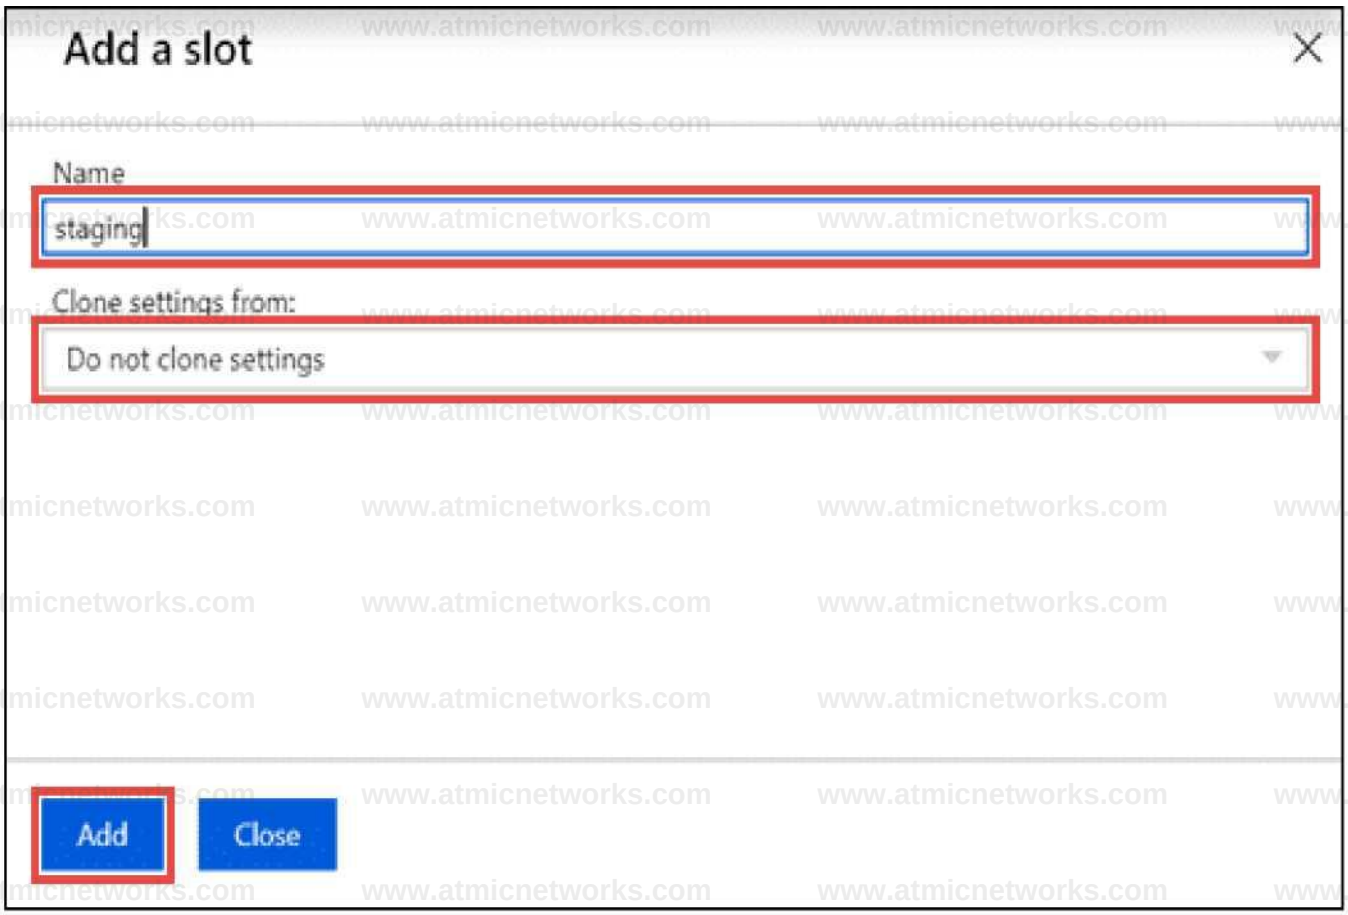

4. After the slot is added, select Close to close the dialogbox. The new slot is now shown on the Deployment slots page.

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.

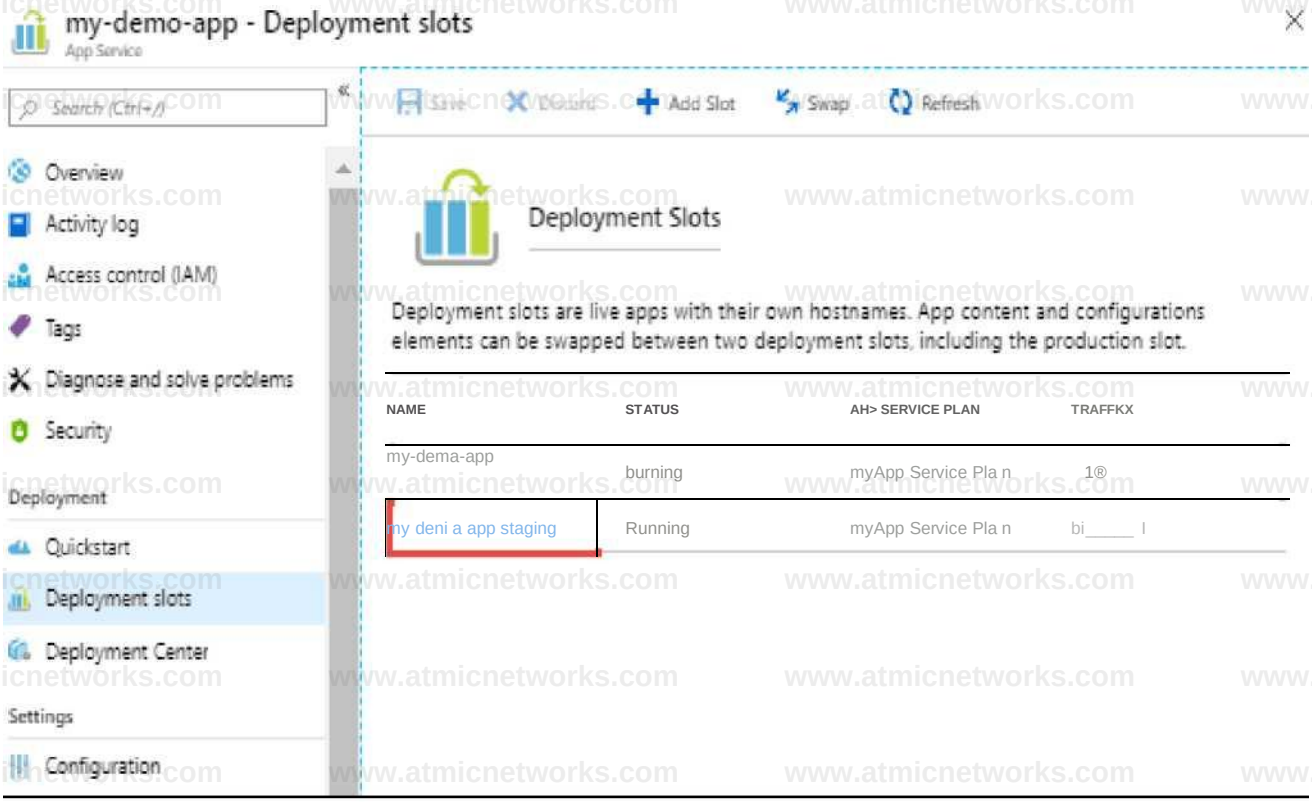

6. Change TRAFFIC % to 10

Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots

Question: 512

SIMULATION

You plane to store signed images in anAzure 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.

Answer: See solution below.

Explanation:

1. Open Microsoft Azure Portal, and select the Azure Container Registry instance named az4009940427acr1.

2. Under Policies, select Content Trust > Enabled > Save.

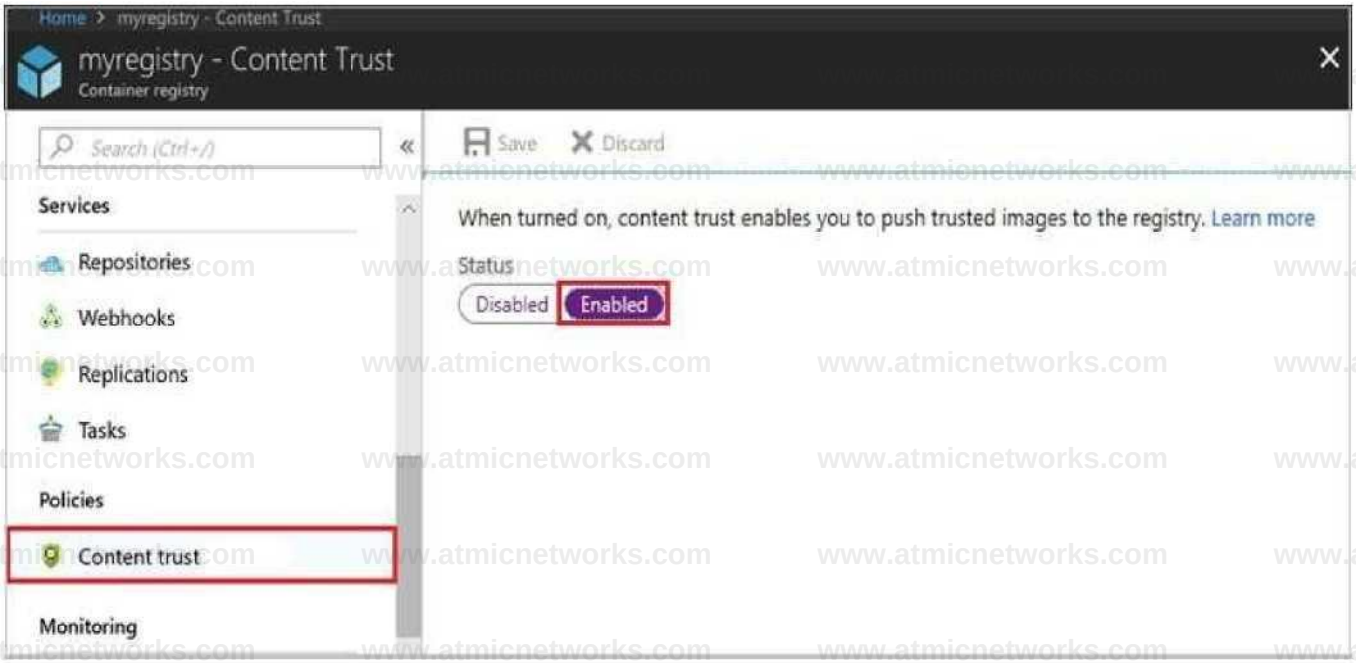

Reference:
https://docs.microsoft.com/en-us/azure/container-registry/container-registry-content-trust

Question: 513

SIMULATION

You need to create a virtual machine template in an Azure DevTestLabs environment named az400- 9940427-dtl1. The template must be based on Windows Server 2016 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.

Answer: See solution below.

Explanation:

1. Open Microsoft Azure Portal

2. Select All Services, and then select DevTest Labs in the DEVOPS section.

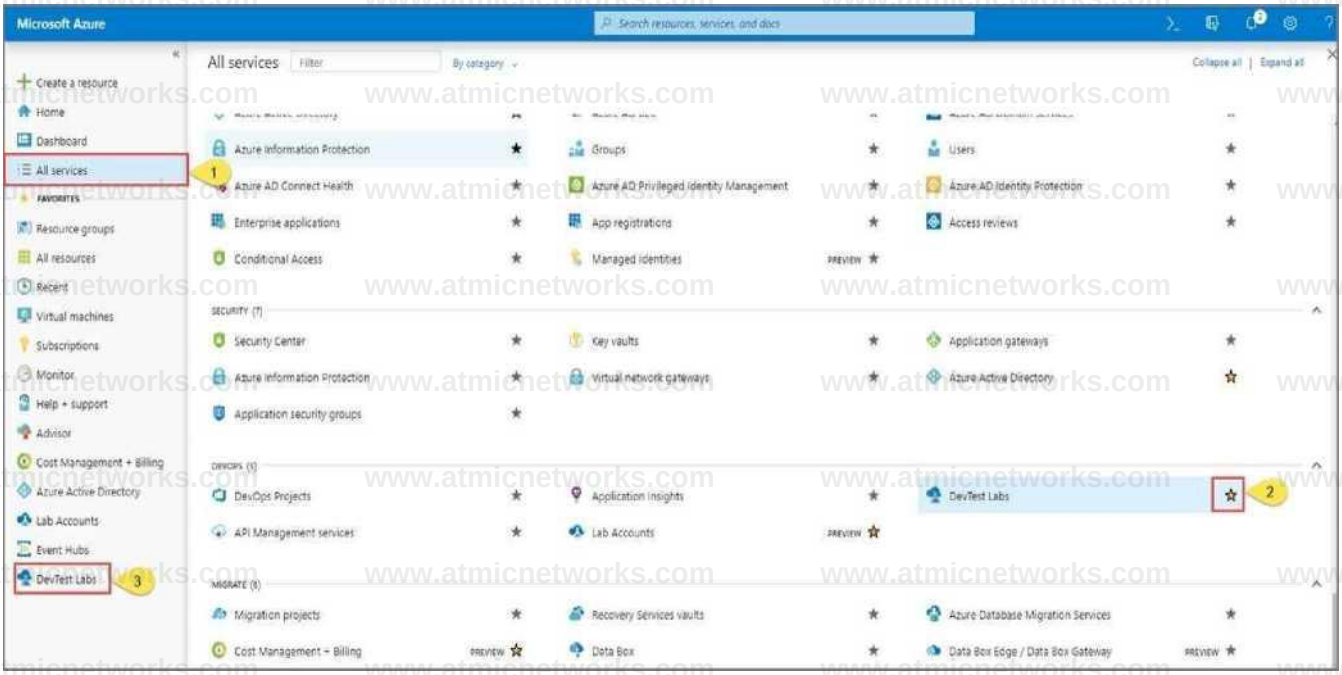

3. From the list of labs, select the az400-9940427-dtl1 lab
4. On the home page for your lab, select + Add on the toolbar.
5. Select the Windows Server 2016 Datacenter base image for the VM.
6. Select automation options at the bottom of the page above the Submit button.
7. You see the Azure Resource Manager template for creating thevirtual 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: 514

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.

Answer: See solution below.

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: Bydefault, 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: 515

SIMULATION

You plan to deploy a template named D:\Deploy.json to a resource group named Deploy-lod9940427.

You need to modify the template to meet the following requirements, and then to deploy the template:
The address space must be reduced to support only 256 total IP addresses.
The subnet address space must be reduced to support only 64 total IP addresses.

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

Answer: See solution below.

Explanation:

1. Sign in to the portal,
2. Choose template Deploy-lod9940427
3. Select Edit template, andthen paste your JSON template code into the code window.
4. Change the ASddressPrefixes to 10.0.0.0/24 in order to support only 256 total IP addresses.

addressSpace":{"addressPrefixes": ["10.0.0.0/24"]},

5. Change the firstSubnet addressprefix to 10.0.0.0/26 to support only 64 total IP addresses.

"subnets":[
{
"name":"firstSubnet",
"properties":{
"addressPrefix":"10.0.0.0/24"
}

