vce trainer



Reference Solution for AZ-400

AZ-400

AZ-400

Score: 800/1000

Version:

Time Limit: 120 Minutes

Licensed to Berend Post berend.post@it-university.nl

Design a DevOps Strategy

(9 questions)

Question 1

HOTSPOT

Your company uses Team Foundation Server 2013 (TFS 2013).

You plan to migrate to Azure DevOps.

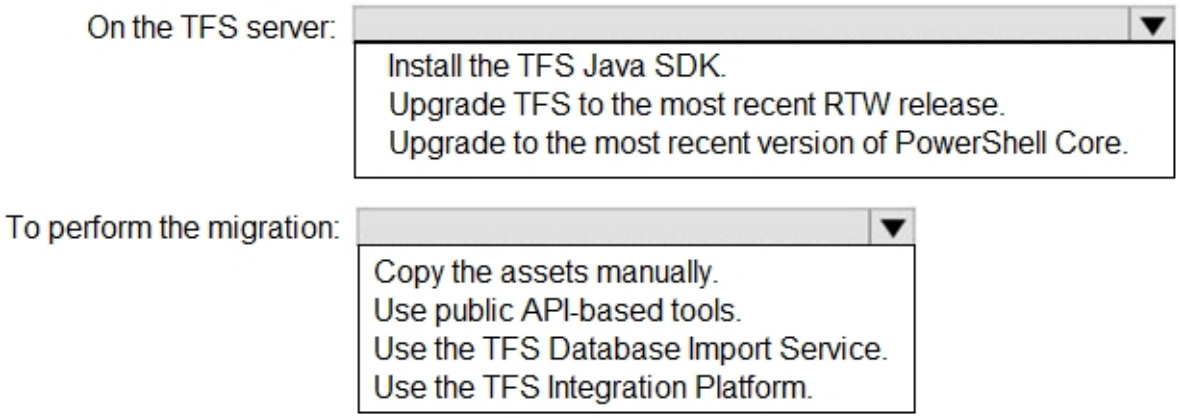
You need to recommend a migration strategy that meets the following requirements:

Preserves the dates of Team Foundation Version Control changesets
Preserves the changes dates of work items revisions
Minimizes migration effort
Migrates all TFS artifacts

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

NOTE: Each correct selection is worth one point.

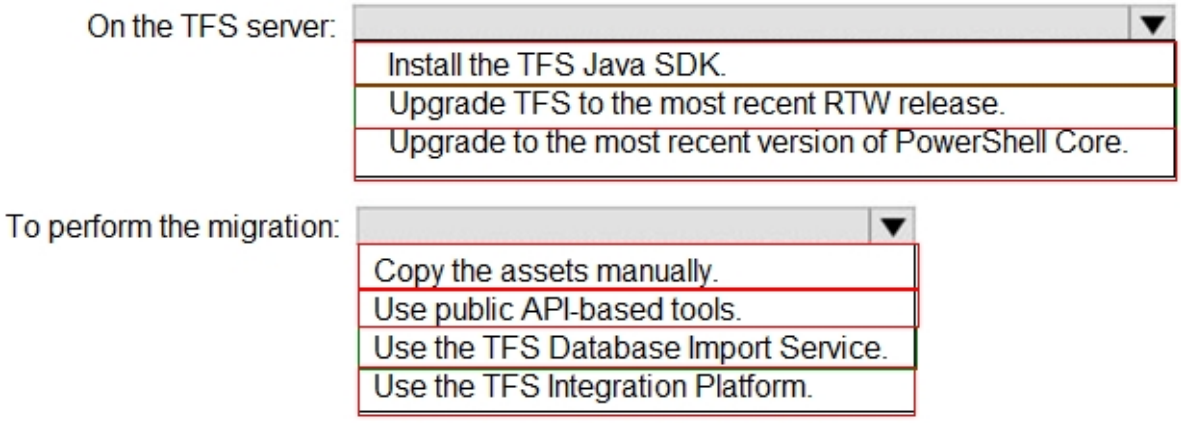
Answer Area



Licensed to Berend Post berend.post@it-university.nl

Solution:

Answer Area



Explanation:

Explanation:

Box 1: Upgrade TFS to the most recent RTM release.

One of the major prerequisites for migrating your Team Foundation Server database is to get your database schema version as close as possible to what is currently deployed in Azure Devops Services.

Box 2: Use the TFS Database Import Service

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

References: Team Foundation Server to Azure Devops Services Migration Guide

Licensed to Berend Post berend.post@it-university.nl

Question 2

You manage build pipelines and deployment pipelines by using Azure DevOps.

Your company has a team of 500 developers. New members are added continually to the team.

You need to automate the management of users and licenses whenever possible.

Which task must you perform manually?

- ☐ modifying group memberships
- ☐ adding users
- ☐ assigning entitlements
- ☒ procuring licenses

Explanation:

Incorrect Answers:

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

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

License

Extensions

Project/Team memberships

References:

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

Licensed to Berend Post berend.post@it-university.nl

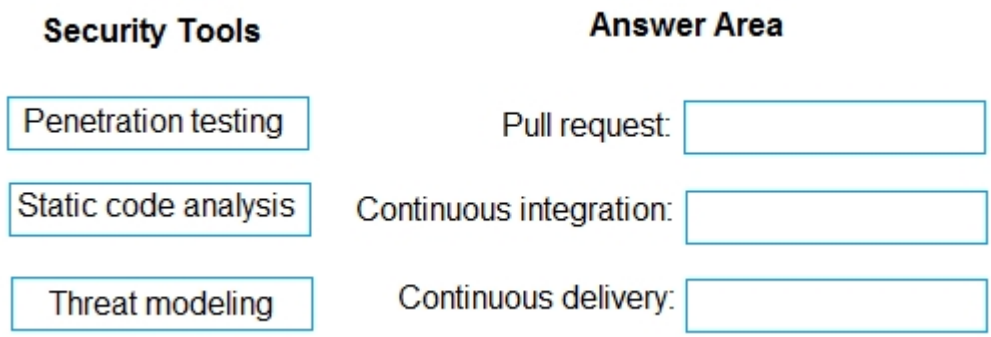
Question 3

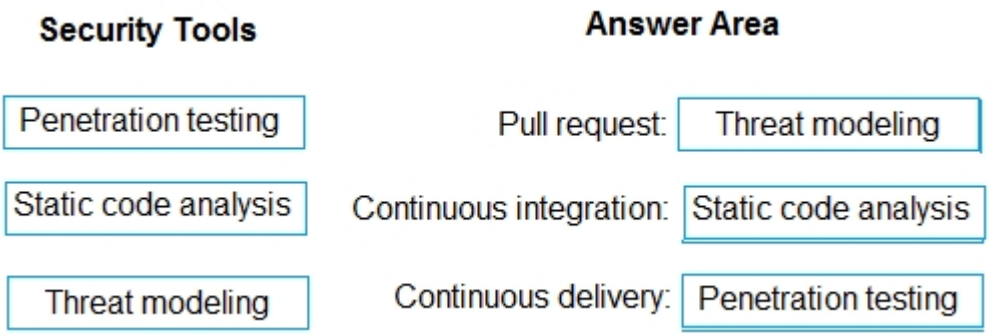
DRAG DROP

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 tools to the correct targets. 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 view content.

NOTE: Each correct selection is worth one point.





Solution:

Explanation:

Explanation:

Box 1: Threat modeling

Threat modeling's motto should be, "The earlier the better, but not too late and never ignore."

Box 2: Static code analysis

Validation in the CI/CD begins before the developer commits his or her code. Static code analysis tools in the IDE provide the first line of defense to help ensure that security vulnerabilities are not introduced into the CI/CD process.

Licensed to Berend Post berend.post@it-university.nl

Box 3: Penetration testing

Once your code quality is verified, and the application is deployed to a lower environment like development or QA, the process should verify that there are not any security vulnerabilities in the running application. This can be accomplished by executing automated penetration test against the running application to scan it for vulnerabilities.

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

Question 4

You are developing a 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 team an email message when the front end fails to return a status code of 200.

Which feature should you use?

- ☐ Service Map in Azure Log Analytics
- ☐ Availability tests in Azure Application Insights
- ☐ Profiler in Azure Application Insights
- ☒ Application Map in Azure Application Insights

Explanation:

Explanation:

Application Map helps you spot performance bottlenecks or failure hotspots across all components of your distributed application. Each node on the map represents an application component or its dependencies; and has health KPI and alerts status.

Incorrect Answers:

A: Service Map automatically discovers application components on Windows and Linux systems and maps the communication between services. You can use it to view your servers as you think of them--interconnected systems that deliver critical services. Service Map shows connections between servers, processes, and ports across any TCP-connected architecture with no configuration required, other than installation of an agent.

References: https://docs.microsoft.com/en-us/azure/azure-monitor/app/app-map

Licensed to Berend Post berend.post@it-university.nl

Question 5

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 following requirements:

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?

- ☐ Microsoft Project
- ☐ Bamboo
- ☐ Microsoft Lync
- ☒ Microsoft Teams

Explanation:

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 to facilitate business-to-business project management.

Teams integrates with Azure DevOps.

References: https://searchunifiedcommunications.techtarget.com/definition/Microsoft-Teams

Licensed to Berend Post berend.post@it-university.nl

Question 6

You use Azure Artifacts to host NuGet packages that you create.

You need to make one of the packages available to anonymous users outside your organization. The solution must minimize the number of publication points.

What should you do?

- ☐ Change the feed URL of the package
- ☒ Create a new feed for the package
- ☐ Promote the package to a release view.
- ☐ Publish the package to a public NuGet repository.

Explanation:

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.

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

Question 7

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 point.

- ☐ a deployment group
- ☐ a Microsoft-hosted agent
- ☐ service hooks
- ☒ a self-hosted agent
- ☒ an External Git service connection

Explanation:

Explanation:

Licensed to Berend Post berend.post@it-university.nl

E: GitLab CI/CD can be used with GitHub or any other Git server such as BitBucket. Instead of moving your entire project to GitLab, you can connect your external repository to get the benefits of GitLab CI/CD.

Note: When a pipeline uses a remote, 3rd-party repository host such as Bitbucket Cloud, the repository is configured with webhooks 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.

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

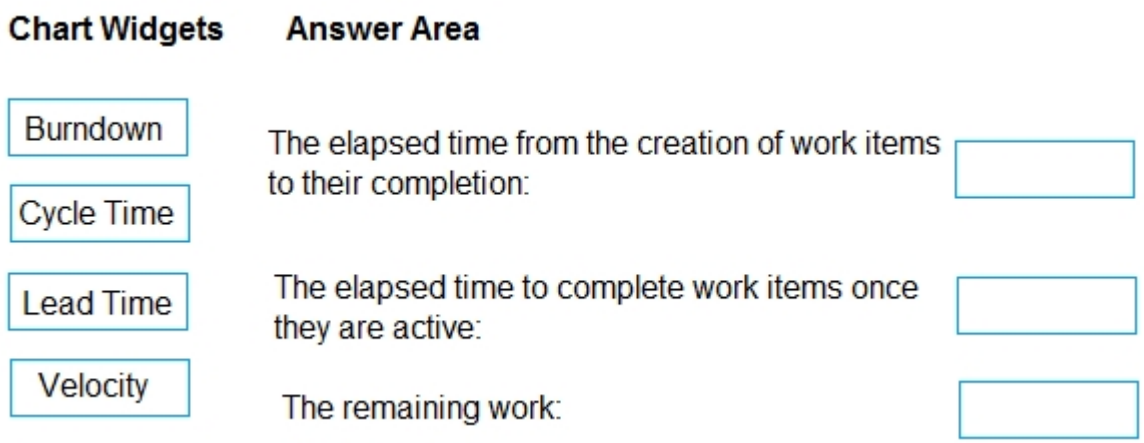
Question 8

DRAG DROP

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

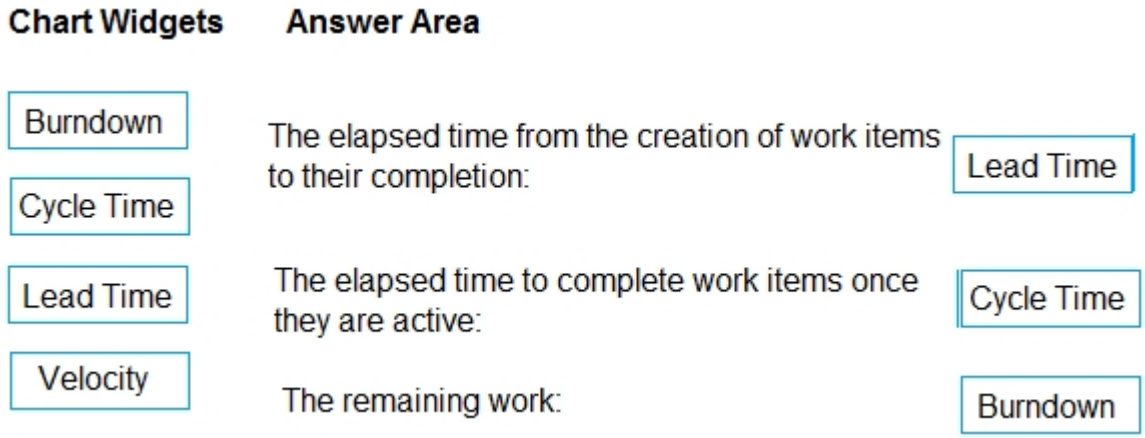
Which chart widgets should you recommend for each metric? To answer, drag the appropriate chart widgets to the correct metrics. Each chart widget may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.



Licensed to Berend Post berend.post@it-university.nl

Solution:



Explanation:

Explanation:

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: Burndown

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

Incorrect Answers:

Velocity provides a useful metric for these activities:

Support sprint planning

Forecast future sprints and the backlog items that can be completed

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

References:

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

Licensed to Berend Post berend.post@it-university.nl

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

Question 9

During a code review, you discover many quality issues. Many modules contain unused variables and empty catch blocks.

You need to recommend a solution to improve the quality of the code.

What should you recommend?

- ☐ In a Grunt build task, select Enabled from Control Options.
- ☒ In a Maven build task, select Run PMD.
- ☐ In a Xcode build task, select Use xcpretty from Advanced.
- ☐ In a Gradle build task, select Run Checkstyle.

Explanation:

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.

Incorrect Answers:

C: xcpretty is a fast and flexible formatter for xcodebuild.

References: https://pmd.github.io/

Licensed to Berend Post berend.post@it-university.nl

Implement DevOps Development Processes

Case Study (3 questions)

Case Study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study
To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Application Architecture

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

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

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

Architectural Issues

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

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

 Licensed to Berend Post berend.post@it-university.nl

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

Customers report that bug reporting is overly complex.

Planned changes

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

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

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

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

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

Technical requirements

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

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

Licensed to Berend Post berend.post@it-university.nl

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

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

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

Question 10

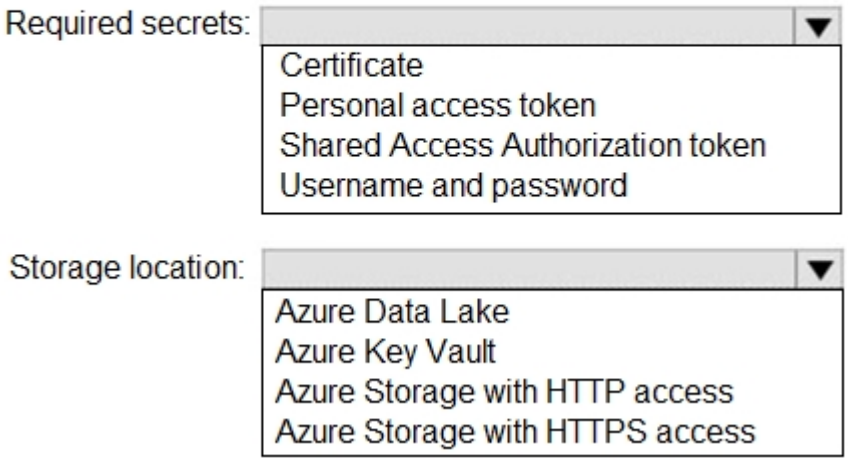
HOTSPOT

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

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

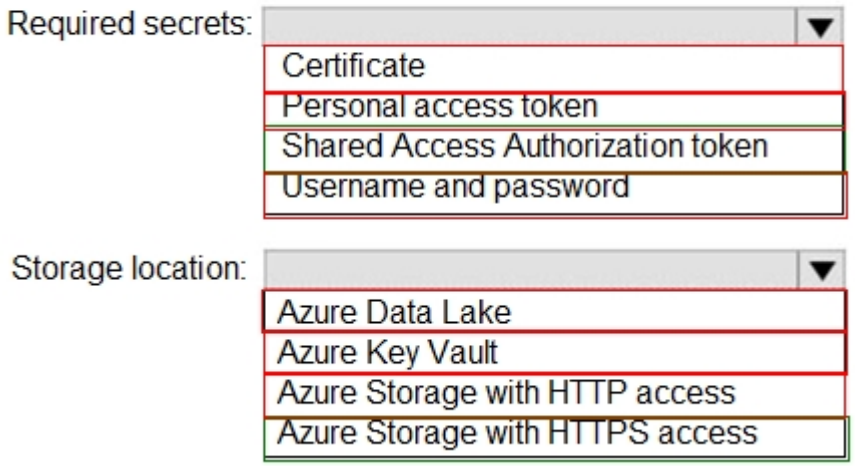
NOTE: Each correct selection is worth one point.