6. Select Save.

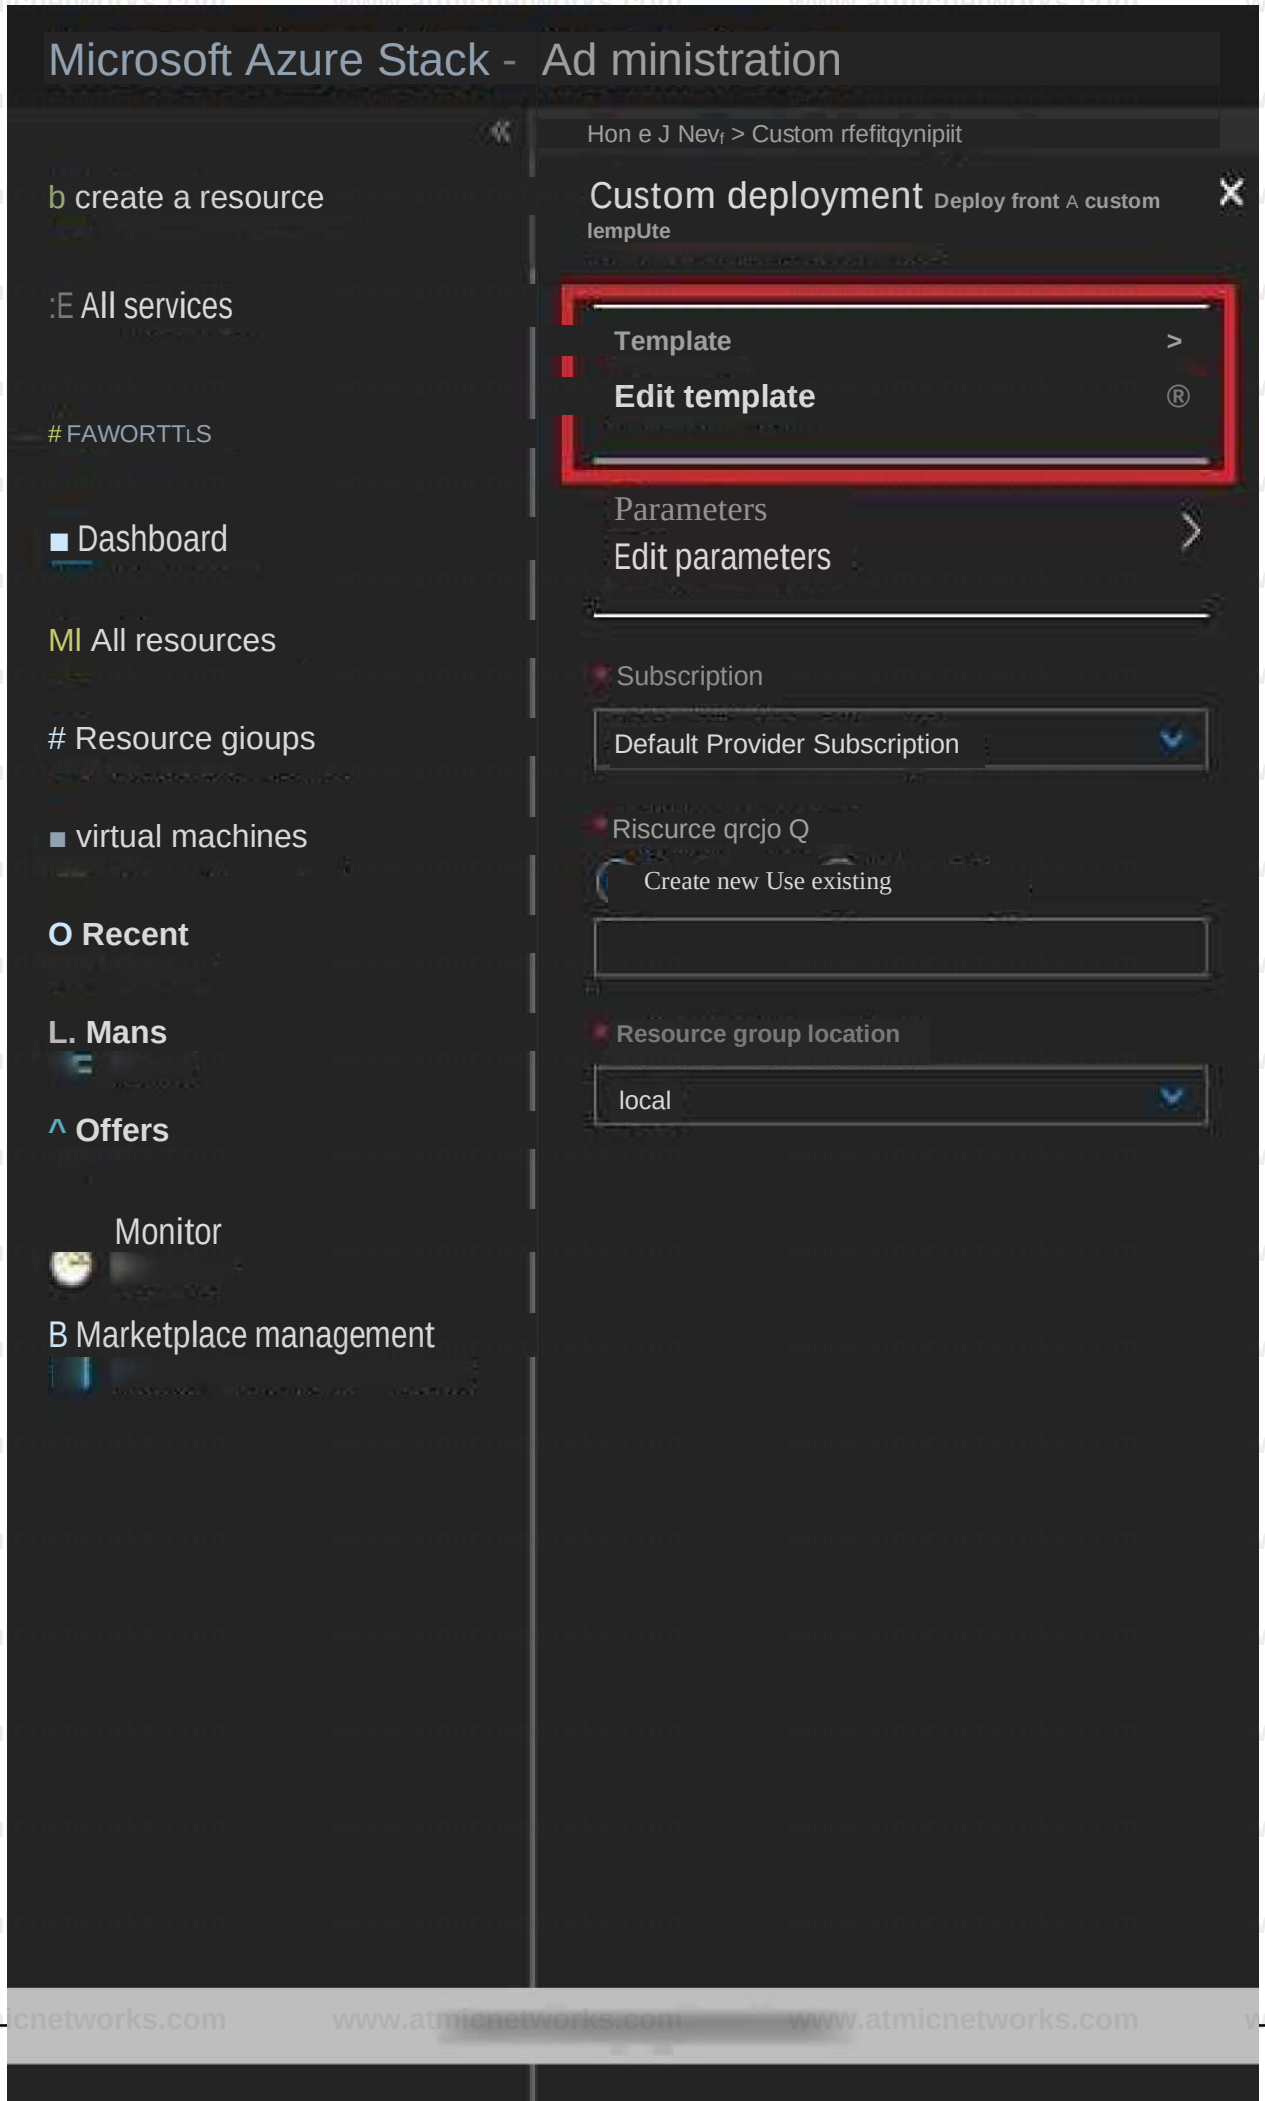
Microsoft Azure Stack - Ad ministration
Hon e J Nevf > Custom rfefitqynipiit
b create a resource
:E All services
FAWORTTLS
■ Dashboard
Ml All resources
Resource gioups
■ virtual machines
O Recent
L. Mans
^ Offers
Monitor
B Marketplace management
Custom deployment Deploy front A custom lempUte
Template >
Edit template ®
Parameters
Edit parameters
Subscription
Default Provider Subscription
Riscurce qrcjo Q
Create new Use existing
Resource group location
local

7. Select Editparameters, provide values for the parameters that are shown, and then select OK.
8. Select Subscription. Choose the subscription you want to use, and then select OK.
9. Select Resource group. Choose an existing resource group or create a new one, and then select OK.

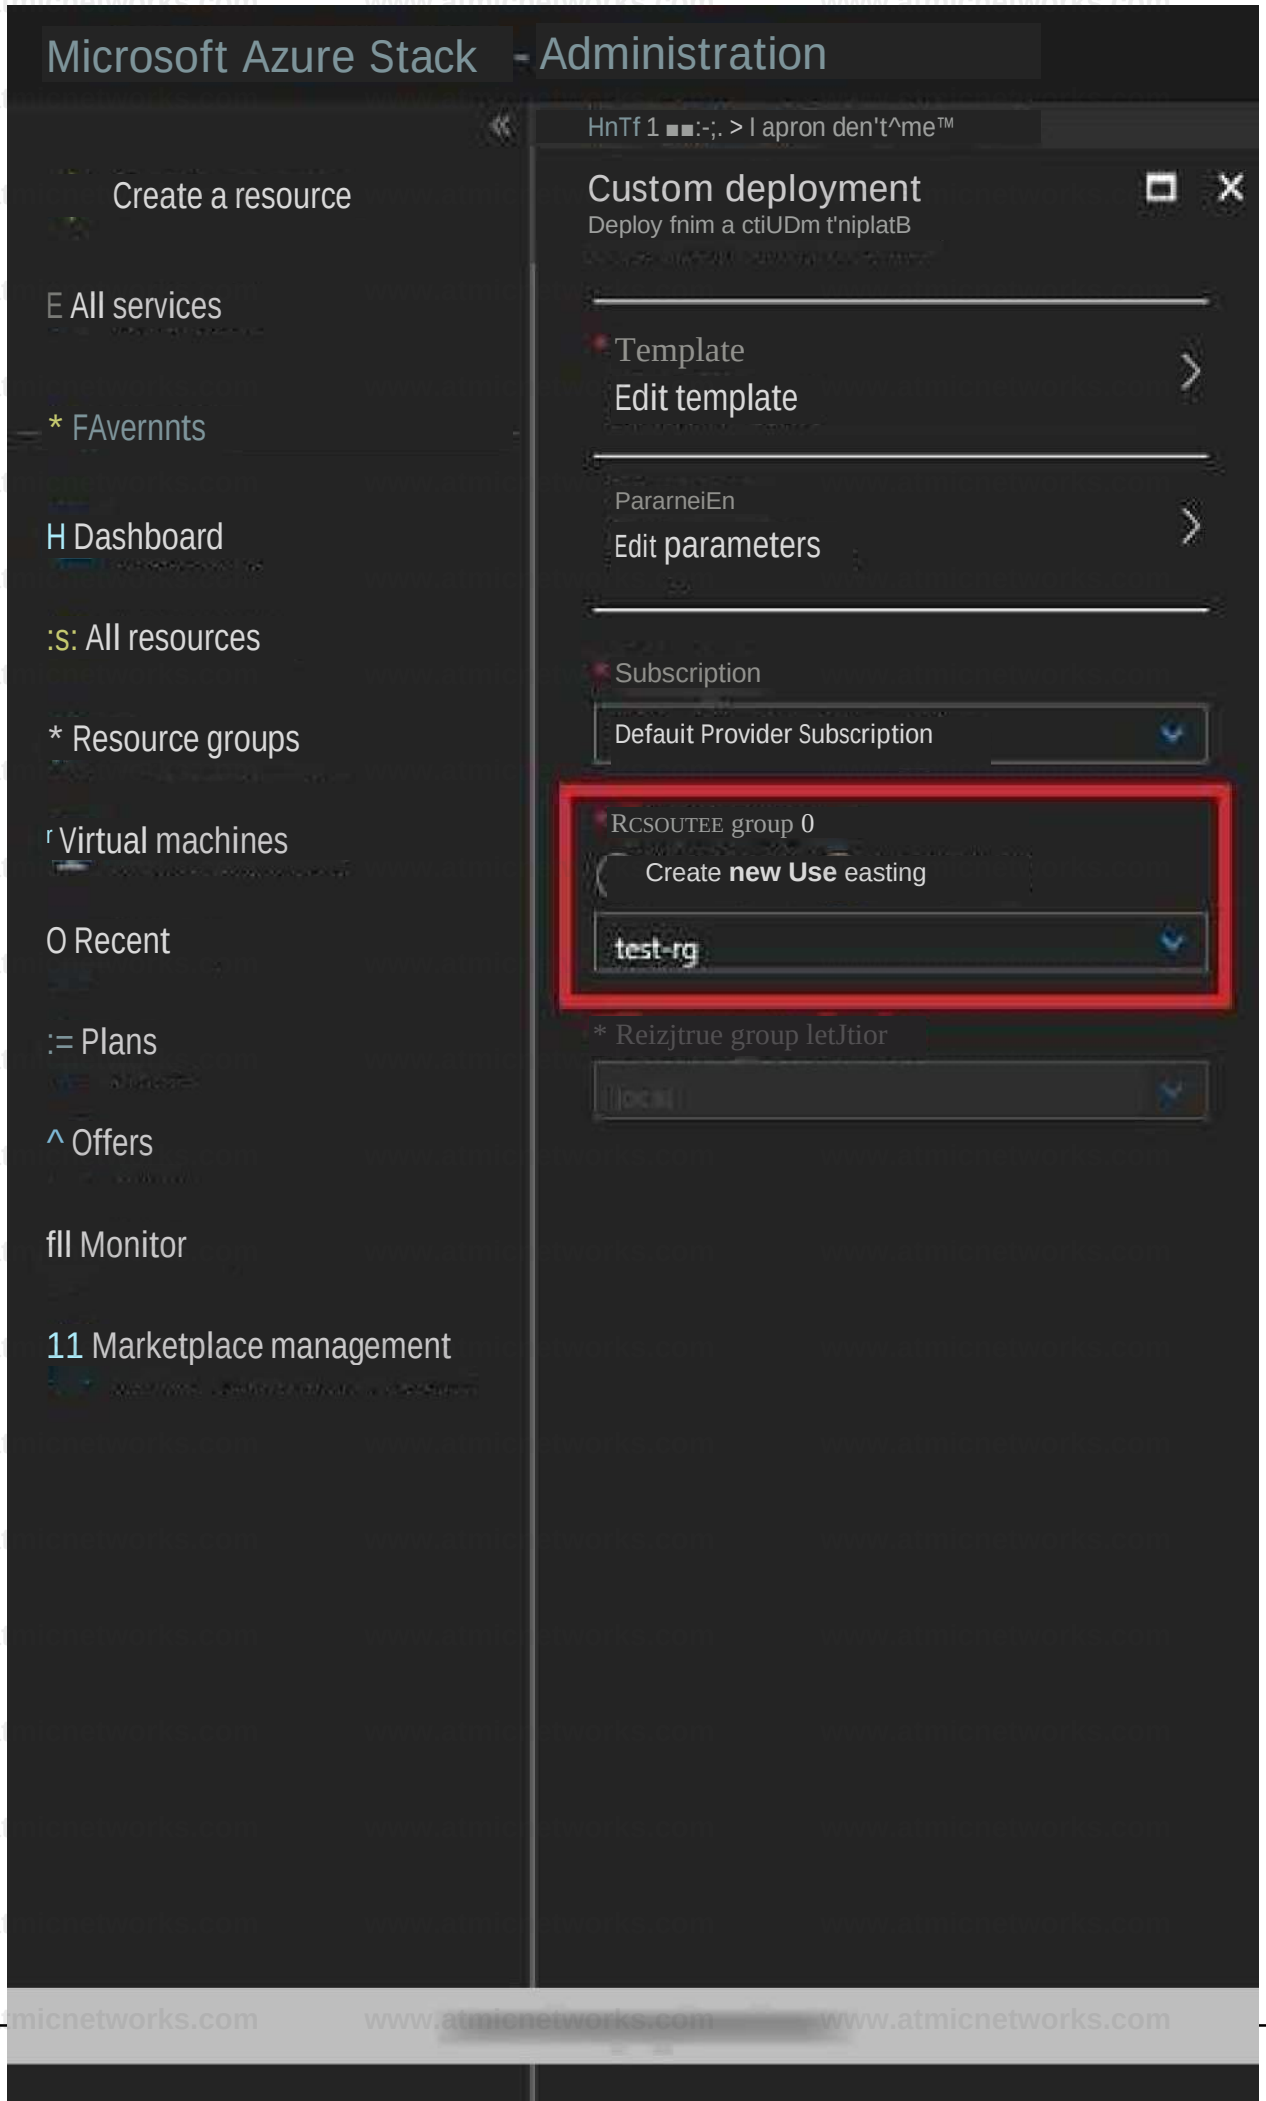
Microsoft Azure Stack - Administration
Create a resource
E All services
* FAvernnts
H Dashboard
:s: All resources
* Resource groups
r Virtual machines
O Recent
:= Plans
^ Offers
fll Monitor
11 Marketplace management
HnTf 1 ■■:-;. > I apron den't^me™
Custom deployment
Deploy fnim a ctiUDm t'niplatB
Template
Edit template
PararneiEn
Edit parameters
Subscription
Defauit Provider Subscription
RCSOUTEE group 0
Create new Use easting
test-rg
* Reizjtrue group letJtior

10. Select Create. A new tile on the dashboard tracks the progress of your template deployment.

Reference:

https://docs.microsoft.com/en-us/azure-stack/user/azure-stack-deploy-template-portal?view=azs-1908

https://docs.microsoft.com/en-us/azure/architecture/building-blocks/extending-templates/update-resource

Question: 516

SIMULATION

You need to configure an Azure web app named az400-9940427-main to contain an environmental variable named "MAX_ITEMS". The environmental variable must have a value of 50.

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

Answer: See solution below.

Explanation:

anation:

1. In the Azure portal, navigate to the az400-9940427-main app's management page. In the app's left menu, click Configuration > Application settings.

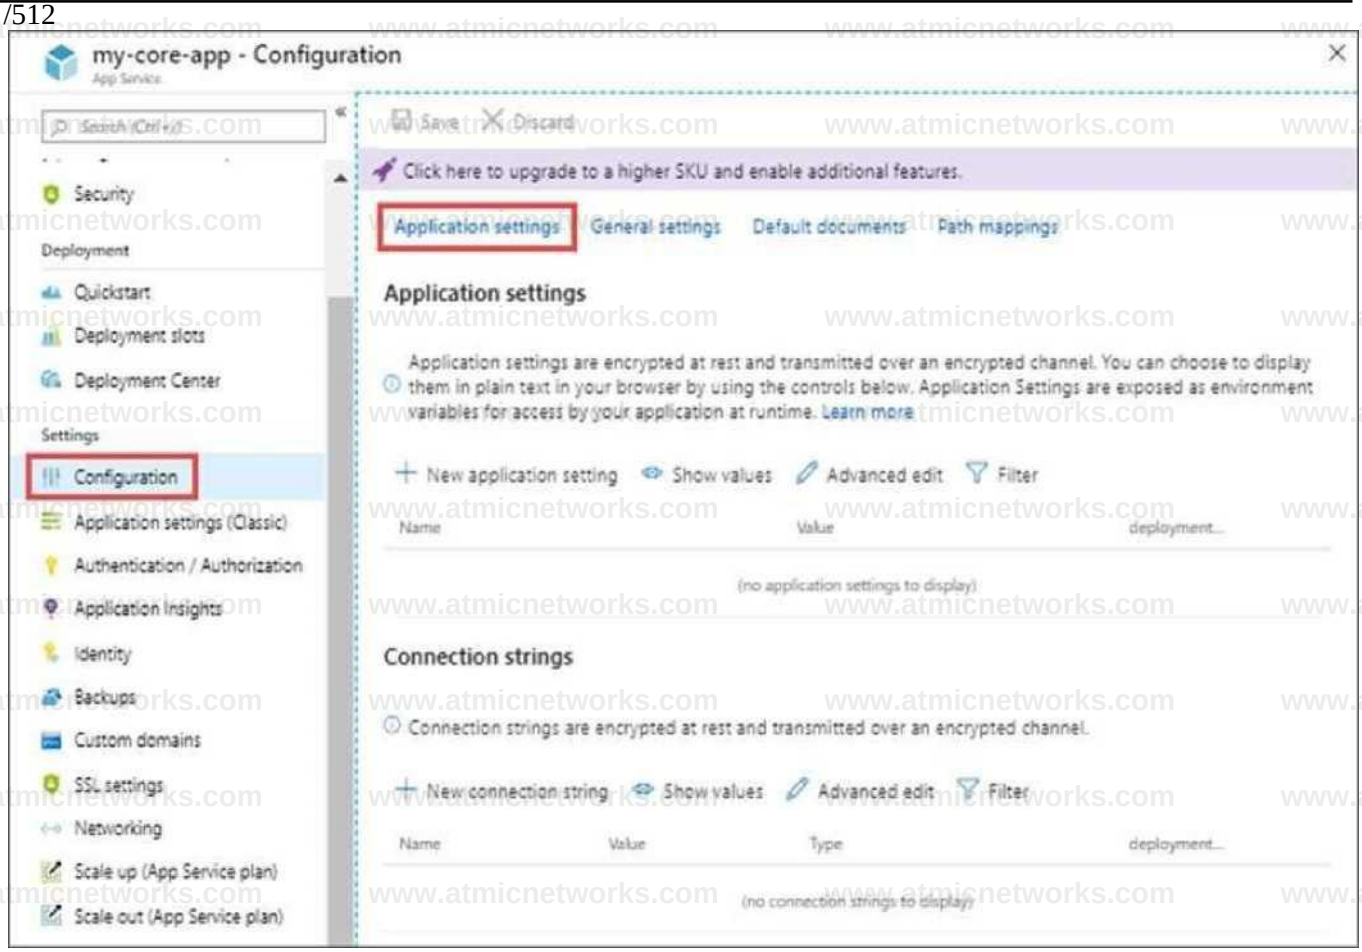

2. Click New Application settings

3. Enter the following: Name: MAX_ITEMS Value: 50

Reference:
https://docs.microsoft.com/en-us/azure/app-service/configure-common

Question: 517

SIMULATION

You need to create a notification if the peak average response time of an Azure web app named az400-9940427-main is more than five seconds when evaluated during a five-minute period. The notification must trigger the "https://contoso.com/notify" webhook.