Answer Area



Licensed to Berend Post berend.post@it-university.nl

Answer Area



Solution:

Explanation:

Explanation:

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

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

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

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

Licensed to Berend Post berend.post@it-university.nl

Question 11

To resolve the current technical issue, what should you do to the Register-AzureRmAutomationDscNode command?

- ☒ Change the value of the ConfigurationMode parameter.
- ☐ Replace the Register-AzureRmAutomationDscNode cmdlet with
- ☐ Register-AzureRmAutomationScheduledRunbook
- ☐ Add the AllowModuleOverwrite parameter.
- ☐ Add the DefaultProfile parameter.

Explanation:

Explanation:

Change the ConfigurationMode parameter from ApplyOnly to ApplyAndAutocorrect.

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

Scenario: Current Technical Issue

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

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

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

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

Licensed to Berend Post berend.post@it-university.nl

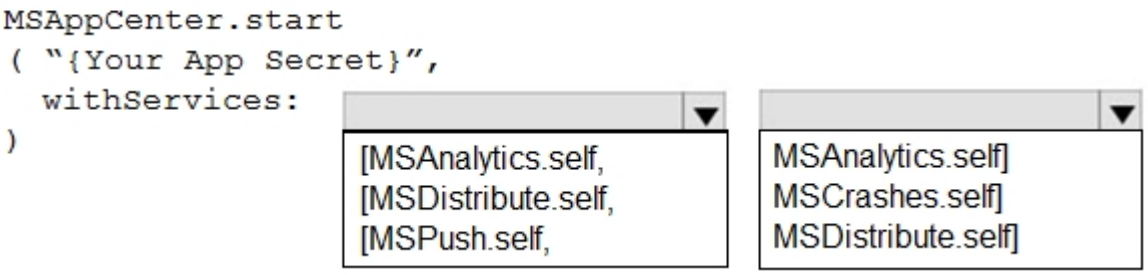
Question 12

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.

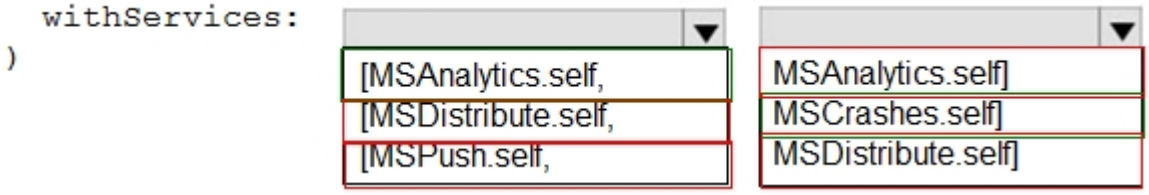
NOTE: Each correct selection is worth one point.

Answer Area



Solution:

Answer Area



Explanation:

Explanation:

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

In order to use App Center, you need to opt in to the service(s) that you want to use, meaning by default no services are started and you will have to explicitly call each of them when starting the SDK.

Licensed to Berend Post berend.post@it-university.nl

Insert the following line to start the SDK in your app's AppDelegate class in the didFinishLaunchingWithOptions method.

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

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

Licensed to Berend Post berend.post@it-university.nl

Case Study (3 questions)

Case Study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study
To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Background

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

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

The Azure DevOps organization includes:

The Docker extension
A deployment pool named Pool7 that contains 10 Azure virtual machines that run Windows Server 2016

The Azure subscription contains an Azure Automation account.

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

Licensed to Berend Post berend.post@it-university.nl

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

Technical requirements

Contoso identifies the following technical requirements:

Implement build agents for Project1.
Whenever possible, use Azure resources.
Avoid using deprecated technologies.
Implement a code flow strategy for Project2 that will:
-Enable Team2 to submit pull requests for Project2.
-Enable Team2 to work independently on changes to a copy of Project2.
-Ensure that any intermediary changes performed by Team2 on a copy of Project2 will be subject to the same restrictions as the ones defined in the build policy of Project2.
Whenever possible implement automation and minimize administrative effort.
Implement Project3, Project5, Project6, and Project7 based on the planned changes
Implement Project4 and configure the project to push Docker images to Azure Container Registry.

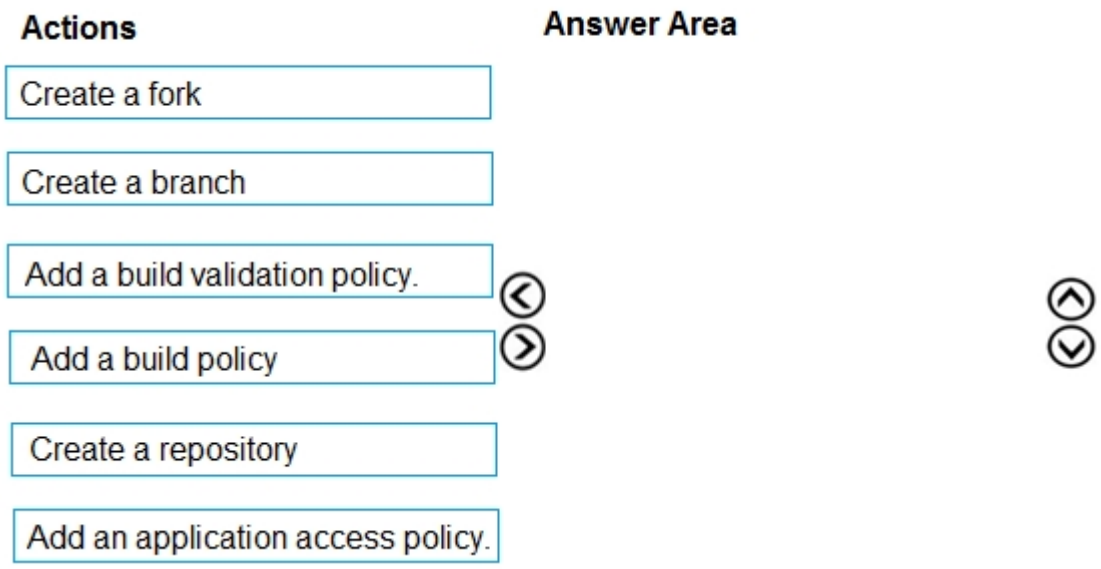
Question 13

DRAG DROP

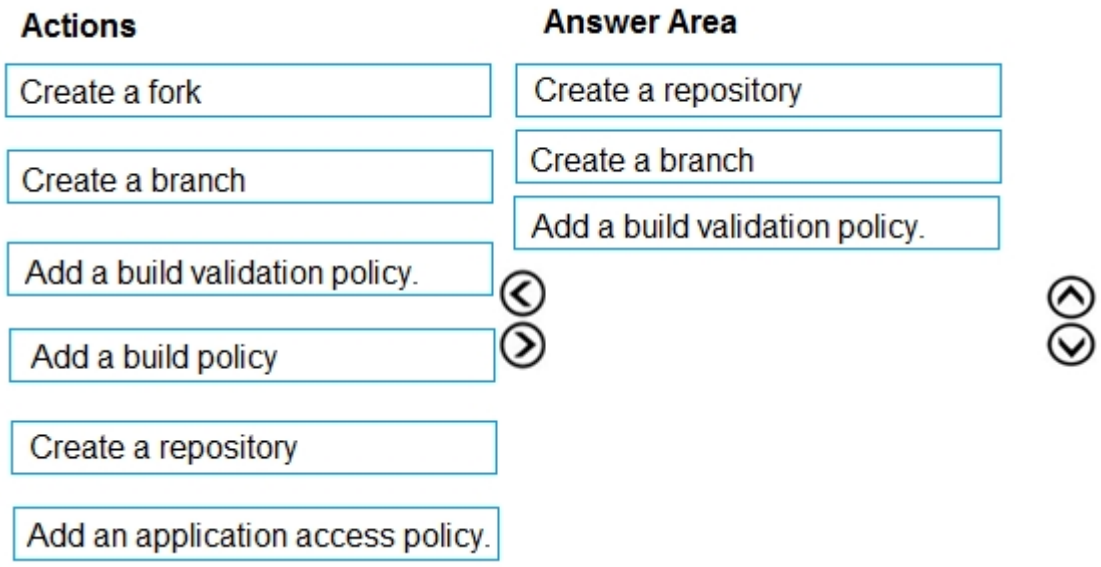
You need to implement the code flow strategy for Project2 in Azure DevOps.

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 them in the correct order.

Licensed to Berend Post berend.post@it-university.nl



Solution:



Explanation:

Explanation:

Step 1: Create a 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.

Licensed to Berend Post berend.post@it-university.nl

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.

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

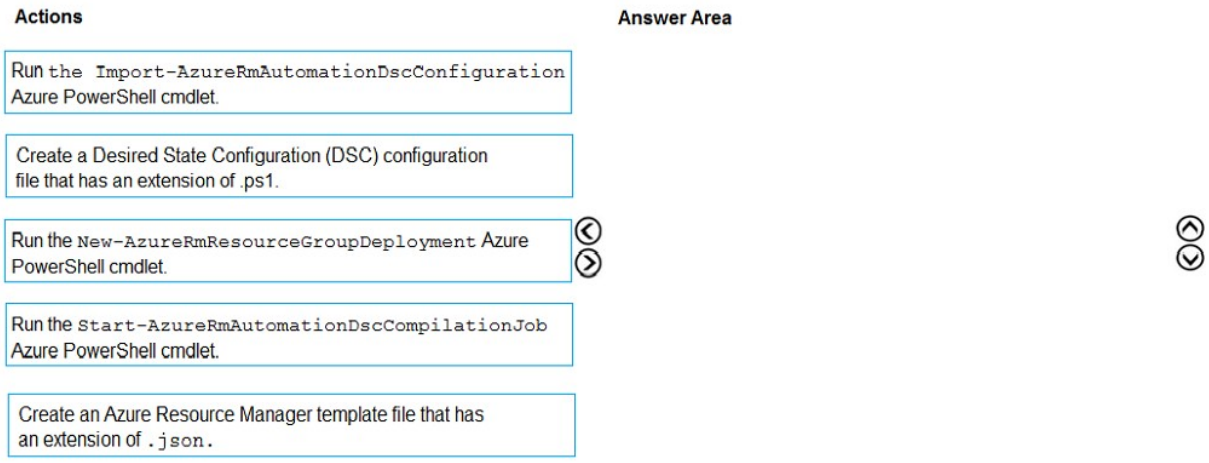
References: https://docs.microsoft.com/en-us/azure/devops/repos/git/manage-your-branches

Question 14

DRAG DROP

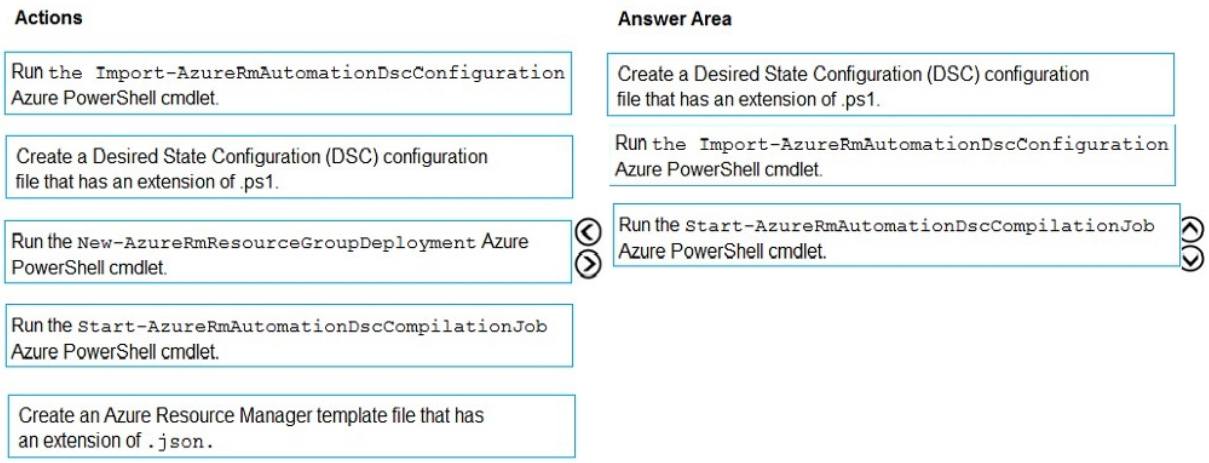
You need to configure Azure Automation for the computers in Group7.

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



Licensed to Berend Post berend.post@it-university.nl

Solution:



Explanation:

Explanation:

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

Step 2: Run the Import-AzureRmAutomationDscConfiguration Azure Powershell cmdlet

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

Example:

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

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

Step 3: Run the Start-AzureRmAutomationDscCompilationJob Azure Powershell cmdlet

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

References:

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

Licensed to Berend Post berend.post@it-university.nl

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

Question 15

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

You need to configure the computers in Group7.

What should you do?

- ☐ Modify the RefreshMode property of the Local Configuration Manager (LCM).
- ☒ Run the Register-AzureRmAutomationDscNode Azure Powershell cmdlet.
- ☐ Modify the ConfigurationMode property of the Local Configuration Manager (LCM)
- ☐ Install PowerShell Core.

Explanation:

Explanation:

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

Scenario: The Azure DevOps organization includes:

The Docker extension

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

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

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

Licensed to Berend Post berend.post@it-university.nl

(11 questions)

Question 16

Your company builds a multi-tier web application.

You use 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 will use to test new features.

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

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

What should you do?

- ☐ In Azure Cloud Shell, run Azure CLI commands to create and delete the new virtual machines in a staging resource group.
- ☐ In Azure DevOps, configure new tasks in the release pipeline to deploy to Azure Cloud Services.
- ☐ From Azure Cloud Shell, run Azure PowerShell commands to create and delete the new virtual machines in a staging resource group.
- ☒ In Azure DevOps, configure new tasks in the release pipeline to create and delete the virtual machines in Azure DevTest Labs.

Explanation:

Explanation:

You can use the Azure DevTest Labs Tasks extension 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.

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

Licensed to Berend Post berend.post@it-university.nl

Question 17

Your company uses cloud-hosted Jenkins for builds.

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

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

NOTE: Each correct selection is worth one point.

- ☐ Create a webhook in Jenkins.
- ☒ Add the Team Foundation Server (TFS) plug-in to Jenkins.
- ☐ Add a domain to your Jenkins account.
- ☒ Create a personal access token in your Azure DevOps account.
- ☒ Create a service hook in Azure DevOps.

Explanation:

Explanation:

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

D: Use Azure DevOps/ Visual Studio Team Services to create an access token, and use th

E: For those who need tighter integration, Team Services provides two additional ways to achieve it: 1) the Jenkins Service Hook, and 2) Jenkins build and release tasks.)

References:

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

Licensed to Berend Post berend.post@it-university.nl

Question 18

DRAG DROP

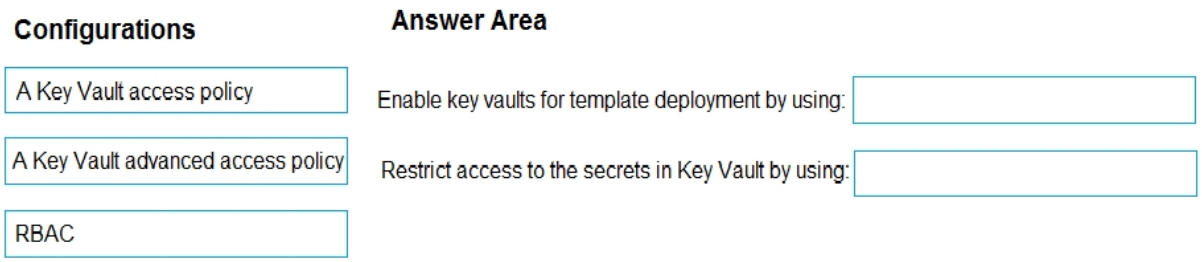
Your company has a project in Azure DevOps.

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.

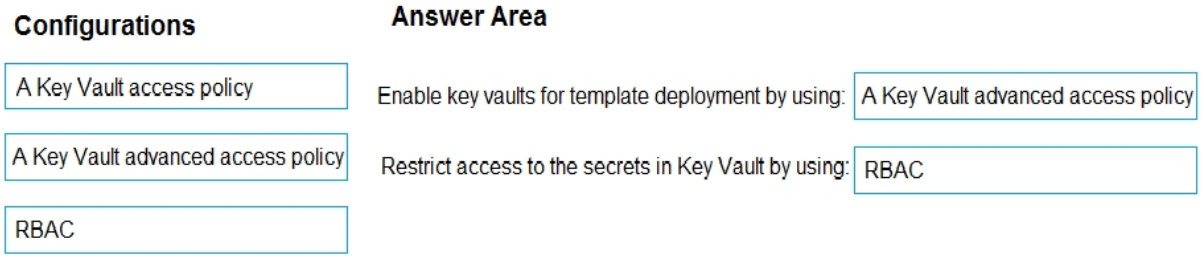
You need to recommend a solution for accessing the secrets stored in the key vault during deployments. The solution must use the principle of least privilege.