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

Answer: See solution below.

Explanation:

1. Open Microsoft Azure Portal

2. Log intoyour Azure account and go to App Service and look under Monitoring then you will see Alert.

3. Select Add an alert rule

4. Configure the alert rule as per below and click Ok.

Source: Alert on Metrics
Resource Group: az400-9940427-main
Resource: az400-9940427-main
Threshold: 5
Period: Over the last 5 minutes
Webhook: https://contoso.com/notify

Reference:
https://azure.microsoft.com/es-es/blog/webhooks-for-azure-alerts/

Question: 518

SIMULATION

You need to create an instance of Azure Application Insights named az400-9940427-main and configure the instance to receive telemetry data from an Azure web app named az400-9940427- main.

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

Answer: See solution below.

Explanation:

Step 1: Create an instance of Azure Application Insights

1. Open Microsoft Azure Portal
2. Log into your Azure account, Select Create a resource > Developer tools > Application Insights.

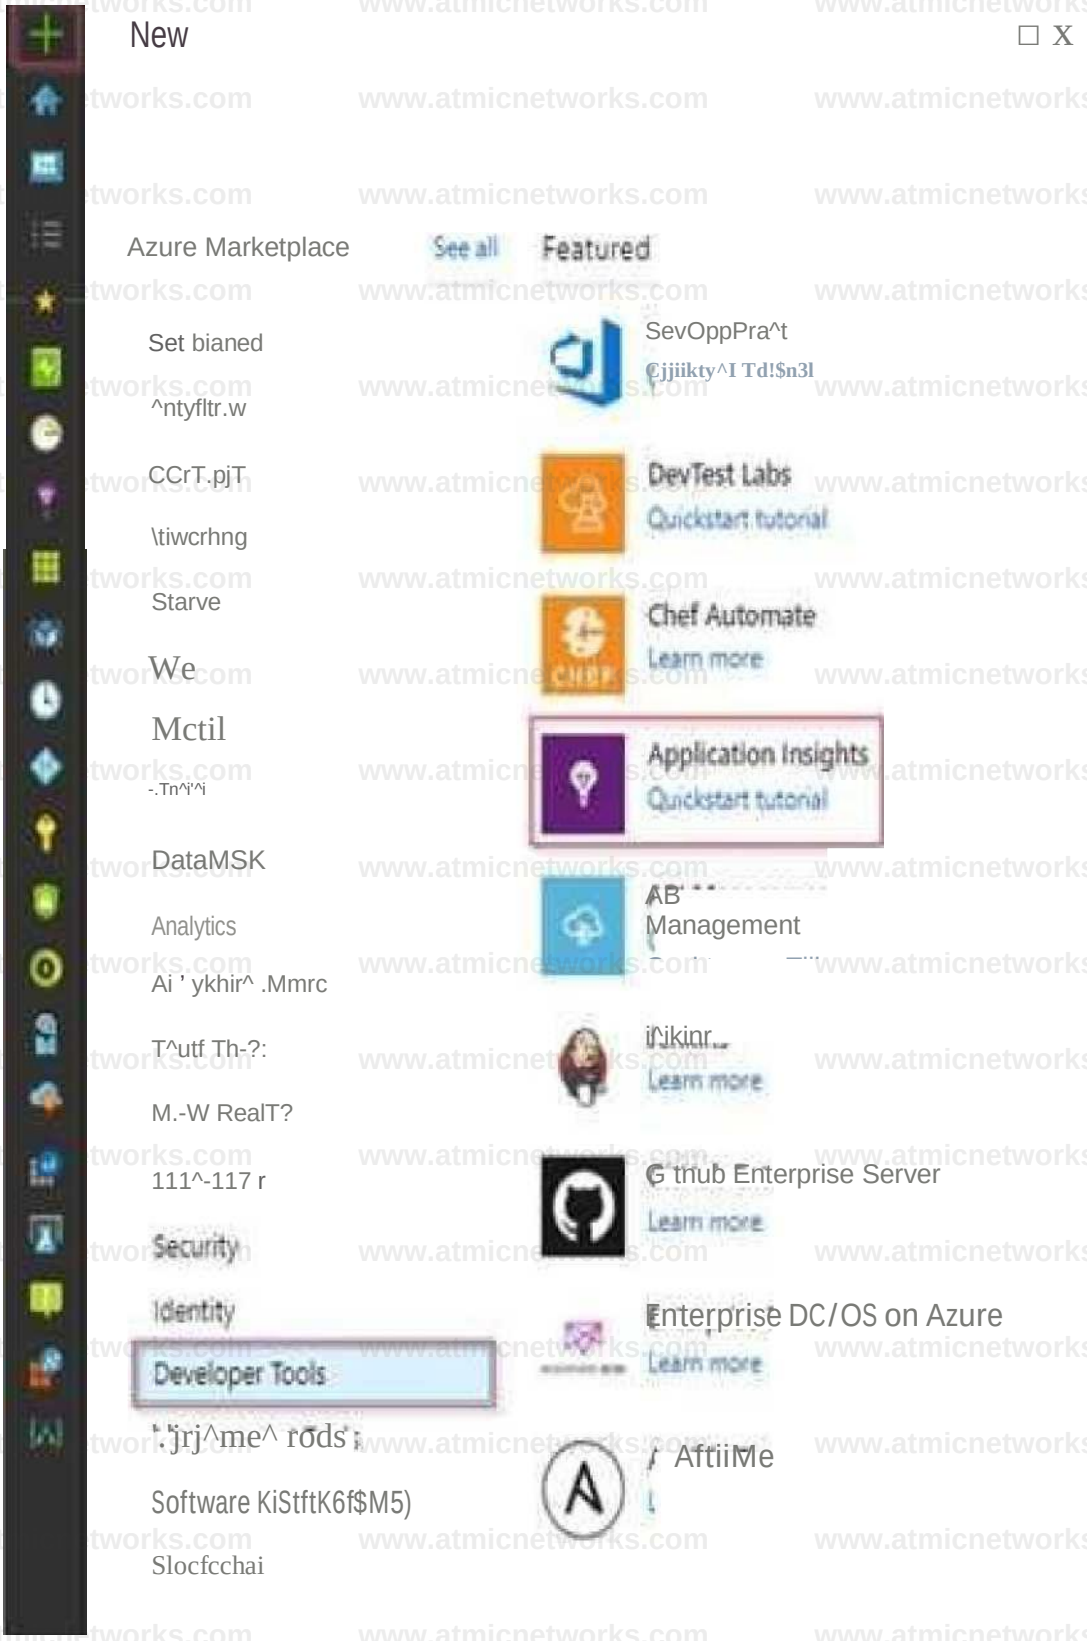

3. Enter the following settings, and then select Review + create. Name: az400-9940427-main

Step 2: Configure App Insights SDK

4. Open your ASP.NET Core Web App project in Visual Studio > Right-click on the AppName in theSolution Explorer > Select Add > Application Insights Telemetry.

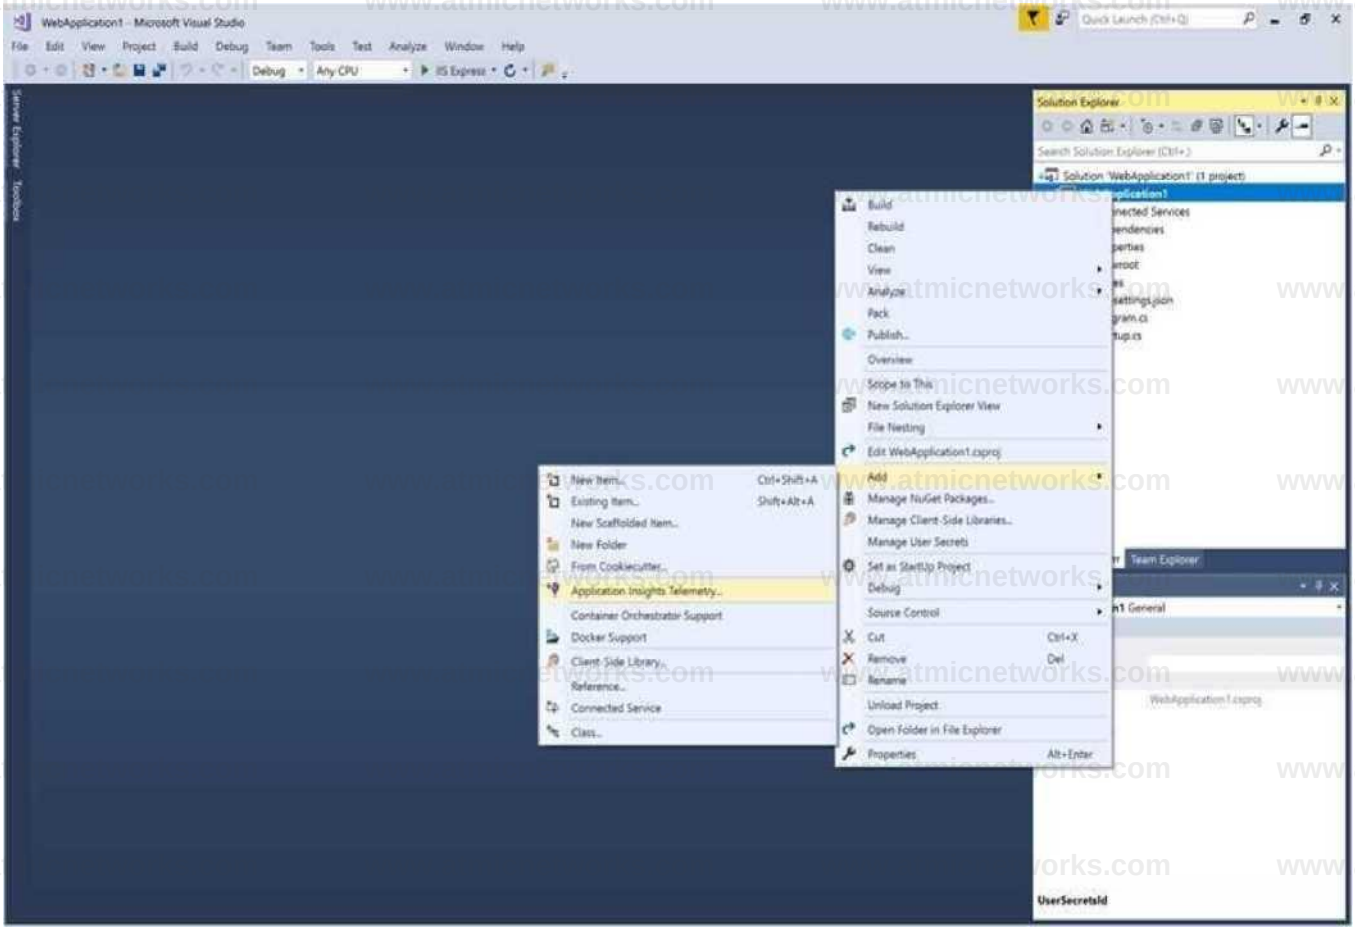

5. Click the Get Started button

6. Select your account and subscription > Select the Existing resource you created in the Azure portal > Click Register.

Reference:
https://docs.microsoft.com/bs-latn-ba/azure/azure-monitor/learn/dotnetcore-quick-start?view=vs- 2017

Question: 519

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.

Answer: See solution below.

Explanation:

Sign in to the Azure portal.

Navigate to the container registry az40010480345acr1.

Under Services, select Webhooks.

Select the existing webhook https://contoso.com/statushook, and double-click on itto get its properties.

For Trigger actions select image push

Example web hook:

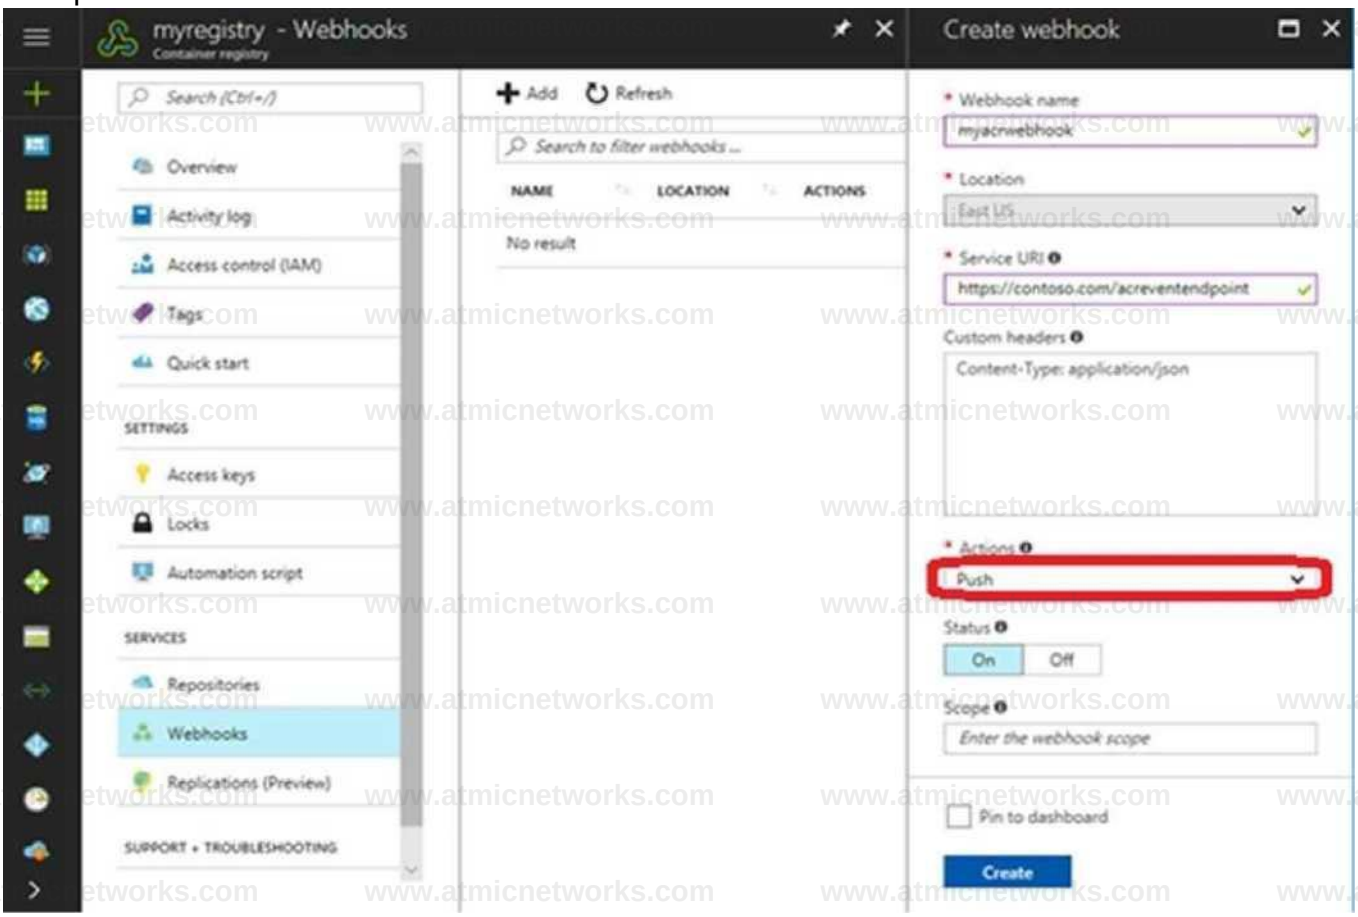

Reference:
https://docs.microsoft.com/en-us/azure/container-registry/container-registry-webhook

Question: 520

SIMULATION

You need to ensure that Microsoft Visual Studio 2017 can remotelyattach to an Azure Function named fa-11566895.

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

Answer: See solution below.

Explanation:

Enable Remote Debugging
Before we start a debugging session to our Azure Function app we need to enable the functionality.

Navigate in the Azure portal to your function app fa-11566895
Go to the “Application settings”
Under “Debugging” set Remote Debugging to On and set Remote Visual Studio version to 2017.

Reference:
https://www.locktar.nl/uncategorized/azure-remote-debugging-manually-in-visual-studio-2017/

Question: 521
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.

Answer: See solution below.

Explanation:

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

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

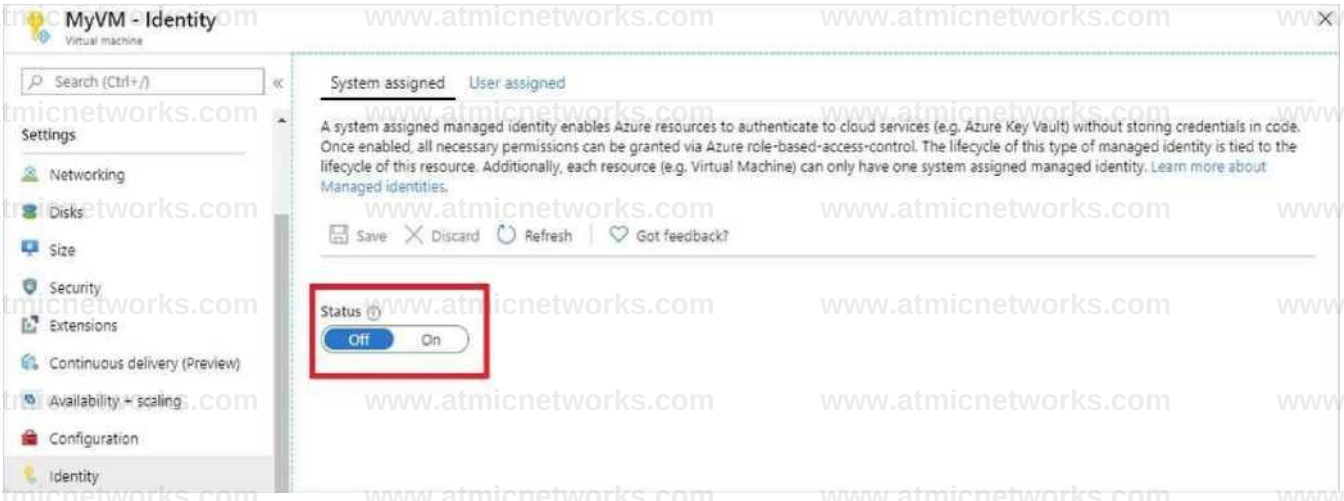

Note:Enabling a system-assigned managed identity is a one-click experience. You can either enable it during the creation of a VM or in the properties of an existing VM.

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

Question: 522

SIMULATION

O: 43

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 accountin your resource group.

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

Answer: See solution below.

Explanation:

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

Open the App Serviceaz400-11566895-main, which you want to protect, in the Azure Portal and

browse to Settings > Backups. Click Configure and a Backup Configuration blade should appear.

Select the storage account.

Click + to create a private container. You could name this container after the web app or App Service. Select the container.

If you want to schedule backups, then set Scheduled Backup to On and configure a schedule: every five hours

Select your retention. Note that 0 means never delete backups.

Decide if at least onebackup should always be retained.

Choose if any connected databases should be included in the web app backup.

Click Save to finalize the backup configuration.

Reference:

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

Question: 523

SIMULATION

You plan to deploy a runbook that will create Azure AD user accounts.

You need to ensure that runbooks can run the AzurePowerShell cmdlets for Azure Active Directory.

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

Answer: See solution below.

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.

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 yourorganization." 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. Multifactor 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: ConfigureAzure 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.

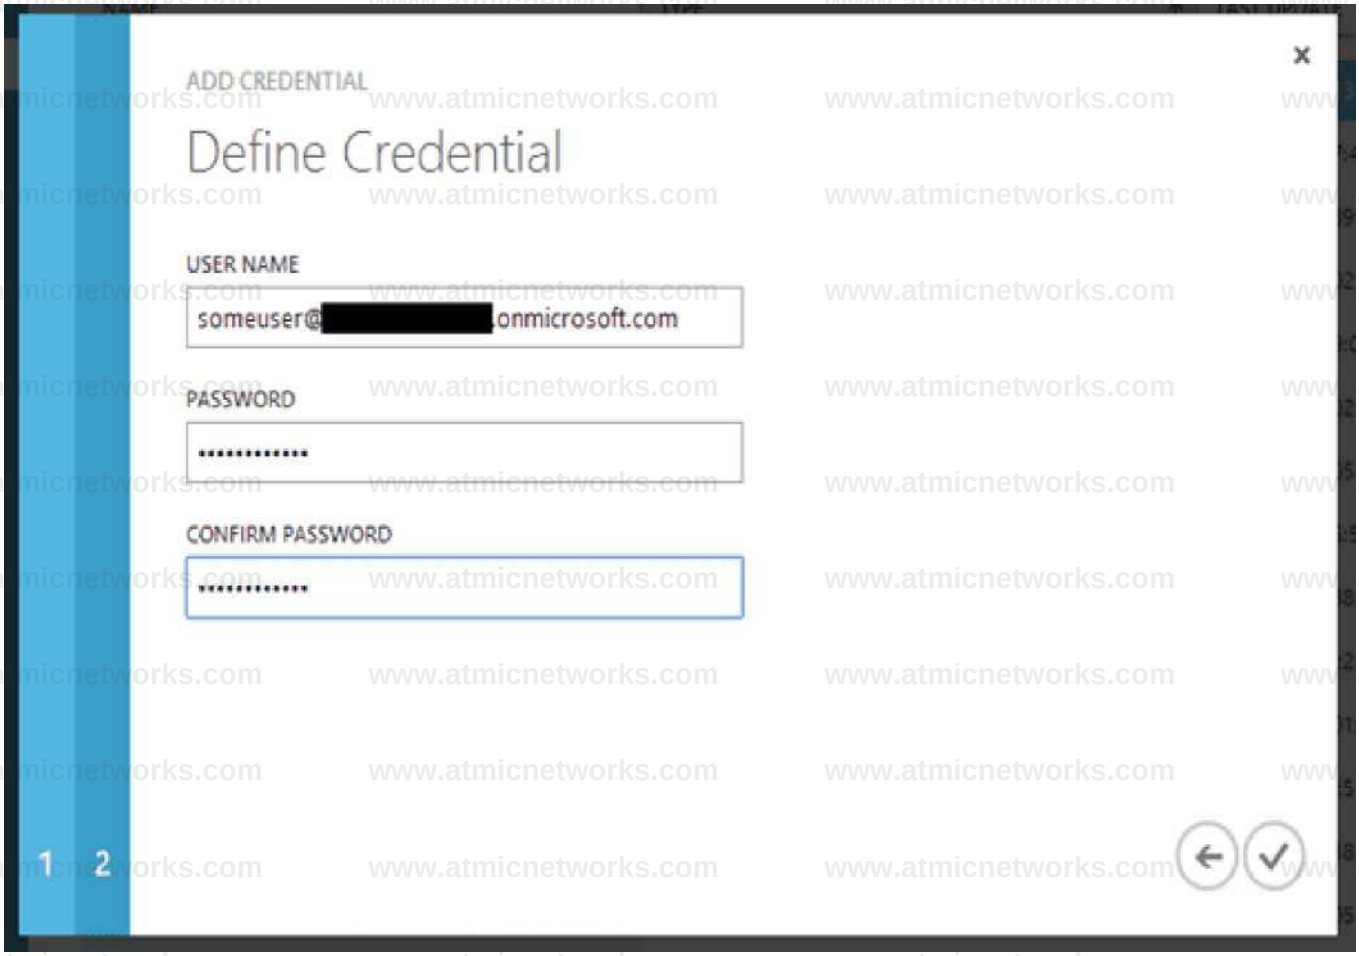

Note: Once you have set up the Azure Active Directory credential in Azure and Azure Automation, you can now manage Azure from Azure Automation runbooks using this credential.

Reference:
https://azure.microsoft.com/sv-se/blog/azure-automation-authenticating-to-azure-using-azure- active-directory/

Question: 524

SIMULATION

You manage a website that uses an Azure SQL Database named db1 in a resource group named RG1lod11566895.

You need tomodify the SQL database to protect against SQL injection.

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

Answer: See solution below.

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:

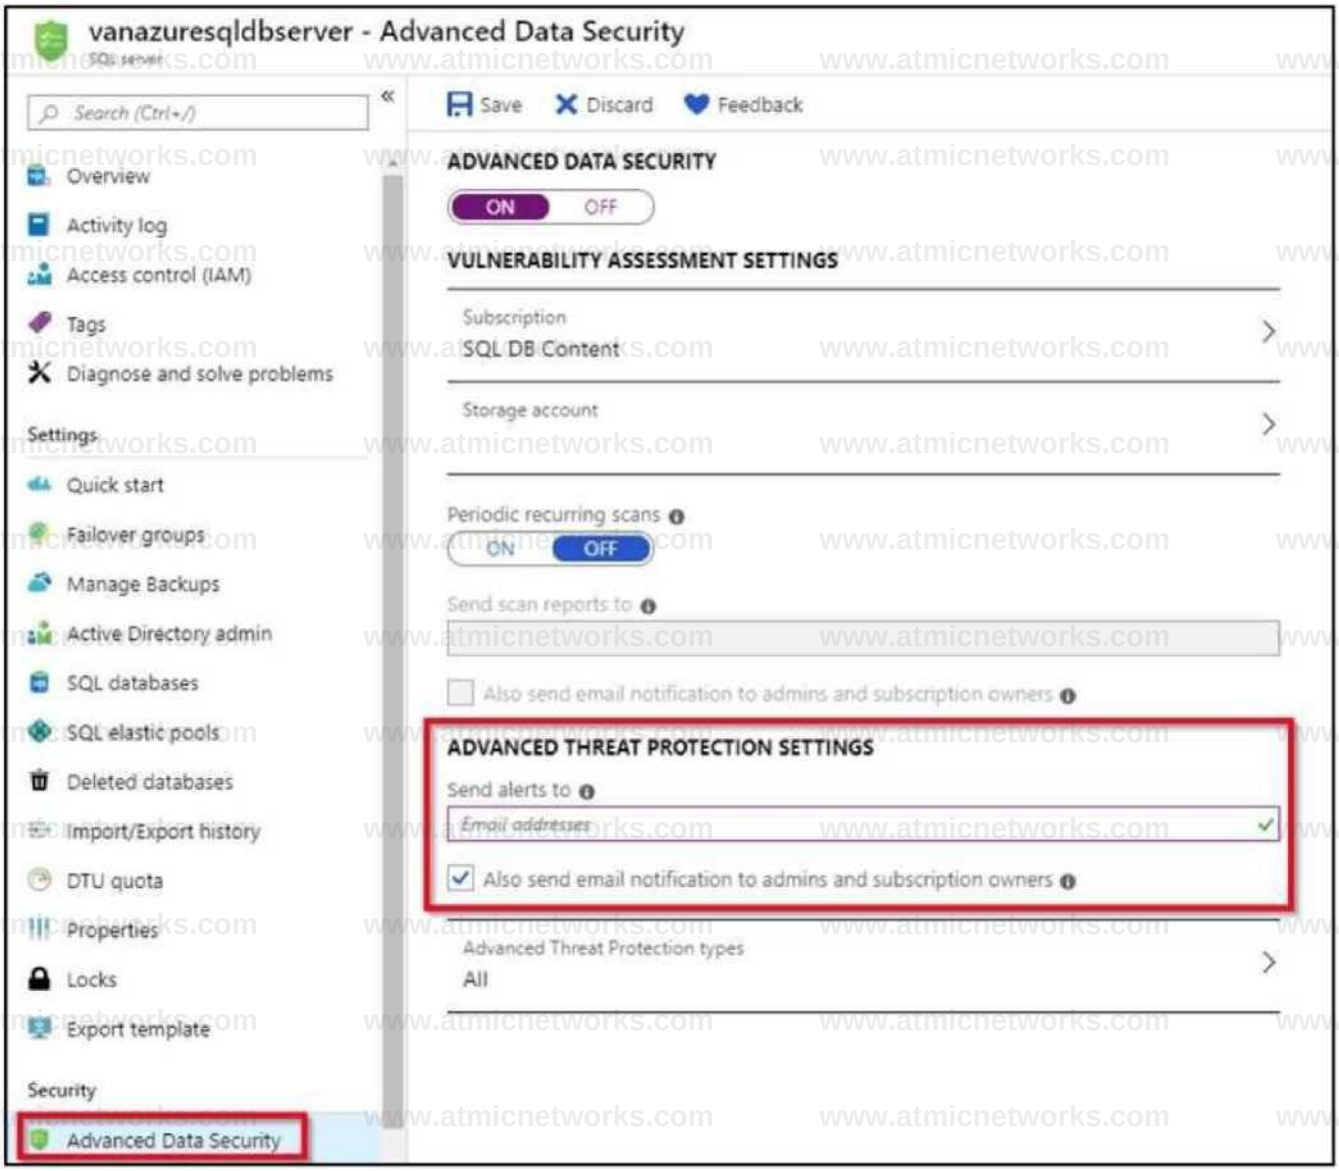

4. Enable Advanced Data Security on the server.

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: 525

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 foraz400-11566895-main.

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

Answer: See solution below.

Explanation:

Add a slot

1. In the Azure portal, search for and select App Services and select your app az400-11566895-main.

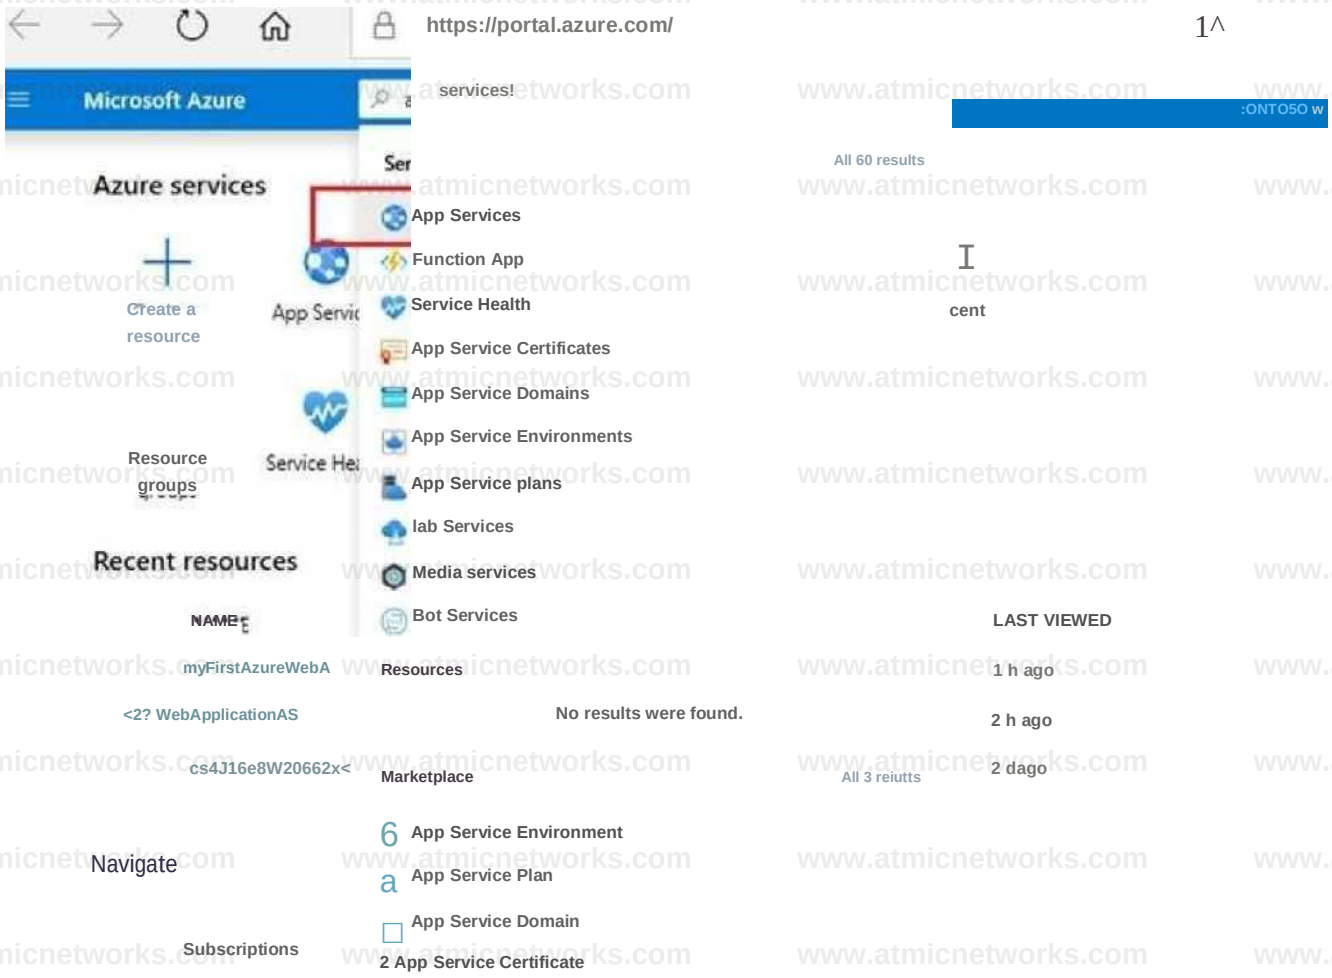

2. In the left pane, select Deployment slots > Add Slot.

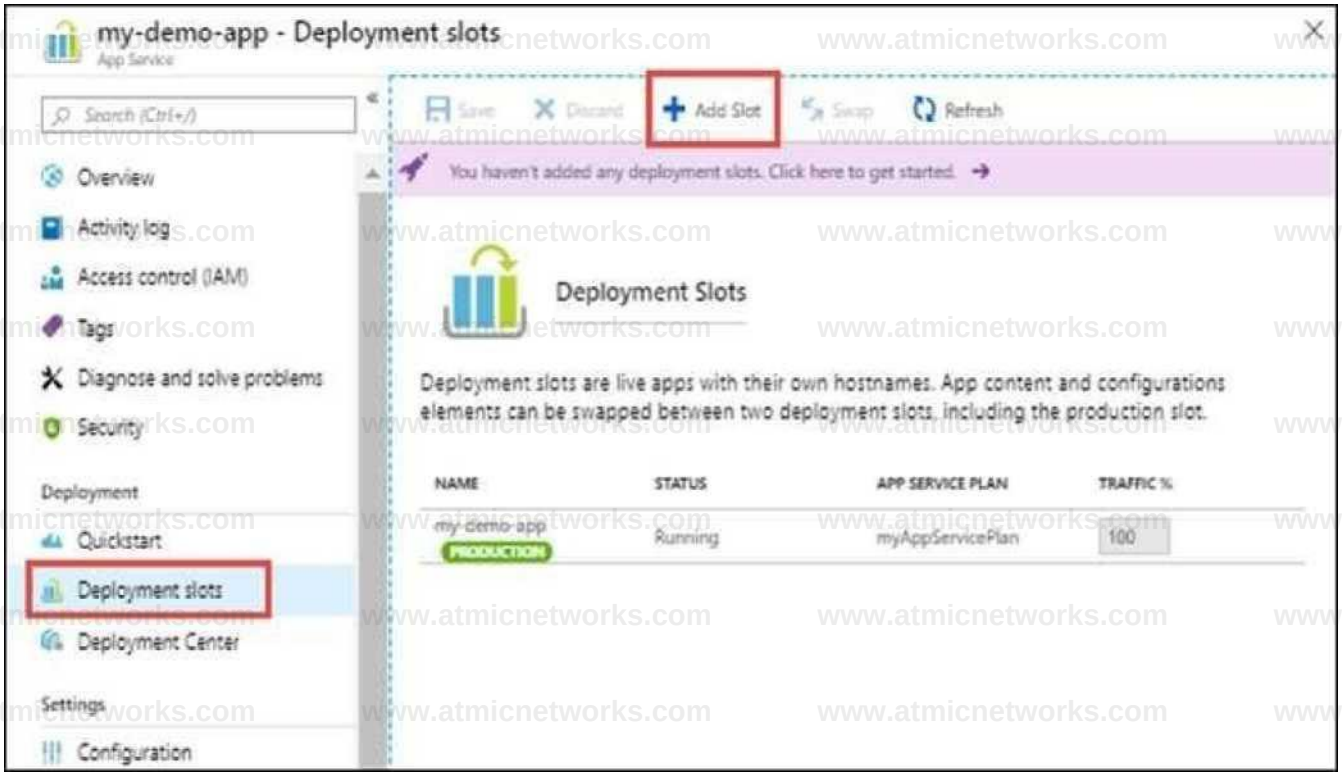

3. In the Add a slot dialog box, give the slot a name, and select whether to clone an app configuration from another deployment slot. Select Add to continue.