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.



Solution:

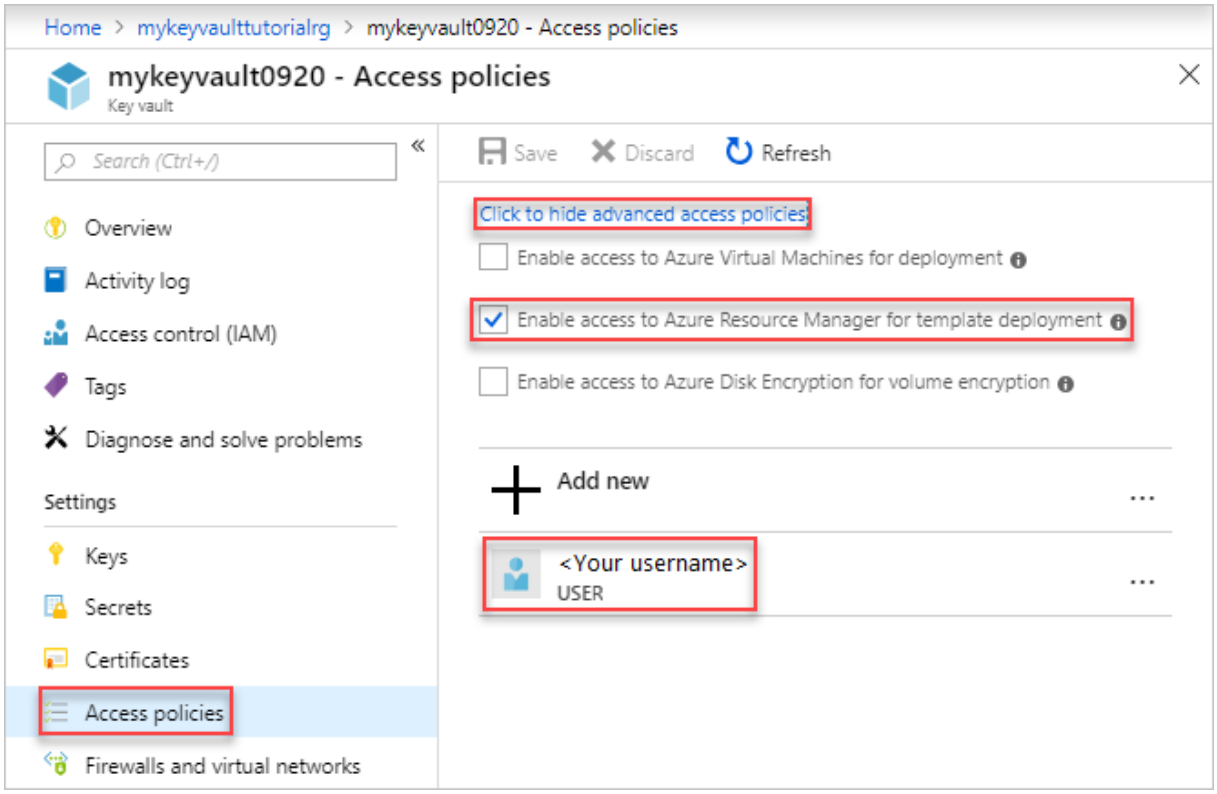


Explanation:

Explanation:

Box 1: A key Vault advanced access policy

Licensed to Berend Post berend.post@it-university.nl



Box 2: 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.

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

Licensed to Berend Post berend.post@it-university.nl

Question 19

You have a branch policy in a project in Azure DevOps. The policy requires that code always builds successfully.

You need to ensure that a specific user can always merge changes to the master branch, even if the code fails to compile. The solution must use the principle of least privilege.

What should you do?

- ☐ Add the user to the Build Administrators group.
- ☐ Add the user to the Project Administrators group.
- ☐ From the Security settings of the repository, modify the access control for the user.
- ☒ From the Security settings of the branch, modify the access control for the user.

Explanation:

Explanation:

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

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

Licensed to Berend Post berend.post@it-university.nl

Question 20

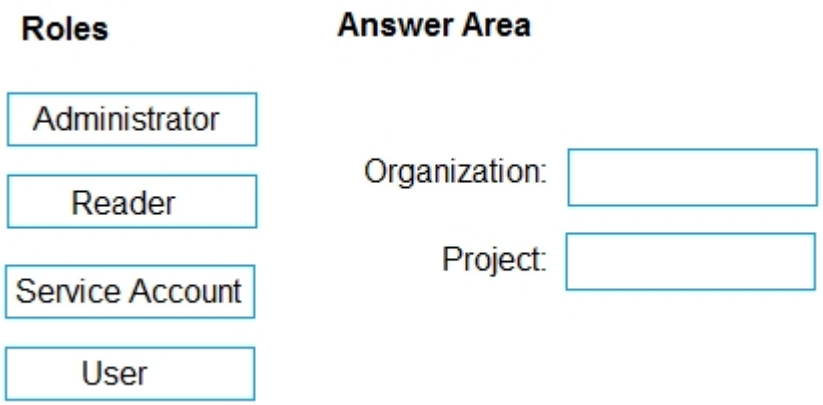
DRAG DROP

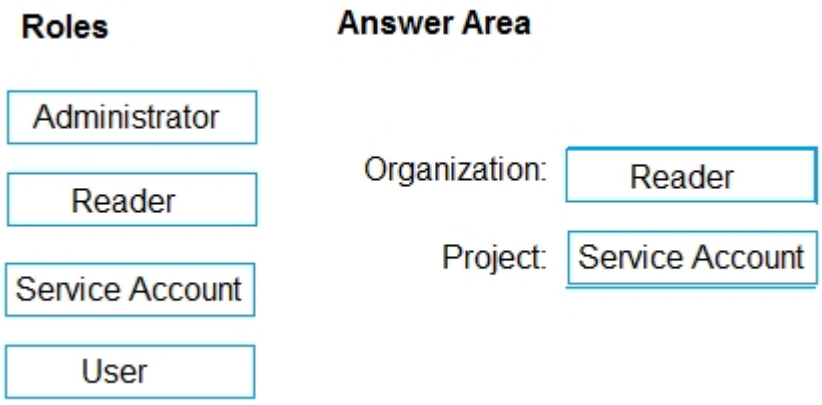
You need to configure access to Azure DevOps agent pools to meet the following requirements:

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.





Solution:

Explanation:

Explanation:

Box 1: Reader

Licensed to Berend Post berend.post@it-university.nl

Members of the Reader role can view the organization agent 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 agent pool 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.

Incorrect Answers:

In addition to all the permissions given the Reader and the Service Account role, members of the administrator role can register or unregister agents from the organization agent pool. They can also refer to the organization agent pool when creating a project agent pool in a project. Finally, they can also manage membership for all roles of the organization agent pool. The user that created the organization agent pool is automatically added to the Administrator role for that pool.

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

Question 21

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

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

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

Consolidate commit histories.
Merge the changes into a single commit.

Which merge strategy should you use in the branch policy?

- ☒ squash merge
- ☐ fast-forward merge
- ☐ Git fetch
- ☐ no-fast-forward merge

Explanation:

Explanation:

Licensed to Berend Post berend.post@it-university.nl

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 branch being added to the history of the default branch, a squash merge takes all the file changes and adds them to a single new commit on the default branch.

A simple way to think about this is that squash merge gives you just the file changes, and a regular merge gives you the file changes and the commit history.

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

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

Question 22

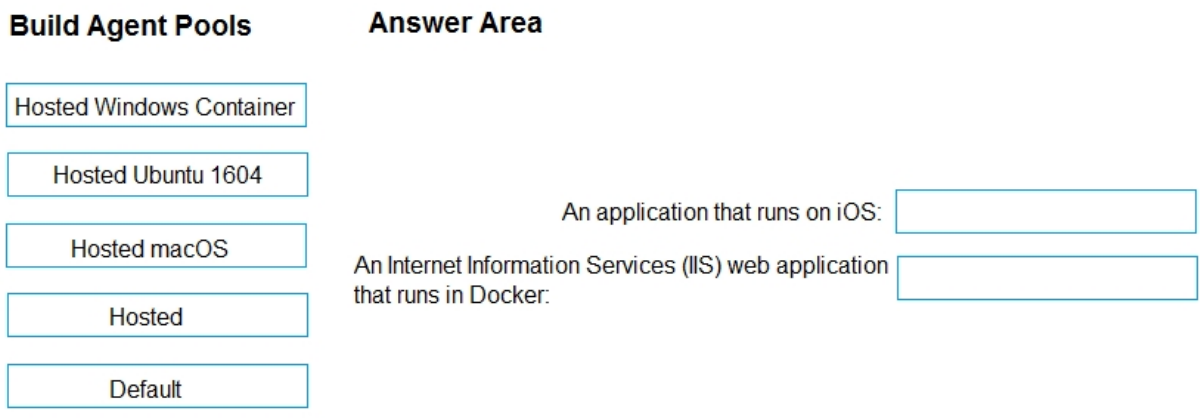
DRAG DROP

You are configuring Azure DevOps build pipelines.

You plan to use hosted build agents.

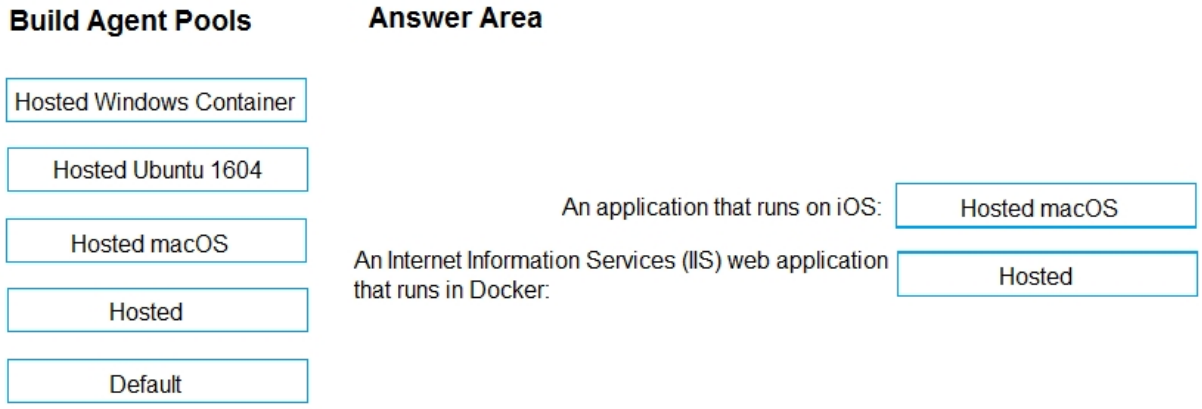
Which build agent pool should you use to compile each application type? To answer, drag the appropriate build agent pools to the correct application types. Each build agent pool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.



Licensed to Berend Post berend.post@it-university.nl

Solution:



Explanation:

Explanation:

Box 1: Hosted macOS

Hosted macOS 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.

Incorrect Answers:

Default pool: Use it to register self-hosted agents that you've set up.

Hosted Windows Container pool (Azure Pipelines only): Enabled you to build and release inside Windows containers. Unless you're building using containers, Windows builds should run in the Hosted VS2017 or Hosted pools.

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

Licensed to Berend Post berend.post@it-university.nl

Question 23

DRAG DROP

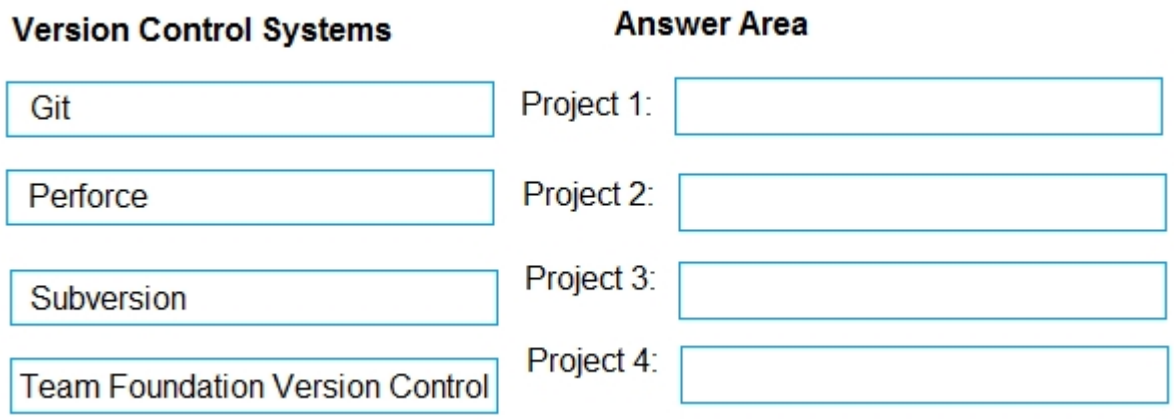
Your company has four projects. The version control requirements for each project are shown in the following table.

Project	Requirement
Project 1	Project leads must be able to restrict access to individual files and folders in the repository.
Project 2	The version control system must enforce the following rules before merging any changes to the main branch: • 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.

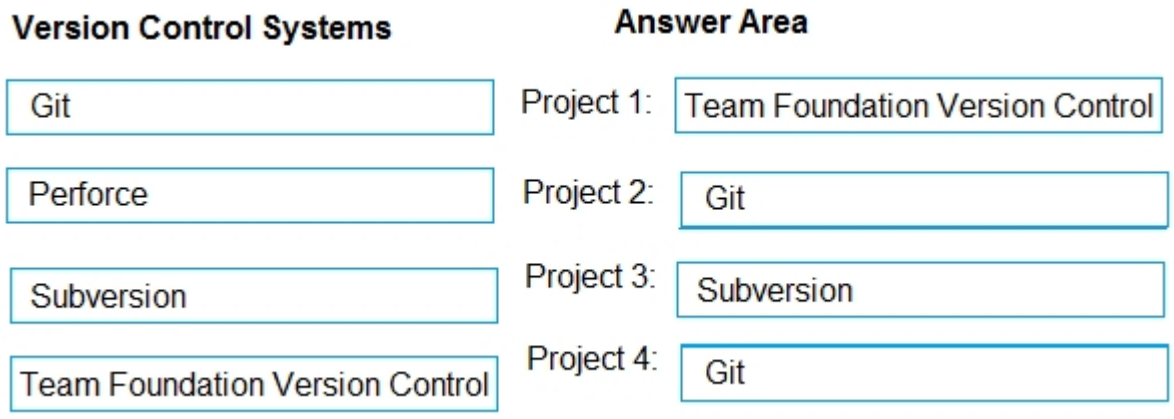
Which version control system should you use for each project? To answer, drag the appropriate version control systems to the correct projects. Each version control system may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.



Licensed to Berend Post berend.post@it-university.nl

Solution:



Explanation:

Explanation:

Box 1: Team Foundation Version Control

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

Box 2: Git

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

Box 3: Subversion

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

Box 4: Git

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

References:

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

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

Licensed to Berend Post berend.post@it-university.nl

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

Question 24

HOTSPOT

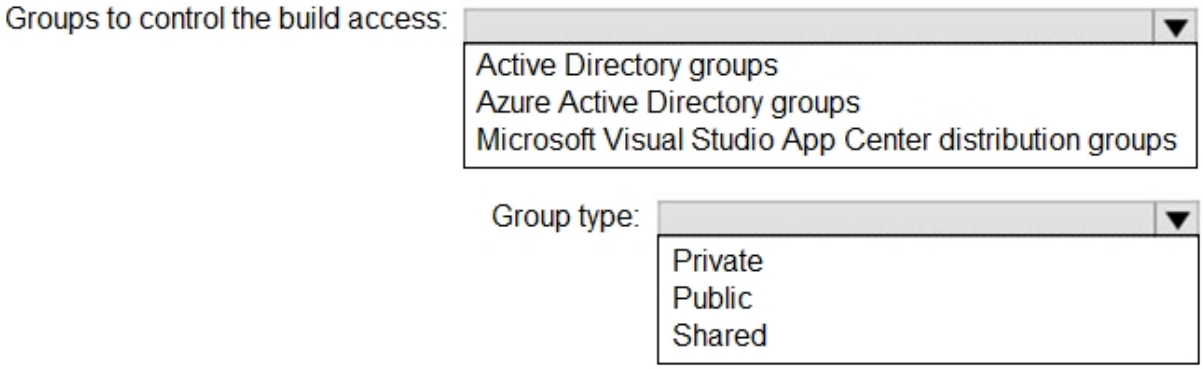
Your company is creating a suite of three mobile applications.

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

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

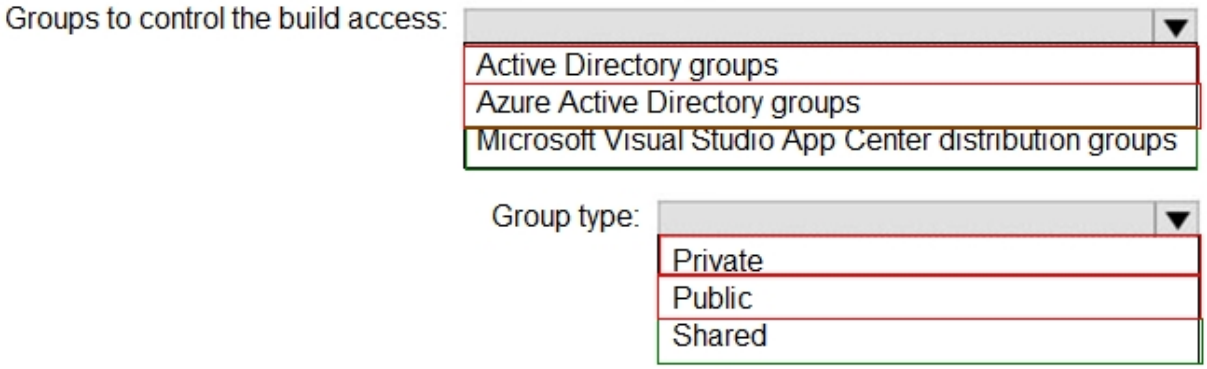
NOTE: Each correct selection is worth one point.

Answer Area



Solution:

Answer Area



Licensed to Berend Post berend.post@it-university.nl

Explanation:

Explanation:

Box 1: Microsoft Visual Studio App Center distribution Groups

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

Box 2: Shared

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

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

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

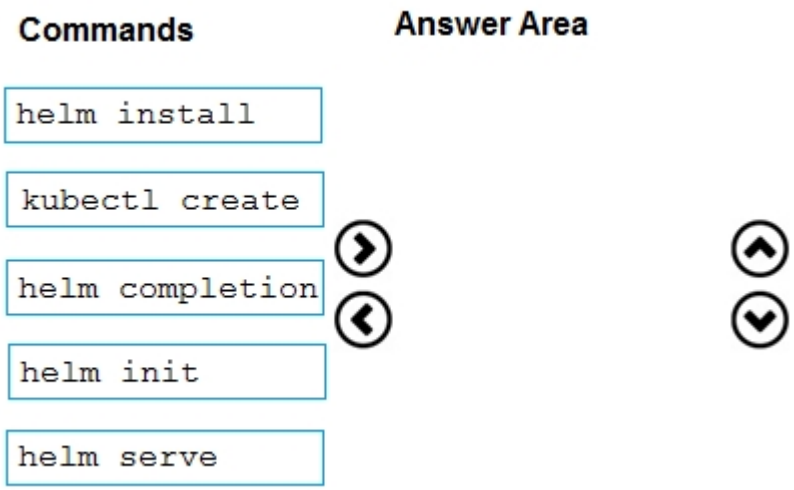
Licensed to Berend Post berend.post@it-university.nl

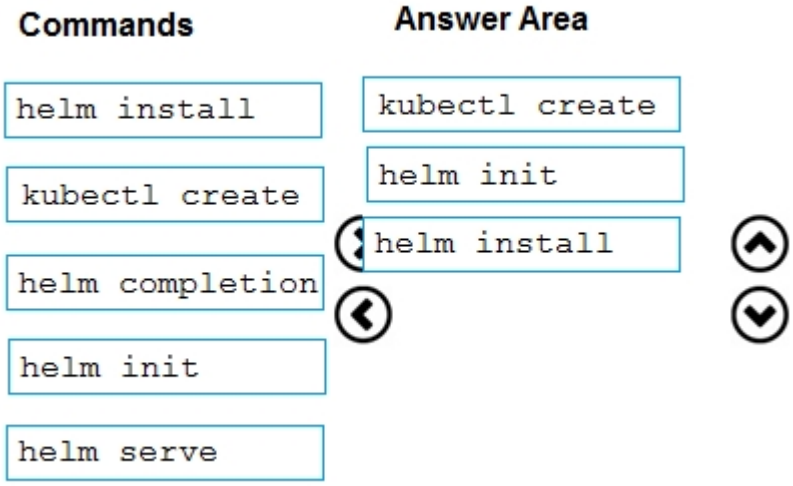
Question 25

DRAG DROP

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

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





Solution:

Explanation:

Explanation:

Step 1: Kubectl create

Licensed to Berend Post berend.post@it-university.nl

You can add a service account to Tiller using the --service-account <NAME> flag while you're configuring Helm (step 2 below). As a prerequisite, you'll have to create a role binding which specifies a role and a service account name that have been set up in advance.

Example: Service account with cluster-admin role

$ kubectl create -f rbac-config.yaml

serviceaccount "tiller" created

clusterrolebinding "tiller" created

$ helm init --service-account tiller

Step 2: helm init

To deploy a basic Tiller into an AKS cluster, use the helm init command.

Step 3: helm install

To install charts with Helm, use the helm install command and specify the name of the chart to install.

References:

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

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

Question 26

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?

- ☒ Azure Key Vault
- ☐ a Web.config file
- ☐ an Appsettings.json file
- ☐ an Azure Storage table
- ☐ an Azure Resource Manager parameter file

Explanation:

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

Licensed to Berend Post berend.post@it-university.nl

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.

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

Licensed to Berend Post berend.post@it-university.nl

Implement Continuous Integration

Case Study (2 questions)

Case Study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study
To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Background

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

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

The Azure DevOps organization includes:

The Docker extension
A deployment pool named Pool7 that contains 10 Azure virtual machines that run Windows Server 2016

The Azure subscription contains an Azure Automation account.

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

Licensed to Berend Post berend.post@it-university.nl

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

Technical requirements

Contoso identifies the following technical requirements:

Implement build agents for Project1.
Whenever possible, use Azure resources.
Avoid using deprecated technologies.
Implement a code flow strategy for Project2 that will:
-Enable Team2 to submit pull requests for Project2.
-Enable Team2 to work independently on changes to a copy of Project2.
-Ensure that any intermediary changes performed by Team2 on a copy of Project2 will be subject to the same restrictions as the ones defined in the build policy of Project2.
Whenever possible implement automation and minimize administrative effort.
Implement Project3, Project5, Project6, and Project7 based on the planned changes
Implement Project4 and configure the project to push Docker images to Azure Container Registry.

Question 27

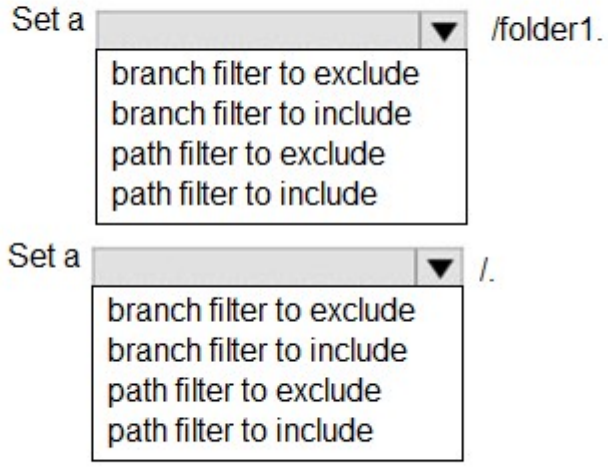
HOTSPOT

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

NOTE: Each correct selection is worth one point.

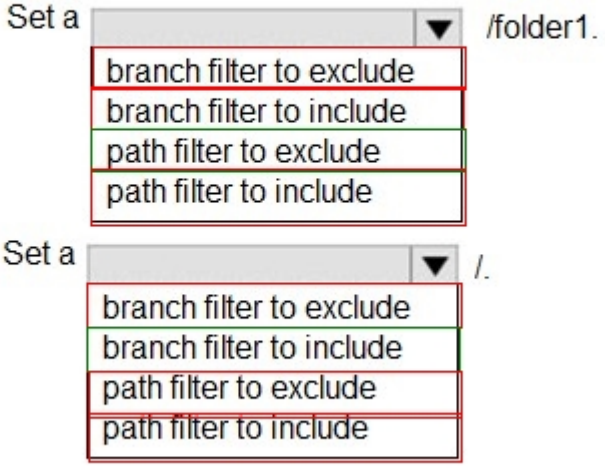
Licensed to Berend Post berend.post@it-university.nl

Answer Area



@

Answer Area



@

Solution:

Explanation:

Explanation:

Scenario:

Project5 will contain a Git repository in Azure Reports and a continuous integration trigger that will initiate a build in response to any change except for changes within /folder1 of the repository.

Licensed to Berend Post berend.post@it-university.nl

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

Question 28

In Azure DevOps, you create Project3.

You need to meet the requirements of the project.

What should you do first?

- ☐ From Azure DevOps, modify the build definition.
- ☐ From SonarQube, obtain an authentication token.
- ☒ From Azure DevOps, create a service endpoint.
- ☐ From SonarQube, create a project.