4. After the slot is added, select Close to close the dialog box. The new slot is now shown on the Deployment slots page.

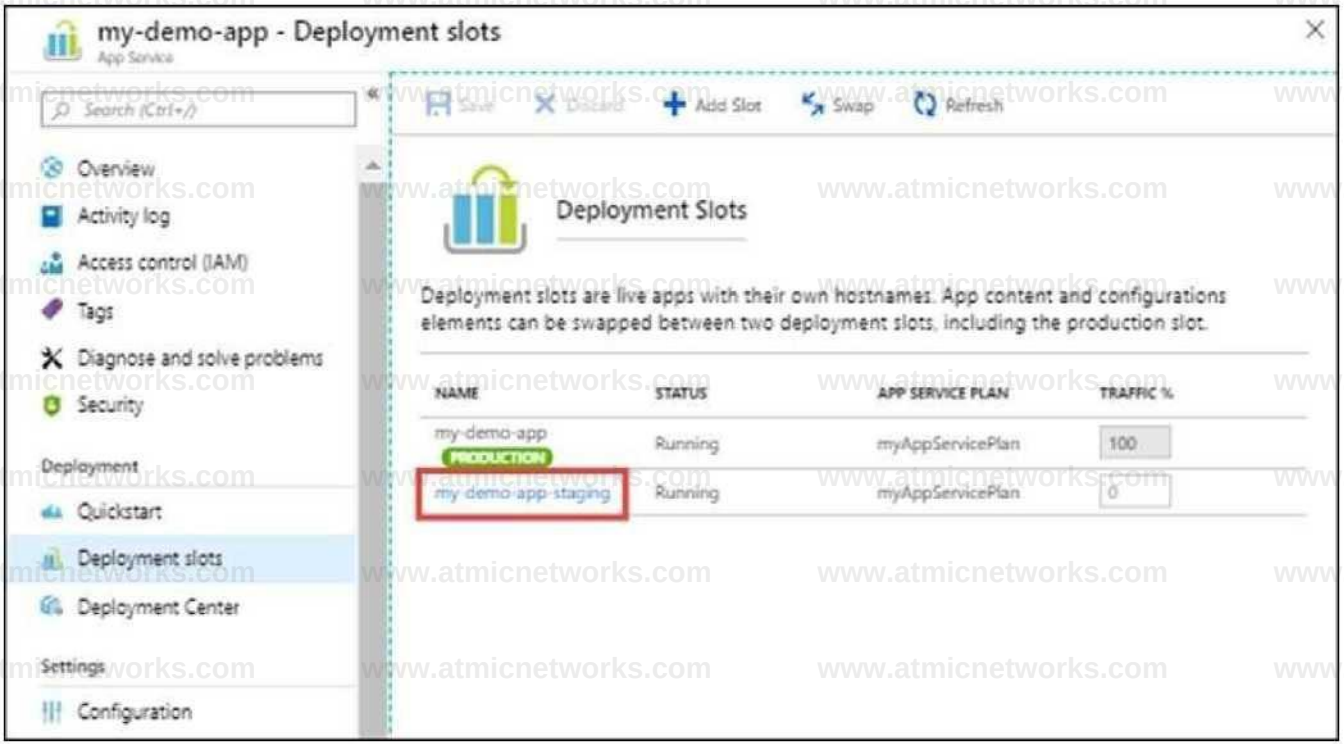

Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots

Question: 526

SIMULATION

You have several apps that use an Azure SQLDatabase 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.

Answer: See solution below.

Explanation:

1. To enable automatictuning on a single database, navigate to the database in the Azure portal and select Automatic tuning.

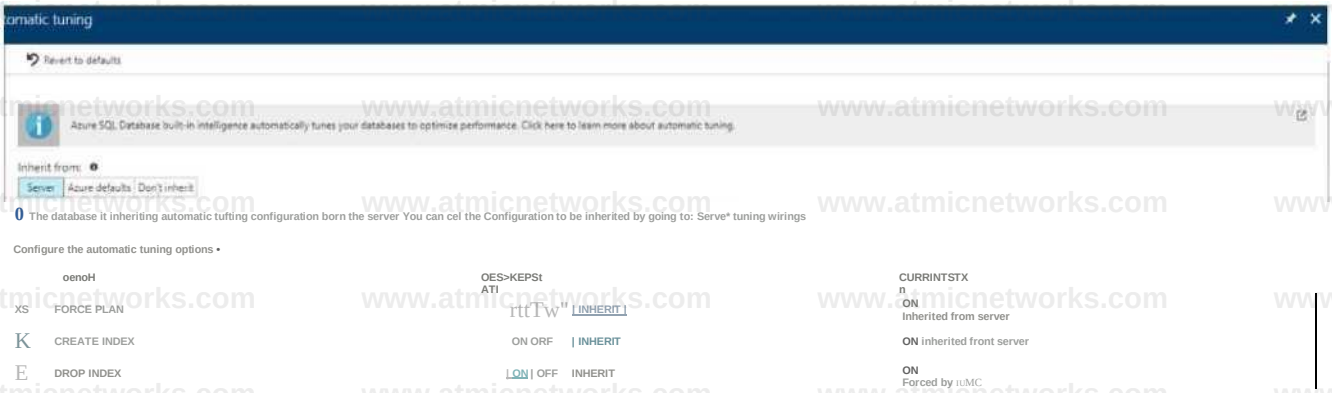

2. Select the automatic tuning options you want to enable and select Apply.

Note: Individual automatic tuning settings can be separately configuredfor 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: 527

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.

Answer: See solution below.

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

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: 528

SIMULATION

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

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

Answer: See solution below.

Explanation:

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

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

Step 2: Enable boot diagnostics on existing virtual machine
To enable Boot diagnostics on an existing virtual machine, follow these steps:

1. Sign in to the Azure portal, and then select the virtual machine VM1.
2. In the Support + troubleshooting section, select Boot diagnostics, then select the Settings tab.

3. In Boot diagnostics settings, change the status to On, and from the Storage account drop-down list, select the storage account az400lod11566895stor.
4. Save the change.

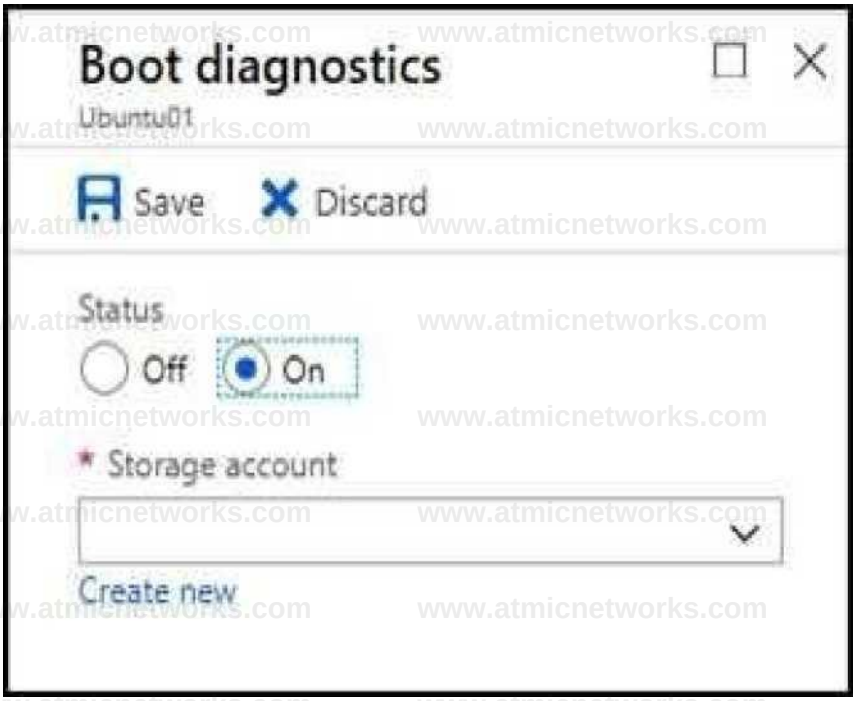

You must restart the virtual machine for the change to take effect.

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

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

Question: 529

SIMULATION

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

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

To complete this task, signin to the Microsoft Azure portal.

Answer: See solution below.

Explanation:

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

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

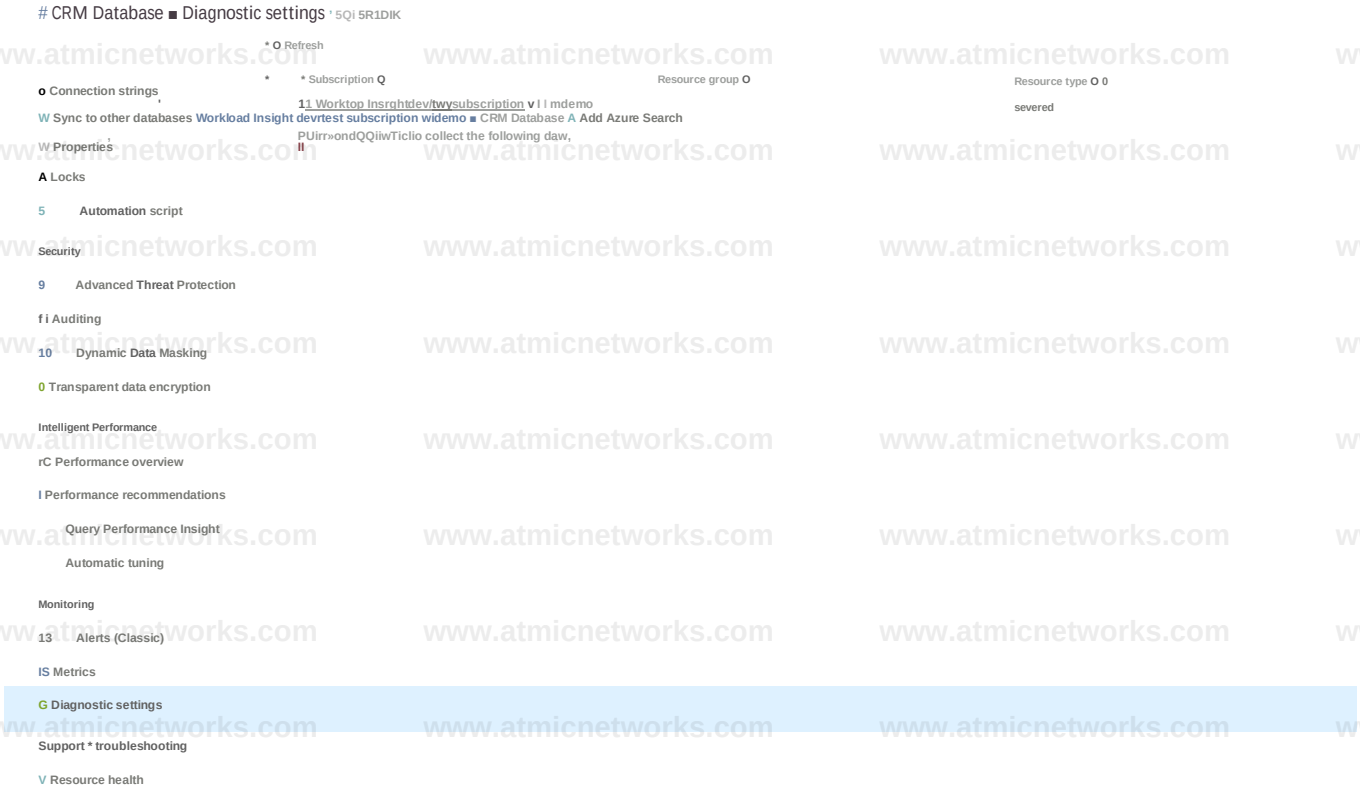

5. Enter a setting name for your own reference.
6. Select a destination resource for the streaming diagnostics data: Archive to storage account, Stream to an event hub, or Send to Log Analytics.
7. For the standard, event-based monitoring experience, select the following check boxes for database diagnostics log telemetry: QueryStoreRuntimeStatistics

Diagnostics settings X

R Save X Discard C - /_

* Name

service

Archive to a storage account

J Stream to an event hub

LOG

- SQLInsights
- AutomaticTuning
- QueryStoreRuntimeStatistics
- QueryStoreWaitStatistics
- Errors
- DatabaseWaitStatistics
- Timeouts
- Blocks
- Deadlocks

METRIC

0 Send to Log Analytics

Subscription

Workload insight dev/test subscription	
Log Analytics Workspace	
sqlanalytics356 (westcentralus)	

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

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

Question: 530

SIMULATION

You need to create an instance of Azure Application Insights named az400-9940427-main and configure the instance to receive telemetry data from an Azure web app named az400-9940427- main.

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

Answer: See explanation below

Explanation:

Step 1: Create an instance of Azure Application Insights

1. Open Microsoft Azure Portal
2. Log into your Azure account, Select Create a resource > Developer tools > Application Insights.
3. Enter the followingsettings, and then select Review + create.

Name: az400-9940427-main

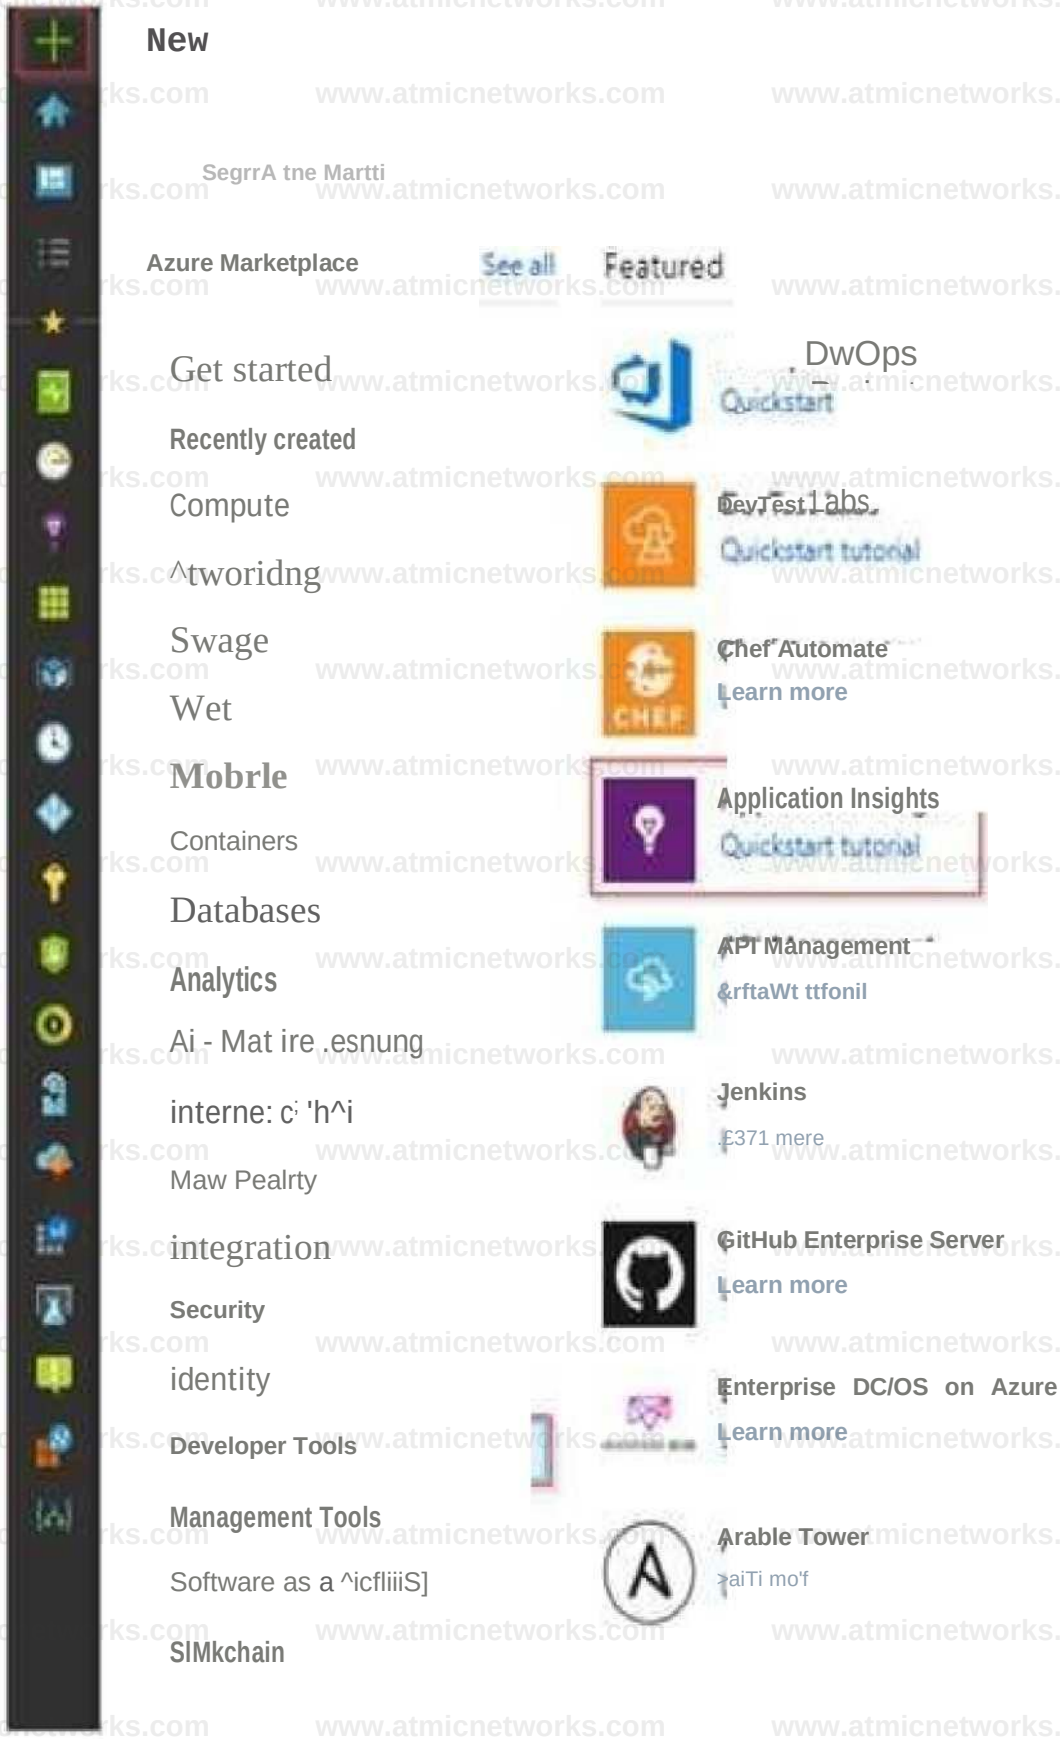

Step 2: Configure App Insights SDK

1. Open your ASP.NET Core Web App project in Visual Studio > Right-click on the AppName in the Solution Explorer > Select Add > Application InsightsTelemetry.

2. Click the Get Started button
3. Select your account and subscription > Select the Existing resource you created in the Azure portal > Click Register.

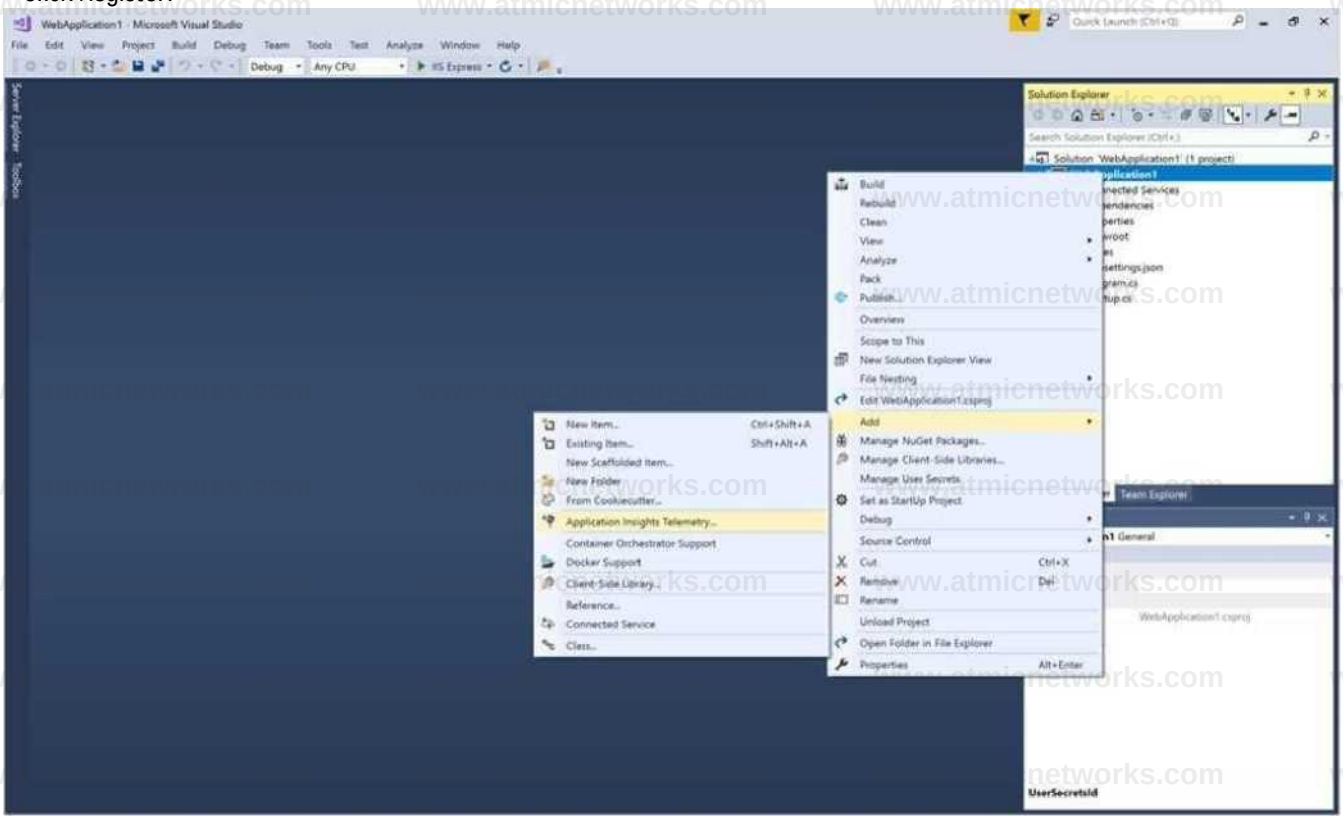

Reference:
https://docs.microsoft.com/bs-latn-ba/azure/azure-monitor/learn/dotnetcore-quick-start?view=vs- 2017

Question: 531

SIMULATION

You need to create a notification if the peak averageresponse time of an Azure web app named az400-9940427-main is more than five seconds when evaluated during a five-minute period. The notification must trigger the "https://contoso.com/notify" webhook.

To complete this task, sign in to the Microsoft Azureportal.

Answer: See explanation below

Explanation:

1. Open Microsoft Azure Portal
2. Log into your Azure account and go to App Service and look under Monitoring then you will see Alert.
3. Select Add an alert rule
4. Configure the alert rule as per belowand click Ok.

Source: Alert on Metrics
Resource Group: az400-9940427-main
Resource: az400-9940427-main

Threshold: 5
Period: Over the last 5 minutes
Webhook: https://contoso.com/notify

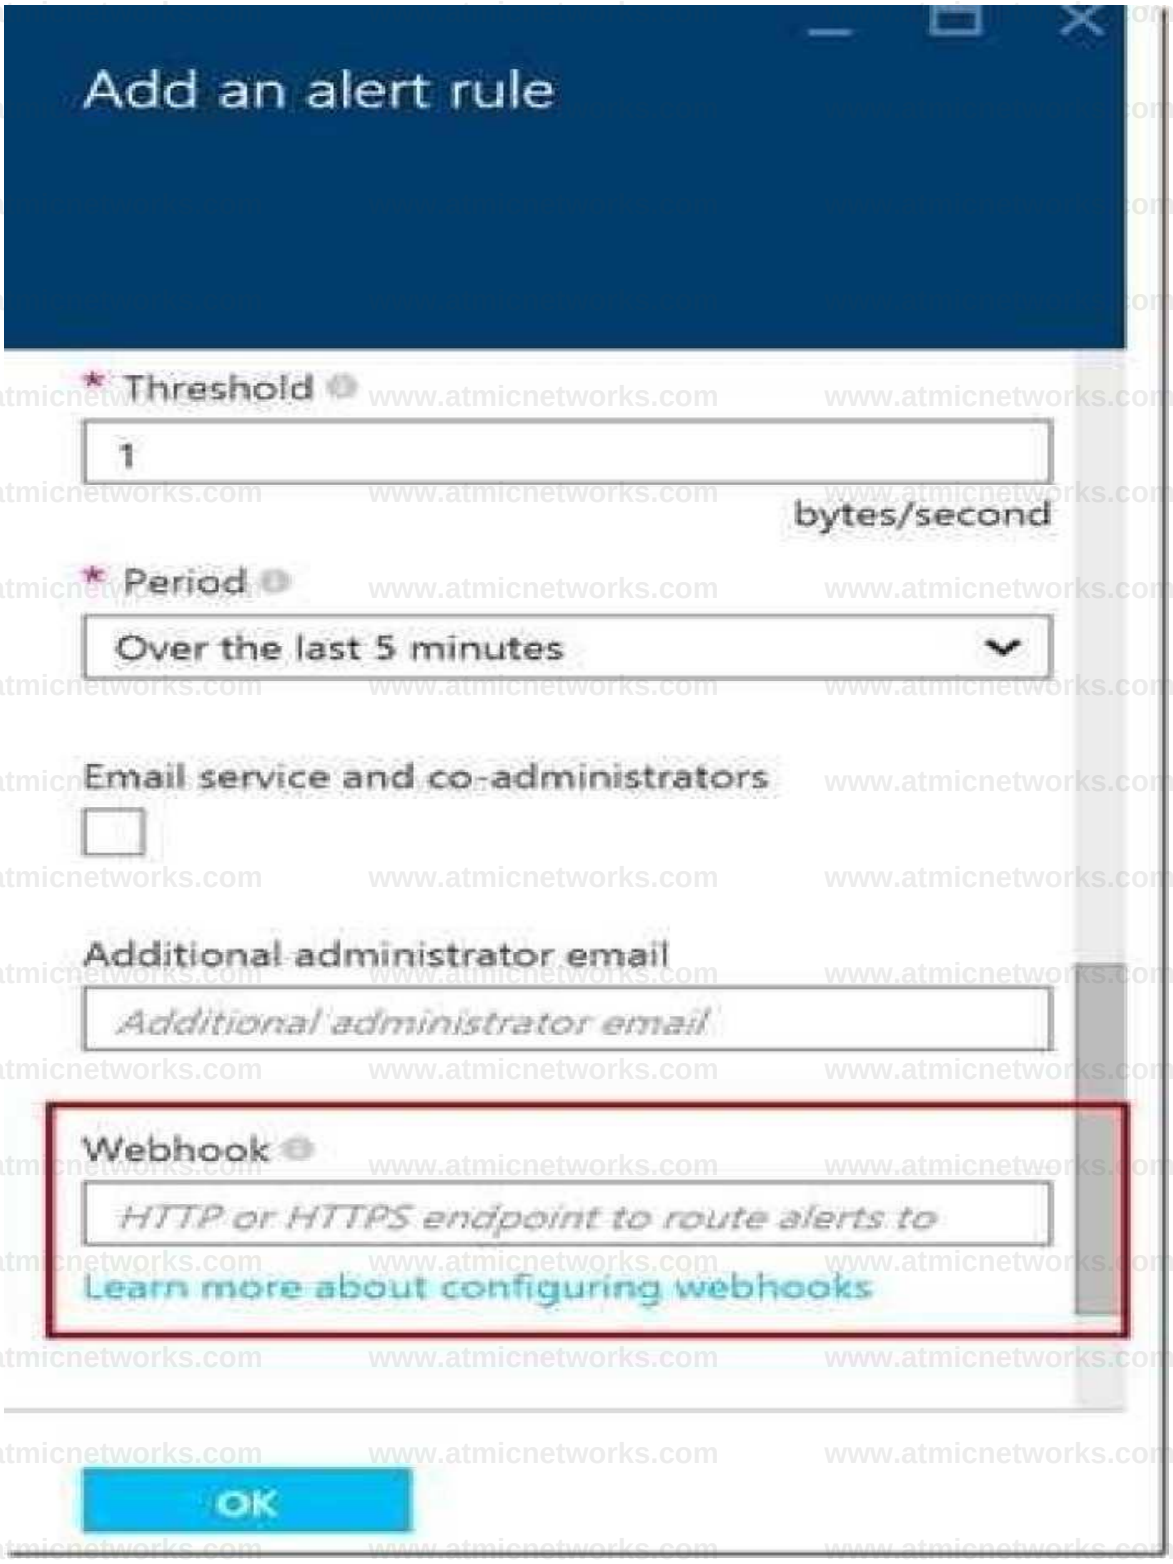

Reference:
https://azure.microsoft.com/es-es/blog/webhooks-for-azure-alerts/

Question: 532

SIMULATION

SIMULATION

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

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

Answer: See explanation below

Explanation:

Step 1: To create a general-purpose v2 storage account in the Azure portal, follow these steps: On the Azure portal menu, select All services. In the list of resources, type Storage Accounts. As you begin typing, the list filters based on your input. Select Storage Accounts.

On the Storage Accounts window that appears, choose Add.

Select the subscription in which to create thestorage account.

Under the Resource group field, select RG1lod11566895

Next, enter a name for your storage account named: az400lod11566895stor

Select Create.

Step 2: Enable boot diagnostics on existing virtual machine

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

Sign in to the Azure portal, and then select the virtual machine VM1.

In the Support + troubleshooting section, select Boot diagnostics, then select the Settings tab.

In Boot diagnostics settings, change the statusto On, and from the Storage account drop-down list, select the storage account az400lod11566895stor.

Save the change.

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

Reference:

https://docs.microsoft.com/en-us/azure/storage/common/storage-account-create
https://docs.microsoft.com/en-us/azure/virtual-machines/troubleshooting/boot-diagnostics

Question: 533

SIMULATION

You have a web app that connects to an Azure SQL Database named db1.
You need to configure db1 to send Query Store runtime statistics to Azure Log Analytics.
To complete this task, sign in to the Microsoft Azure portal.

Answer: See explanation below

Explanation:

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

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

{▶ CRM Database - Diagnostic settings " 1 O4UB1I*

[7 . 77] * O w*мн

5. Enter a setting name for your own reference.
6. Select a destination resource for the streaming diagnostics data: Archive to storage account, Stream to an event hub, or Send to Log Analytics.
7. For the standard, event-based monitoring experience, select the following check boxes for database diagnostics log telemetry: Query Store Runtime Statistics
8. For an advanced, one-minute-based monitoring experience, select the check box for Basic metrics.
9. Select Save.

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

Question: 534

SIMULATION

Task 1

Navigate to aex.dev.azure.com and specify the following credentials:

- UserUser1-48901628@cxamUsers.com
- Password: aFpJ2j-6M!

Use the default setting to sign up for Azure DevOps and create an Azure DevOps organization. Once the organization is created. create a private project named Project1.

You need to ensure that the maximum file size for all the Azure DevOps repositories of Project1 is 2 MB.

Answer: See explanation below

Explanation:

Step 1: Access Azure DevOps

Open your browser and navigate to https://aex.dev.azure.com.

You will be prompted to sign in. Enter the following credentials:
Username: UserUser1-48901628@cxamUsers.com
Password: aFpJ2j-6M!
Click Next or Sign in to proceed.
Note: If you see any Microsoft security challenge (like phone or email verification), follow the instructions to complete it.

Step 2: Sign Up for Azure DevOps
After signing in, you'll see a prompt to sign up for Azure DevOps.
Use the default settings provided by Azure DevOps for the sign-up process (region, default name, etc.).
Click Continue or Start free with Azure DevOps.
An Azure DevOps Organization is a container for your projects and resources. Default settings are usually fine for new organizations.

Step 3: Create the Organization
Enter a name for your organization. Azure DevOps will suggest a default name (like yourname or org- name), but you can change it if you prefer.
Verify the region is correct.
Click Continue to create the organization.
Your Azure DevOps organization is now created. It's the top-level container for all your projects.

Step 4: Create a Private Project Named Project1
In your Azure DevOps portal, find the "New Project" button and click it.
Enter the following details:
Project Name: Project1
Description: (optional)
Visibility: Private
Click Create.
A private project means that only users you add can see or access it. This is important for security and privacy.

Step 5: Enforce Maximum File Size for Repositories
By default, Azure DevOps allows files up to 100 MB to be pushed. To ensure the maximum file size for repositories in Project1 is 2 MB, you can enforce this using a pipeline validation.
Azure DevOps does not provide a direct setting for file size limits. However, you can create a pipeline to validate file size on pull requests and block changes if a file exceeds 2 MB.
Here's how to do it:
Option 1: Enforce with a pipeline validation for pull requests
In your Project1, create a new pipeline.
Use the following YAML for the pipeline:

yaml
Copy

```
trigger: none
pr:
branches:
  include:
  -
```

```
jobs:
- job: CheckFileSize
pool:
  vmImage: 'ubuntu-latest'
steps:
- script: |
    echo "Checking file sizes in the PR..."
    git fetch origin +refs/pull/*/merge:refs/remotes/origin/pr/*
    large_files=$(git diff --cached --name-only | xargs -I{} du -b {} | awk '$1 > 2097152 {print $2}')
    if [ -n "$large_files" ]; then
      echo "Error: The following files exceed 2 MB:"
      echo "$large_files"
      exit 1
    fi
  displayName: 'Check file sizes in PR'
```

This pipeline will:

Run on pull requests only.

Check the size of all files in the pull request.

If any file is larger than 2 MB (2,097,152 bytes), it will fail the pipeline and block the PR.

Option 2: Local Git configuration (for each developer)

On each developer's machine, run this command in their local repository: bash

Copy

```
git config --local core.bigFileThreshold 2M
```

This limits the file threshold for checkout, but does not prevent large files from being pushed. Using the pipeline validation (Option 1) is more secure for team environments.

Question: 535

SIMULATION

Task 2

You need to log all audit events for the Azure DevOps organization in the Log Analytics! Logs Analytics workspace.

Answer: See explanation below

Explanation:

Step 1: Understand the Requirements

You want to ensure that audit events (such as user actions, project changes, security settings, etc.) from your Azure DevOps organization are logged and available in a Log Analytics workspace in Azure. This enables centralized monitoring and security compliance.

Step 2: Prerequisites

Before starting, make sure:

You have an Azure subscription.

You have permission to create or use a Log Analytics workspace in the Azure portal.

You are a Project Collection Administrator or Organization Owner in Azure DevOps.

Step 3: Create or Identify a Log Analytics Workspace

Go to the Azure portal.

In the search bar, type Log Analytics workspaces and click the service.

Click + Create to create a new workspace (or select an existing workspace if you have one).

Provide the following:

Subscription: your Azure subscription.

Resource Group: create a new or use an existing one.

Name: a unique name for the workspace (like DevOpsAuditWorkspace).

Region: choose the same region as your Azure DevOps organization if possible.

Click Review + Create, then Create to deploy the workspace.

Step 4: Configure Azure DevOps to Stream Audit Logs

Azure DevOps can stream audit logs to your Log Analytics workspace using the Azure DevOps Audit Stream feature.

In your browser, go to your Azure DevOps organization:

https://dev.azure.com/{YourOrganizationName}

In the bottom-left corner, click on the Organization Settings gear icon.

In the left menu, click on Audit logs.

In the top-right, click on Audit streams.

Click on + Add stream to create a new stream.

In the New audit stream pane, do the following:

Stream type: select Azure Monitor Logs (Log Analytics).

Azure subscription: select the subscription containing your Log Analytics workspace.

Resource group: select the resource group.

Log Analytics workspace: select the workspace created in Step 3.

Click Save.

Step 5: Validate the Audit Stream Connection

Go back to the Audit streams page in Azure DevOps to confirm the stream shows as Connected.

To validate logs:

In the Azure portal, go to your Log Analytics workspace.

In the left menu, click on Logs.

Use the query:

kusto

Copy

AzureDevOpsAuditing

| sort by TimeGenerated desc

You should see audit events from your Azure DevOps organization appear in the results.

Question: 536

SIMULATION

Task 3

Initialize the default main branch, if it floes not already exist.

You need to protect the main branch of the Project I repository to ensure that all changes are processed by using a pull request that has traceability to a work item.

Answer: See explanation below

Explanation:

Step 1: Navigate to the Project1 Repository
In your browser, go to your Azure DevOps organization:
https://dev.azure.com/{YourOrganizationName}
In the left menu, select the project named Project1.
In the left menu under Repos, click on Files.

Step 2: Initialize the Main Branch (if it doesn't exist)
If the main branch doesn't exist (because the repo is empty):
In the Repos > Files view, you'll see a message: "This repository is empty."
Click on the Initialize button.
Choose the default branch name as main and create a README file (or any file to initialize).
Click Initialize.
This creates the main branch in your repository.