Explanation:

Explanation:

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

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

(7 questions)

Question 29

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

You create a new public project in Azure DevOps.

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

Which authentication type should you use?

- ☒ OAuth
- ☐ GitHub App
- ☐ a personal access token
- ☐ SAML

Explanation:

Explanation:

You can authenticate as a GitHub App.

Licensed to Berend Post berend.post@it-university.nl

References: https://developer.github.com/apps/building-github-apps/authenticating-with-github-apps/

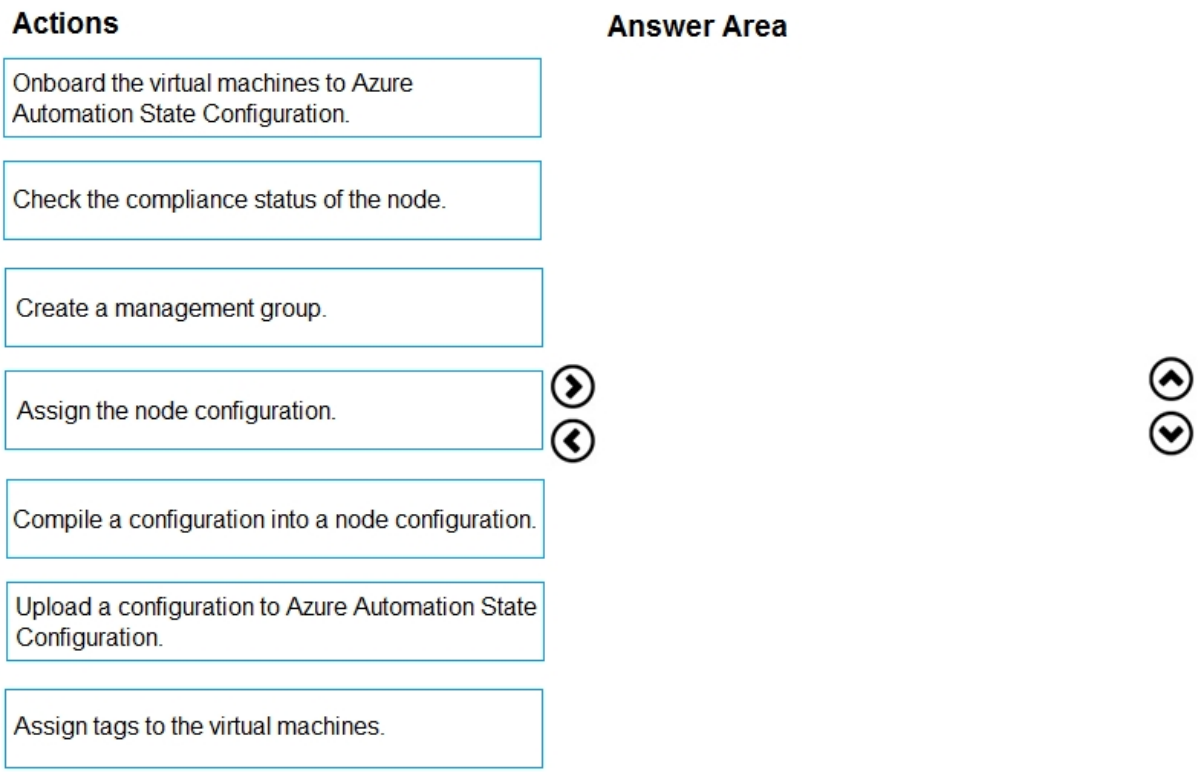
Question 30

DRAG DROP

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

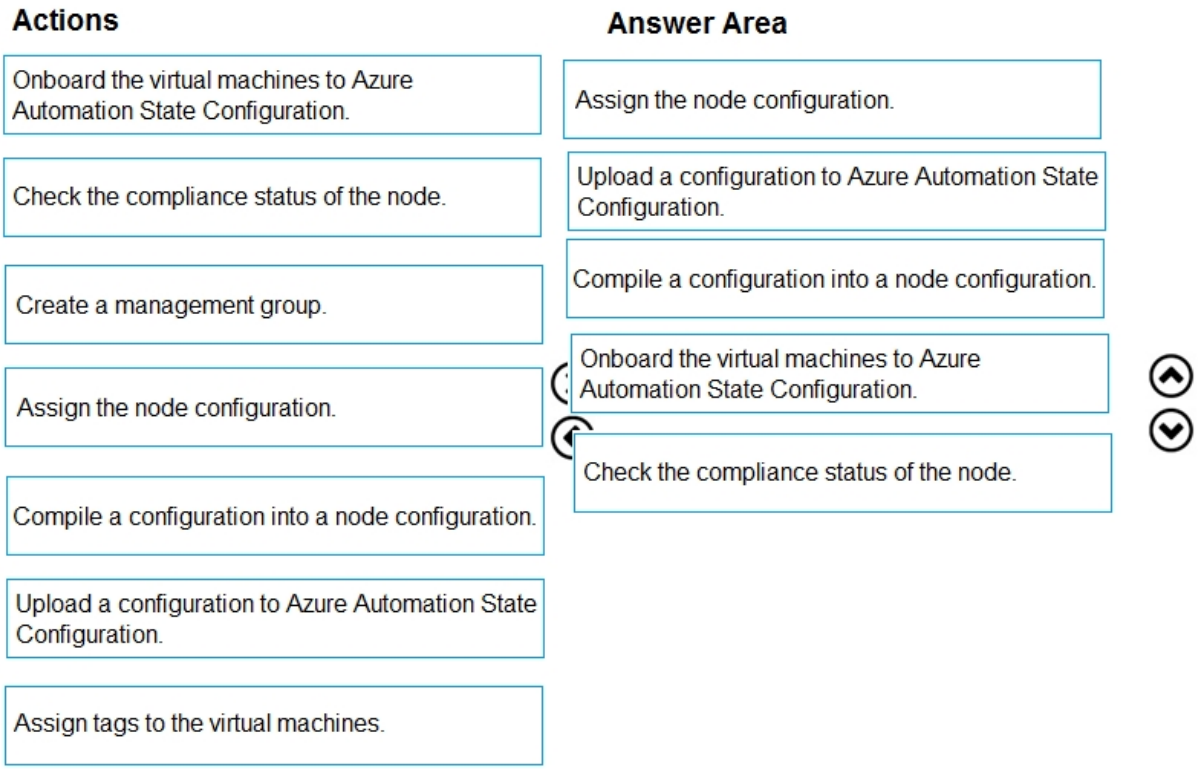
Which five actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.



Licensed to Berend Post berend.post@it-university.nl

Solution:



Explanation:

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

Licensed to Berend Post berend.post@it-university.nl

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).

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

Question 31

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

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

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

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

Does this meet the goal?

- ☐ Yes
- ☒ No

Explanation:

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

Licensed to Berend Post berend.post@it-university.nl

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

Question 32

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

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

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

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

Does this meet the goal?

- ☒ Yes
- ☐ No

Explanation:

Explanation:

On the Triggers tab you specify the events that will trigger the build. You can use the same build pipeline for both CI and scheduled builds.

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

Question 33

You have 50 Node.js-based projects that you scan by using WhiteSource. Each project includes Package.json, Package-lock.json, and Npm-shrinkwrap.json files.

You need to minimize the number of libraries reports by WhiteSource to only the libraries that you explicitly reference.

What should you do?

- ☐ Configure the File System Agent plug-in.
- ☒ Add a devDependencies section to Package-lock.json.
- ☐ Configure the Artifactory plug-in.
- ☐ Delete Package-lock.json.

Explanation:

Explanation:

Licensed to Berend Post berend.post@it-university.nl

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.

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

Question 34

Your company 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.

- ☐ Configure post-deployment approvals in the deployment pipeline.
- ☒ Configure pre-deployment approvals in the deployment pipeline.
- ☒ Integrate Azure DevOps and SonarQube.
- ☐ Integrate Azure DevOps and Azure DevTest Labs.

Licensed to Berend Post berend.post@it-university.nl

Question 35

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?

- ☒ 1
- ☐ 3
- ☐ 4
- ☐ 7

Explanation:

Explanation:

Multi-stage builds (in Docker 17.05 or higher) allow you to drastically reduce the size of your final image, without struggling to reduce the number of intermediate layers and files.

With multi-stage builds, you use multiple FROM statements in your Dockerfile. Each FROM instruction can use a different base, and each of them begins a new stage of the build. You can selectively copy artifacts from one stage to another, leaving behind everything you don't want in the final image.

References: https://docs.docker.com/develop/develop-images/multistage-build/#use-multi-stage-builds

Licensed to Berend Post berend.post@it-university.nl