Step 3: Set Up Branch Protection for the Main Branch
Branch protection ensures that all changes go through a pull request (PR), and that PRs are linked to work items (for traceability).
In the Repos view, click on Branches in the left menu.
Locate the main branch in the list.
Click on the three dots (...) next to main and select Branch policies.

Step 4: Configure Branch Policies for main
In the Branch policies for the main branch:
Require a minimum number of reviewers:
Enable this option and set to 1 or more as per your team's standards.
Check for linked work items:
Enable the Check for linked work items option.
This ensures that every PR has at least one linked work item.
Require a linked work item:
This enforces the traceability requirement you mentioned.
Optionally, you can also:
Enable build validation if you have a pipeline to validate builds.
Enforce merge strategies (like squash merge).
Click Save changes at the bottom.

Question: 537

SIMULATION

Task 4
For Project1, you need to create a self-hosted agent pod named Pool1.
All the pipelines in Project1 must have access to Pool1.

Answer: See explanation below

Explanation:
Step 1: Understand the Requirements
You need to create a self-hosted agent for Azure Pipelines.
It will be hosted in a pod (for example, in Kubernetes).
All pipelines in Project1 should be able to use this agent pool named Pool1.

Step 2: Create the Agent Pool in Azure DevOps
Go to your Azure DevOps organization:
https://dev.azure.com/{YourOrganizationName}
In the bottom-left corner, click on the Organization settings gear icon.
In the left menu, click on Agent pools.
Click on Add pool.
Provide the following:
Pool name: Pool1
Pool type: Self-hosted
Grant access permission to all pipelines: Checked (so that all Project1 pipelines can use this pool) Click Create.
Now you have an empty agent pool named Pool1.

Step 3: Prepare the Self-hosted Agent as a Pod
You need to deploy an Azure Pipelines agent container as a pod. Here's how to do it:
Option 1: Using Kubernetes directly (YAML deployment)
Create a Kubernetes deployment YAML file for the agent pod:

```
apiVersion: apps/v1
kind: Deployment
metadata:
name: azure-pipelines-agent
labels:
  app: azure-pipelines-agent
spec:
replicas: 1
selector:
  matchLabels:
   app: azure-pipelines-agent
template:
  metadata:
   labels:
     app: azure-pipelines-agent
  spec:
   containers:
   - name: agent
     image: mcr.microsoft.com/azure-pipelines/vsts-agent:latest
```

```
env:
- name: AZP_URL
  value: "https://dev.azure.com/{YourOrganizationName}"
- name: AZP_TOKEN
  valueFrom:
    secretKeyRef:
      name: azure-pipelines-token
      key: token
- name: AZP_POOL
  value: "Pool1"
- name: AZP_AGENT_NAME
  value: "agent-pod"
```

Create a Kubernetes secret for the personal access token (PAT):

```
kubectl create secret generic azure-pipelines-token --from-literal=token=<your-PAT>
```

Note:

Replace <your-PAT> with a PAT that has “Agent Pools (Read & manage)” scope.

Replace {YourOrganizationName} with your actual Azure DevOps organization.

Deploy the pod:

bash

Copy

```
kubectl apply -f azure-pipelines-agent.yaml
```

Step 4: Validate the Agent Connection

Go back to Organization Settings > Agent pools > Pool1 in Azure DevOps.

You should see the agent pod connected and listed as online.

Step 5: Use the Pool in Pipelines

By default, all pipelines in Project1 will now have access to the Pool1 agent pool because you granted access during the pool creation.

In your pipeline YAML files, specify the pool name to use the self-hosted agent:

yaml

Copy

```
pool:
name: Pool1
```

Question: 538

SIMULATION

Task 5

You need to create an instance of Azure App Configuration and a feature flag named feature 1.

Feature1 must be enabled and must expire in one week.

Answer: See explanation below

Explanation:

Step 1: Log in to the Azure Portal
Go to https://portal.azure.com.
Log in with your Azure account.

Step 2: Create an Azure App Configuration Instance
In the search bar at the top, type "App Configuration" and select App Configuration from the search results.
Click + Create.
Fill in the required fields:
Subscription: Your Azure subscription.
Resource Group: Create a new one or select an existing one.
Name: Provide a globally unique name (for example, myappconfig-instance).
Location: Choose a region close to your application.
Click Review + Create, then Create.
The instance will take a few seconds to deploy.

Step 3: Access the App Configuration Instance
Once the deployment is complete, click Go to resource.
You are now in the App Configuration resource.

Step 4: Enable Feature Management
In the left menu, click on Feature Manager (Preview).
Click + Add to create a new feature flag.

Step 5: Create the Feature Flag
Provide the following details:
Feature flag name: feature1
Label: (optional)
Description: (optional)
State: Set to Enabled.
Click on the Add filters link if you want to add targeting filters (optional).
Click Save to create the feature flag.

Step 6: Set the Expiration Date for the Feature Flag
In the Feature Manager list, find the newly created feature1 flag.
Click on the three dots (context menu) next to feature1 and select Edit.
In the Edit feature pane, look for the Expiration date setting.
If there is no direct UI for expiration date, you can store it as a custom key-value (metadata).
Here's how to store expiration as metadata:
In the App Configuration left menu, click on Configuration Explorer.
Locate the key: .appconfig.featureflag/feature1.
Click on the key to edit.
Add a new label or add a custom key in the Content type or tags section to store expiration metadata:
Key: expiration

Value: yyyy-MM-dd (set to today + 7 days)
For example, if today is June 4, 2025, set the expiration as 2025-06-11.

Question: 539

SIMULATION

Task 6

For Project 1. you need to create a service connection that can be used to deploy resources to the RGHod489Q1628 resource group.
The service connection must use the ManagedJd1 identity and workload identity federation.

Answer: See explanation below

Explanation:

Task 6: Create a Service Connection for Resource Group Deployment using Managed Identity and Workload Identity Federation

Step 1: Understand the Requirements
You want to deploy resources in the RGHod489Q1628 resource group.
The service connection must:
Use the ManagedJd1 managed identity.
Use workload identity federation (OIDC-based authentication for enhanced security).

Step 2: Verify Prerequisites
You need to ensure:
The ManagedJd1 managed identity exists in your Azure subscription.
Your Azure DevOps project (Project1) is linked to an Azure Active Directory tenant (for OIDC support).
You have the Owner or User Access Administrator role on the RGHod489Q1628 resource group.

Step 3: Assign Role to Managed Identity
Go to the Azure Portal.
In the search bar, type Managed Identities and select Managed Identities.
Locate and click on the ManagedJd1 identity.
In the left menu, click Azure role assignments.
Click + Add role assignment.
Set the following:
Scope: Resource Group
Subscription: Your subscription
Resource Group: RGHod489Q1628
Role: Contributor (or appropriate role)
Click Save.
This step ensures ManagedJd1 has permissions to deploy resources to RGHod489Q1628.

Step 4: Create a Federated Credential for Workload Identity Federation
In the Azure Portal, navigate to the ManagedJd1 managed identity.
In the left menu, click Workload identity federation (preview).
Click + Add a federated credential.
Configure as follows:
Federated credential name: devops-oidc
Issuer: https://vstoken.actions.githubusercontent.com (or use the default https://pipelines.actions.githubusercontent.com for Azure DevOps)
Subject identifier: Use the following format for Azure DevOps:
css
Copy
system:azuredevops:{organizationName}:{projectName}
For example:
css
Copy
system:azuredevops:{YourOrganizationName}:{Project1}
Audience: api://AzureADTokenExchange
Click Add.
This federated credential establishes trust between your Azure DevOps project and the managed identity.

Step 5: Create a Service Connection in Azure DevOps
Go to your Azure DevOps project (Project1) in the browser.
In the left menu, click Project settings.
Under Pipelines, click Service connections.
Click New service connection.
Choose Azure Resource Manager.
Choose the authentication method:
Select Workload identity federation.
Configure the service connection:
Scope level: Resource Group.
Resource Group: RGHod489Q1628.
Subscription: Your subscription.
Authentication method: Managed Identity with workload identity federation.
Managed Identity: Enter the client ID or select ManagedJd1.
Service connection name: e.g., Project1-RGHod489Q1628-Conn.
Grant access permission to all pipelines (recommended).
Click Save.

Step 6: Validate the Service Connection
After creation, click on the new service connection to Verify it.
Ensure the connection test is successful.
You can now use this service connection in your pipelines for deploying resources to RGHod489Q1628.

Question: 540

SIMULATION
Task 7
You need to create a pipeline to deploy a Docker image. The commit must be created in a new branch named

azure-pipelines.
The pipeline must be created by using the predefined Docker template.

Answer: See explanation below

Explanation:

Task 7: Create a Pipeline to Deploy a Docker Image in a New Branch

Step 1: Create a New Branch
In your Azure DevOps Project (Project1), navigate to Repos.
In the left menu, click on Branches.
Click New branch.
Enter:
Branch name: azure-pipelines
Based on: main (or your default branch)
Click Create.
This creates a new branch named azure-pipelines.

Step 2: Switch to the New Branch
In Repos, click on Files.
In the top-left branch selector, choose azure-pipelines.

Step 3: Use the Predefined Docker Pipeline Template
Azure Pipelines has predefined YAML templates for common tasks, including Docker.
Here's how to add it:
In the azure-pipelines branch, click New file.
Name the file: azure-pipelines.yml.
Click the Show templates button in the YAML editor (if available).
Search for the Docker template in the list of predefined templates.
Select the Docker template to auto-populate the YAML file.
If the editor doesn't show the template picker, you can manually add the following basic Docker pipeline template: yaml
Copy

```
- Predefined Docker pipeline template
trigger:
- main

resources:
- repo: self
variables:
imageName: 'mydockerimage'

stages:
- stage: Build
displayName: Build and push stage
```

```
jobs:
- job: Build
  pool:
    vmImage: 'ubuntu-latest'
  steps:
  - task: Docker@2
    displayName: Build and push an image to container registry
    inputs:
      command: buildAndPush
      repository: $(imageName)
      dockerfile: '**/Dockerfile'
      containerRegistry: '<your-container-registry-service-connection>'
      tags: |
        latest
```

Replace <your-container-registry-service-connection> with your actual Azure Container Registry (ACR) or DockerHub service connection name.

Step 4: Commit the Pipeline to the New Branch

In the YAML editor, review and adjust if needed (like setting the correct container registry service connection).

In the bottom commit message box:

Message: e.g., Add Docker deployment pipeline

Ensure the branch is set to azure-pipelines.

Click Commit (not commit to main).

This creates the pipeline YAML file in the new azure-pipelines branch.

Step 5: Create the Pipeline in Azure DevOps

In the left menu, click Pipelines.

Click New pipeline.

Choose:

Azure Repos Git.

Select your repo.

Select Existing Azure Pipelines YAML file.

In the branch selector, choose the azure-pipelines branch.

Select the YAML file path (azure-pipelines.yml).

Click Continue and then Run to validate the pipeline.

Question: 541

SIMULATION

Task 8

In Project1, you need to create a service hook that will integrate with Azure Storage. Use the queue1 queue from the storage48901628 storage account.

Answer: See explanation below

Explanation:

Task 8: Create a Service Hook to Integrate with Azure Storage (queue1)

Step 1: Verify the Storage Account and Queue

In the Azure portal (https://portal.azure.com):

Go to Storage accounts.

Find the storage48901628 storage account.

In the left menu, click on Queues.

Ensure that a queue named queue1 exists.

If not, create it by:

Clicking + Queue.

Enter Name: queue1.

Click OK.

Step 2: Prepare the Storage Account's Shared Access Signature (SAS)

Azure DevOps uses SAS tokens to send messages to a storage queue.

In the Azure portal, go to the storage48901628 storage account.

In the left menu, click Shared access signature.

Set the following:

Allowed services: Queue.

Allowed resource types: Object.

Allowed permissions: Add, Write.

Start and expiry date/time: Adjust as needed (e.g., 1 year).

Click Generate SAS and connection string.

Copy the Queue SAS URL (e.g., https://storage48901628.queue.core.windows.net/queue1?sv=...).

Step 3: Create the Service Hook in Azure DevOps

In your Azure DevOps Project (Project1):

In the bottom-left, click Project settings.

In the left menu, click Service hooks.

Click + Create subscription.

Step 4: Configure the Azure Storage Service Hook

In the Create a subscription wizard, choose Azure Storage Queues and click Next.

Choose the Trigger:

For example: Code pushed (or other events you want to send to the queue).

Click Next.

Configure the Action:

Queue endpoint URL: The Queue SAS URL you copied in Step 2 (including the queue1 endpoint and SAS token).

Example:

bash

Copy

https://storage48901628.queue.core.windows.net/queue1?sv=...

Step 5: Finalize and Test the Integration

Click Test to ensure Azure DevOps can send a message to the queue.

If the test succeeds, click Finish to create the service hook.

Step 6: Validate in the Azure Portal

In the Azure portal, go to storage48901628 > Queues > queue1.

You should see new messages appear in the queue whenever the configured Azure DevOps event (like push or PR) occurs.

Question: 542

SIMULATION

Task 9

Initialize the default main branch, if it does not exist already.

You need to create a pipeline that will execute a .NET Core build task for the src directory.

The pipeline code must be stored in a new branch named azure-pipelinesT.

The pipeline must run automatically when a pull request against the mam branch is created or updated.

Answer: See explanation below

Explanation:

Task 9: Initialize the main branch (if it doesn't exist) and create a pull request-triggered pipeline for .NET Core build

Step 1: Initialize the main Branch

In your Azure DevOps Project (Project1), navigate to Repos.

In the left menu, click Files.

If you see a message like "This repository is empty," click Initialize.

Create a README.md file (or another file) and set the default branch name to main.

Click Initialize.

If the main branch already exists, you can skip this step.

Step 2: Create a New Branch for the Pipeline

In Repos > Files, click on the Branches tab.

Click New branch.

Enter:

Branch name: azure-pipelinesT

Based on: main

Click Create.

This creates a new branch named azure-pipelinesT from main.

Step 3: Create the Pipeline YAML File in the New Branch

Switch to the azure-pipelinesT branch.

In Files, click New file.

Name the file: azure-pipelines.yml.

Step 4: Write the YAML Pipeline for .NET Core Build

Here's a sample YAML file:

```
trigger: none

pr:
branches:
  include:
    - main

pool:
vmImage: 'windows-latest'

variables:
buildConfiguration: 'Release'

steps:
- task: UseDotNet@2
displayName: 'Install .NET Core SDK'
inputs:
  packageType: 'sdk'
  version: '8.x' # Or specify the exact version of .NET Core SDK you need
  installationPath: $(Agent.ToolsDirectory)/dotnet

- script: dotnet build src --configuration $(buildConfiguration)
displayName: 'Build .NET Core project'
```

This pipeline:

Runs only when a pull request is created or updated targeting the main branch.

Uses the windows-latest agent.

Installs the .NET SDK and builds the project in the src directory.

Step 5: Commit the Pipeline File

In the bottom commit message box:

Message: Add PR-triggered pipeline for .NET Core build

Ensure the branch is set to azure-pipelinesT.

Click Commit.

Step 6: Create the Pipeline in Azure DevOps

In the left menu, click Pipelines.

Click New pipeline.

Select:

Azure Repos Git.

Select your repository.

Choose Existing Azure Pipelines YAML file.

In the branch selector, choose the azure-pipelinesT branch.

Select the YAML file path (azure-pipelines.yml).

Click Continue and Run to validate.

Step 7: Test the Pipeline Trigger

Create a new branch (e.g., feature/test-pr).

Push a commit and create a pull request targeting main.

Verify that the pipeline runs automatically for this PR in Pipelines.

Question: 543

SIMULATION

Task 10

You need to create a variable group named varGroup1 that will contain the following variables and default values:

- username: UserI-48901628
- password: aFpJ2j+6M!

The PASSWORD variable must be secret.

Answer: See explanation below

Explanation:

Task 10: Create a Variable Group in Azure DevOps

Step 1: Navigate to Library

In your Azure DevOps Project (Project1), click on the Pipelines section in the left menu.

In the Pipelines submenu, click on Library.

Step 2: Create a New Variable Group

Click on the + Variable group button.

Enter the Name:

nginx

Copy

varGroup1

Step 3: Add Variables

Add the first variable:

Click Add.

In the Name field, enter:

nginx

Copy username

In the Value field, enter:

Copy

UserI-48901628

Add the second variable:

Click Add again.

In the Name field, enter:

nginx

Copy password

In the Value field, enter:

Copy aFpJ2j+6M!

Check the box Keep this value secret.

Step 4: Save the Variable Group

Click Save at the top to save the variable group.

Question: 544

SIMULATION

Task 11

You need to create a new artifact feed named artifact_feed.

Only 10 versions of the packages must be retained, and downloaded packages must be retained for 90 days.

Answer: See explanation below

Explanation:

Task 11: Create an Artifact Feed with Specific Retention Policies

Step 1: Navigate to Artifacts

In your Azure DevOps Project (Project1), click on the Artifacts section in the left menu.

Step 2: Create a New Feed

Click on the + New feed button.

In the Name field, enter:

nginx

Copy

artifact_feed

Choose the visibility of the feed:

Project scoped (only accessible to Project1).

Or Organization scoped (accessible across your Azure DevOps organization).

Leave other settings at default unless you have specific permissions or upstream sources.

Click Create.

This creates a new feed named artifact_feed.

Step 3: Configure Retention Policies

In the Artifacts section, click on the feed named artifact_feed to open its details page.

In the top-right corner, click on the three dots (...) and select Feed settings.

Step 4: Set the Maximum Number of Versions to Retain

In the Feed settings page, find the Retention policies section.

Set Maximum number of versions per package to:

Copy 10

Click Save.

This ensures that only the latest 10 versions of any package are retained in the feed.

Step 5: Set the Downloaded Package Retention Period

In the Feed settings page, under Retention policies, find Downloaded package retention.

Set Number of days to retain downloaded packages to:

Copy 90

Click Save.

This ensures that downloaded packages will be retained in the cache for 90 days before deletion.

Question: 545

SIMULATION

Task 12

In Project 1. you need to ensure that you can add the following iterations to items added tc Az-fe Ecards:

- Sprint 1
- Sprint 2
- Sprint 3

Sprint 3 must have a start date of January 1 of next year and an end date of Januan / 31 of next year.

Answer: See explanation below

Explanation:

Task 12: Configure Iterations in Project1 for Azure Boards

Step 1: Navigate to Project Settings

In your Azure DevOps Project (Project1), click on the Project settings gear icon in the lower-left corner.

Step 2: Manage Project Iterations (Sprints)

In the Boards section of the Project settings, click Project configuration.

In the left menu, click on Iterations.

Step 3: Add Iterations

Click on the + New child link to add a new iteration under the root node (e.g., Project1).

Enter the following details:

Name: Sprint 1

Leave the start and end dates empty (or set them as needed for your planning).

Click Save and close.

Repeat the process:

Click + New child again.

Name: Sprint 2

Leave the start and end dates empty.

Click Save and close.

Finally, create Sprint 3 with specific dates:

Click + New child.

Name: Sprint 3

Start date: January 1 of next year (e.g., 2026-01-01).

End date: January 31 of next year (e.g., 2026-01-31).

Click Save and close.

Step 4: Assign Iterations to Work Items

With these iterations created, you can now assign them to work items (e.g., tasks, bugs, features) using Azure Boards.

In Boards, click on Work items.

Open any work item.

In the Iteration Path field, select one of your newly created iterations (Sprint 1, Sprint 2, or Sprint 3).
This ensures that work items can be tracked under the appropriate sprint.

Question: 546

HOTSPOT

You ate using Agile process methodologies in Azure DevOps.
You need to deploy a dashboard that will provide progress reports for the following work items:

- How long it took to close a work item after it was created.
- How long it took to close a work item after the work was started.

Which type of widget should you use for each work item? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

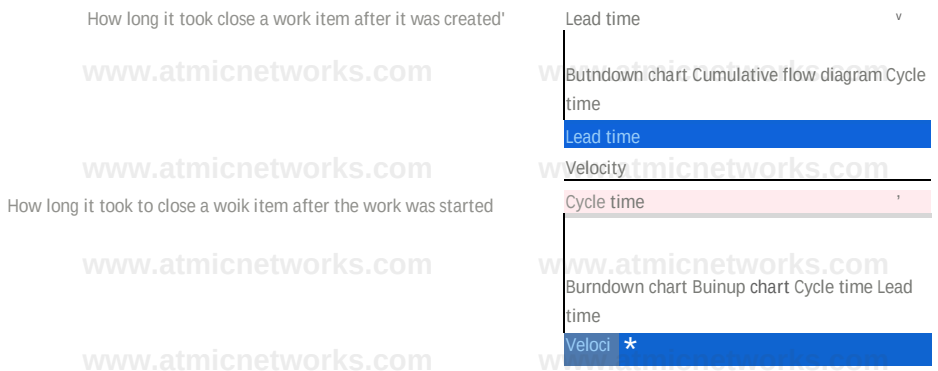

Answer:

Explanation:

Answer Area

How long it took close a work item after it was created' Lead time

How long it took to close a work item after the work was started Cycle time

Question: 547

HOTSPOT

You have a Bicep file named Main.bicep.
You run the following command.

```
iconfig • j{
    Name = 'StorageDeployment'
    ResOurceGroupUame * 'RG1'
     Templatefile = ' Amain, bicep' whatif * (true }

New ■ Az Resour ceGroupOeployraent ^config
```

You receive the following response

```
The deployment will update the following scope:

Scope: /subscriptions/mySubscription/resourceGroups/R61

  + Microsoft.Storage/storageAccounts/storageAccountl [2023-05 01]
      apiversion:                    "2023-05-01'
      id:
"/subscriptions/mySubscription/resourceGroups/RGl/providers/Microsoft.Storage/storageAccounts/storageAccounti" kind:
                                      "Blobstorage"
      location:                      "westeurope"
      name:                          "storageAccountl"
      properties.acresslier:         "Hot"
      properties.allowBlobPublicAccess: false properties.niniBunUsVersions    'TIS1_2'
      situ, name:                    "$tandard_GRS"
      type:                          "Microsoft.Storage/storageAc counts"

  -   Microsoft.Storage/storageAccounts/storageAccount2 12023-05-31]
  -     ■ properties.accessTier:     "Hot" => "Cold"
  -     properties.allowBlobPublicAccess: true => false
  -     properties.minimumTlsVersion: "TIS1 0" -> "TLS1 2"
    ' sku.name:                      "Standard GR5" s> "Standard IRS"

  * Microsoft.Storage/storageAccounts/storageAccount*

  Resource changes: 1 to create, 1 to modify, 1 to ignore.
```

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

NOTE: Each correct selection is worth one point.

Answer Area

Answer

Explanation:

Answer Area

StorageAccount3 will be [answer choice]. created

The command will [answer choice] the resources in RG1. evaluate

Question: 548

You are designing a build and release pipeline.
You need to implement a shift-left testing strategy for the pipeline.
Which types of tests should you include?

A. integration, user acceptance, and regression
B. unit, smoke, and integration
C. formal acceptance, user acceptance, and smoke
D. functional, regression, and formal acceptance

Answer: B

Explanation:

Question: 549

You are designing a build and release pipeline.
You need to implement continuous integration and continuous delivery (CI/CD) for the pipeline Which automated activities should you include in the pipeline?

A. unit tests, integration tests, and acceptance tests
B. functional tests, production tests, and production release
C. functional tests, regression tests, and formal assessments
D. unit tests, production release, and production deployment

Answer: A

Explanation:

Question: 550

You use GitHub Enterprise to manage source code and deployments.
From GitHub Actions, you create a reusable workflow named workflow 1.
You need to ensure that the secrets defined in the calling workflow are passed to workflow1.

What should you include in the YAML file for the calling workflow? A.

```
with:
  secrets: all
```

B.

```
on:
  workflow call:
    secrets;
      inherit: true
```

C.

```
I  parent workflow: secrets: inherit
```

D.

```
on;
  workflow_dispatch:
    inputs:
      secrets: all
```

Answer: C

Explanation:

Question: 551

HOTSPOT

You manage the Git repository for a large enterprise application.

You need to reduce 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.

Answer Area

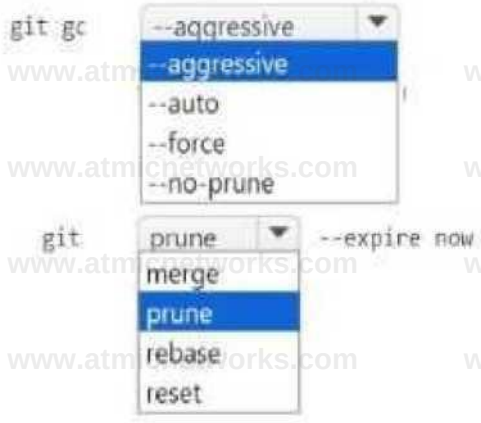

Answer:

Explanation:

Answer Area

git gc aggressive

git prune T --expire now

Question: 552

DRAG DROP

You have a GitHub repository that contains the code for an app named App1.

App1 depends on a library of functions from a repository at https://github.com/contoso/afeed.

You need to keep a clone of a feed repository as a subdirectory of the App1 repository.

How should you complete the Git command? 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.

Answer:

Explanation:

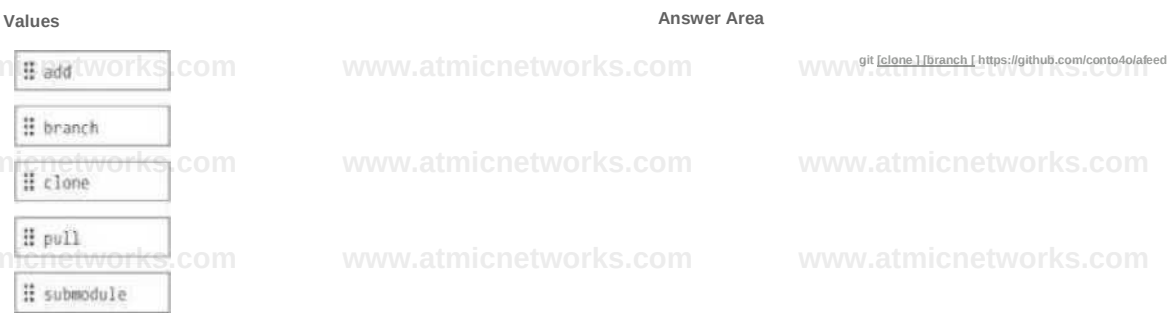

Question: 553

You have an Azure DevOps organization.

You are designing an automated security and compliance solution that will use GitHub Advanced Security for Azure DevOps.

You need to ensure that security checks and compliance requirements are enforced automatically

throughout the development lifecycle.

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

NOTE: Each correct selection is worth one point.

A. automated approvals
B. code linting
C. secret scanning push protection
D. dependency scanning
E. pull request annotation scanning

Answer: C, D

Explanation:

Question: 554

DRAG DROP

You have an Azure subscription that contains a resource group named RG1.

You have a template artifact named template1.

You plan to use Azure Pipelines to deploy an app named App1 to RG1. App1 requires template1.

You need to deploy template1 to RG1.

How should you configure the pipeline task? 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.

Answer Area

```
-   task: AtureRe5OurceManagerTemplateDeployment^3
inputs:
   deploymentscope: 'Resource Group'
   azureRescurceManagerConnection: azureserviceconnection
   subscription^: $(subscriptionld)
   action: 'Create Or Update Resource Group'
   resourceGroupName: 'ARMPipelinesLAMP-rg'
   location: St location)
   templeteLocation:
   csmFile: '$(Build.ArtifactStagingDirectory)/azuredeploy• json'
   csmParametersFile: '$(Build.ArtifactStagingDirectory)/azuredeploy.parameters.jLpn'
   overrideParameters: '-siteName $($iteName) -administratorLogin ${adminUser) -administratorLoginPassword $(ARM_pASS)'
   deploymentMode:
```

Answer:

Explanation:

Values

- azuredeploy.json
- Incremental
- 'Linked artifact'
- 'Management Group'
- 'Resource Group'

Answer Area

```
-   task: AzureR^sourcsWanagerTMplatjDsploynient^
inputs:
    deploymentscope: 'Resource Group'
    azureResourceManagerConnection: azureserviceconnection
    subscriptionld: J(subscription^)
    action: 'Create Or Update Resource Group'
    resourceGroupName: 'ARMPipelinasLAMP-rg'
    location: $(location)
    templateLocation: |'Linked artifact'            |
    csmFile: *$(Build.ArtifactStagingDirectory)/azuredepIoy.json'
    csmParametersFile: '£(Build.ArtifactStagingDirectory)/azuredeploy.parameters .j^n'
    ouerrideParameters: '-siteName $(siteMame) -administratorLogin J(adminUser) -administratorLoginPassword S(ARM_PASS)'
    dsployimntrtodc:   Intramtal
```

Question: 555

You have an Azure subscription that contains an Azure Kubernetes Service (AKS) instance named AKS1. You collect and analyze metrics for AKS1 by using the Azure Monitor managed service for Prometheus. You need to analyze the performance of AKS1.

Which query language should you use?

A. PromQL
B. KQL
C. PL/SQL
D. SparkQL

Answer: A

Explanation:

Question: 556

You have an Azure DevOps pipeline named Pipeline1.
You need to track the stability of code changes in Pipeline1.
Which metric should you use?

A. test pass rate
B. deployment frequency (DT)
C. lead time for changes (LT)
D. time to detect (TTD)

Answer: B

Explanation:

Question: 557

HOTSPOT

You have an Azure Pipelines pipeline named Pipeline1 that has the following YAML definition.

```
pool:
    name: Azure Pipelines demands:
    -   msbuIId - visualstudio

steps:
    task: vsBuildgi
    displayNamo: 'Build solution ••\*.sln'
    inputs:
        solution: ^(Parameters.Solution)'
        platform: '{(BuildPlatform)'
        configuration: '$(8uildConfiguration)'

-   task: CopyFiles^I
    displayName: 'Copy Files to; $(Build.ArtifactStagingDirectory)' inputs:
        sourcefolder: '^System.DefaultMorkingbirectory)*
        contents: '*"\bin\|(Buildconfiguration)\"'
        TargetFolder: S(Build.ArtitactStagingOirectory)'
    condition: succeededOrFailed()
    task: PublishBuildArtifacts^I
    displayName: 'Publish Artifact: drop'
    inputs:
        pathtoPublish: '$(Buiid.ArtifactstagingDirectory)'
    condition: succeededorFailed()
```

For each of the following statements, select Yes if True. Otherwise select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Pipelinel consists of three stages.	Q	
S(BuildPlatfori) is a predefined variable.	O	O
SfSystea.OefaultNorkingDirectory) is a predefined variable		

Answer:

Explanation:

Answer Area

Statements	Yes	No
Pipeline! consists of three stages.	®	
{(BulldPlatfoni) is a predefined variable		
{(system.Defaultwkingbirectory) is a predefined variable	®	

Question: 558

You have an Azure subscription.

You create two Bicep templates named Template1 and Template2 that will be used to create a virtual machine and a website.

You need to create a template named Template3 that will reuse logic from Template1 and Template2.

What should you define first?

A. parameters
B. modules
C. resources
D. outputs

Answer: B

Explanation:

Question: 559

DRAG DROP

You are using Agile process methodologies and Azure Boards in Azure DevOps. You import the work items shown in the following table.

Name	Description
Iteml	The solution was implemented but **NOT** verified. The item must appear on the backlog.
Item2	The work was **NOT** started The item must appear on the backlog
Item 3	The work **is** complete. The item must **NOT** appear on the backlog The 'tern must appear on the board

You need to categorize the items.

Which state should you assign to each item? To answer, drag the appropriate states to the correct isms. Each state may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Question: 560

DRAG DROP

You have a GitHub organization.

You ate creating a GitHub Actions workflow.

You need to perform authenticated AP1 requests by using a GitHub app name App1.

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.

Answer:

Explanation:

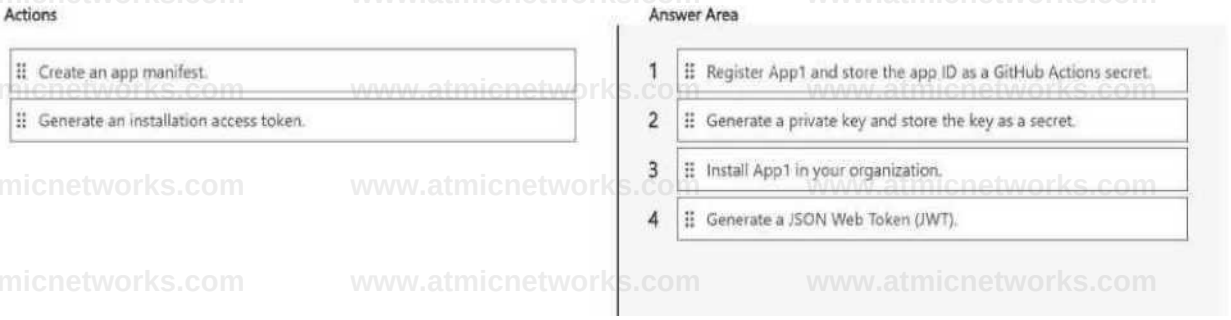

Question: 561

Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals. More than one solution in the set might solve the problem. It is also possible that none of the solutions in the set solve the problem.

After you answer a question in this section, you will NOT be able to return. As a result, these questions do not appear on the Review Screen.

You have an on-premises web server named Server1.

You need to ensure that you can build .NET apps and deploy the apps to Server1 by using Azure Pipelines.

Solution: You deploy a virtual machine scale set and select a custom Ubuntu image that includes the .NET command-line interface (CU) tools.

Does this meet the goal?

A. Yes

B. No

Answer: A

Explanation:

Question: 562

HOTSPOT

You have a GitHub organization.

You configure the personal access token (PAT) policy shown in the following exhibit.

Restrict access via personal access tokens (classic)

8y default. personal access tokens (classic) can access content owned by your organization via the GitHub API or Git over HTTPS. This includes both public and private resources such as repositories.

0 Allow access via personal access tokens (classic)

API and Grt access will be allowed using an organization member $ personal access token (classic)

9 Restrict access via personal access tokens (classic) *

Organization members will not be allowed to access your organization using a personal access token (classic)

Use the dropdown menus to select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

Answer Area

Classic PATs (answer choice) access the public resources of the organization.

|can

cannot

require additional configurations to

SSH keys created by using classic PATs (answer choice) access the resources of the organization

cannot

I require additional configurations to

Answer:

Explanation:

Answer Area

Classic PATs (answer choice] access the public resources of the organization, can

SSH keys created by using classic PATs [answer choice) access the resources of the can organization.

Question: 563

You have a widget named Appl and an Azure DevOps YAML pipeline named Pipelinel. App1 is released by using Pipelined Pipeline1 contains a single stage and a single job. You need to ensure that Appl is approved before the app is released. Which two actions should you perform? Each correct answer presents part of the solution. NOTE; Each correct selection is worth one point.

A. Add a stage named deployment to the YAML pipeline.

B. Change Pipelinel to a deployment job.

C. Create a service connection.

D. Create an environment an tod add a check.

E. Configure branch control.

Answer: A, D

Explanation:

Question: 564

HOTSPOT

You have a GitHub organization that contains a repository and an organizational secret named DB.Credential. You need to build a workflow that will configure a server by using PowerShell and DB.Credential. 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

```
01 steps:
02   - shell: pwsh
03     env:
04       DB_CREDENTIAL:
05     run:
06       configure_server
```

DB_CREDENTIAL:
- "$DB_CREDENTIAL"
- "$env:DB_CREDENTIAL"
- "$secrets.DB_CREDENTIAL"
- "%DB_CREDENTIAL%"
- "secrets.DB_CREDENTIAL"
- ${{secret.DB_Credential}}

configure_server
- "$DB_CREDENTIAL"
- $env:DB_CREDENTIAL"
- "$secrets.DB_CREDENTIAL"
- "%DB_CREDENTIAL%"
- "secrets.DB_CREDENTIAL"
- $((secrets.DB_Credential))

Answer:

Explanation:

Answer Area

```
01 steps:
02     - shell: pash
03     env:
04         D6_CREO€NTIAL:     ${{ secrets.DB.Credential)}
05     run: |
06         conf igure_server      "$env:D8 .CREDENTIAL"
```

Question: 565

You have a GitHub repository named repol.

You commit a change to repol.

You need to update the commit message. The solution must ensure that the version number remainsthe same.

Which command should you run?

A. git commit --patch

B. git commit --amend

C. git status --renames

D. git restore --patch

Answer: B

Explanation:

Question: 566

You have an Azure subscription.

You need to ensure that users can provision preconfigured cloud-based environments for developing apps. The solution must minimize administrative effort.

What should you include in the solution?

A. Microsoft Dev Box

B. Windows 365 Cloud PC

C. Azure Deployment Environments

D. Azure Virtual Desktop

Answer: C

Explanation:

Question: 567

You have an Azure DevOps release pipeline named Pipeline1. You plan to use Pipelinel to deploy infrastructure as code (laC) resources to Azure. You need to create templates to define which resources to deploy. Which template formats should you use?

A. Azure Resource Manager (ARM) only

B. Biceponly

C. XML only

D. Azure Resource Manager (ARM) and Bicep only

E. Azure Resource Manager (ARM), Bicep. and XML

Answer: D

Explanation: