Question: 114

You develop and deploy an Azure App Service web app named App1. You create a new Azure Key Vault named Vault 1. You import several API keys, passwords, certificates, and cryptographic keys into Vault1.

You need to grant App1 access to Vault1 and automatically rotate credentials Credentials must not be stored in code.

What should you do?
A. Enable App Service authentication for Appt. Assign a custom RBAC role to Vault1.
B. Add a TLS/SSL binding to App1.
C. Assign a managed identity to App1.
D. Upload a self-signed client certificate to Vault1. Update App1 to use the client certificate.

Answer: D

Question: 115

DRAG DROP

You need to ensure disaster recovery requirements are met.

What code should you add at line PC16? To answer, drag the appropriate code fragments to the correct locations. Each code fragment 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.

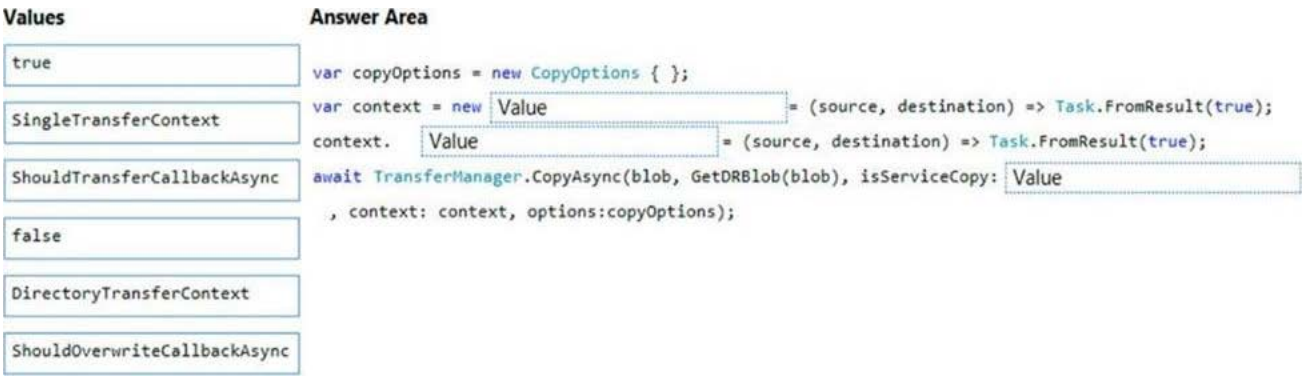

Answer:

Explanation:

Scenario: Disaster recovery. Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Box 1: DirectoryTransferContext

We transfer all files in the directory.

Note: The TransferContext object comes in two forms: SingleTransferContext and DirectoryTransferContext. The former is for transferring a single file and the latter is for transferring a directory of files.

Box 2: ShouldTransferCallbackAsync

The DirectoryTransferContext.ShouldTransferCallbackAsync delegate callback is invoked to tell whether a transfer should be done.

Box 3: False

If you want to use the retry policy in Copy, and want the copy can be resume if break in the middle, you can use SyncCopy (isServiceCopy = false).

Note that if you choose to use service side copy ('isServiceCopy' set to true), Azure (currently) doesn't provide SLA for that. Setting 'isServiceCopy' to false will download the source blob loca

Question: 116

You provide an Azure API Management managed web service lo clients. The back end web service implements HTTP Strict Transport Security (HSTS).

Every request to the backend service must include a valid HTTP authorization header.

You need to configure the Azure API Management instance with an authentication policy.

Which two policies can you uses? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.
A. Certificate Authentication
B. Basic Authentication
C. OAuth Client Credential Grant
D. Digest Authentication

Answer: C,D

Question: 117

HOTSPOT

You need to configure security and compliance for the corporate website files.

Which Azure Blob storage settings should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

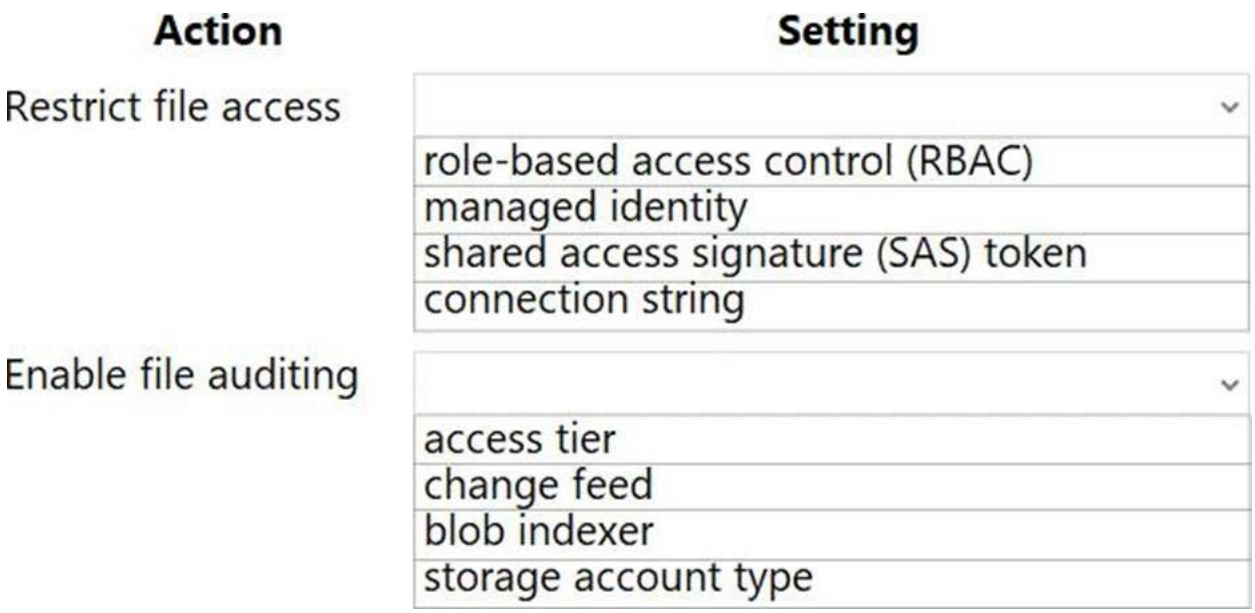

Answer:

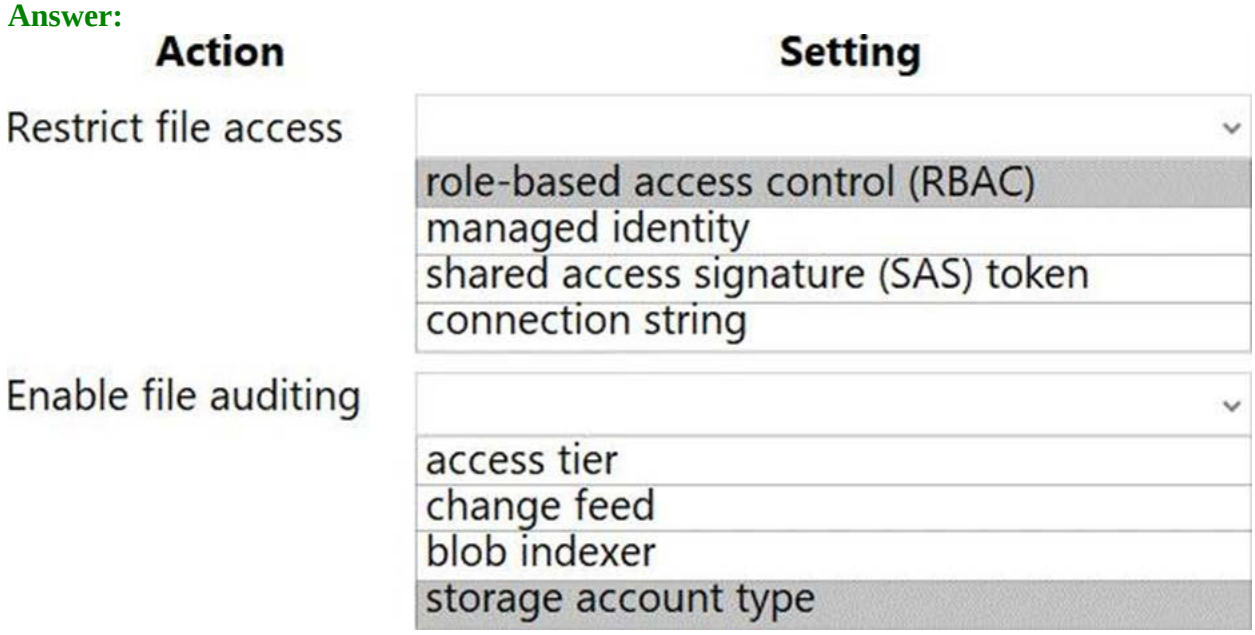

Explanation:

Box 1: role-based access control (RBAC)

Azure Storage supports authentication and authorization with Azure AD for the Blob and Queue services via Azure role-based access control (Azure RBAC).

Scenario: File access must restrict access by IP, protocol, and Azure AD rights.

Box 2: storage account type

Scenario: The website uses files stored in Azure Storage

Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR).

Creating a diagnostic setting:

Question: 118

You are developing a complex workflow by using Azure Durable Functions.

During testing you observe that the results of the workflow differ based on how many instances of the Azure Function are running.

You need to resolve the issue.

What should you do?
A. Ensure that all Orchestrator code is deterministic.
B. Read all state data from the durable function context
C. Configure the Azure Our able f unction to run on an App Service Plan with one instance.
D. Implement the monitor pattern within the workflow.

Answer: A

Question: 119

HOTSPOT

You are creating a CLI script that creates an Azure web app related services in Azure App Service.

The web app uses the following variables:

Variable name	Value
$gitrepo	https://github.com/Contos/webapp
$webappname	Webapp1103

You need to automatically deploy code from GitHub to the newly created web app.

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

```
az group create - -location westeurope - -name myResourceGroup
```

[▼] `--name $webappname - -resource-group myResourceGroup - -sku FREE`

- az webapp create
- az appservice plan create
- az webapp deployment
- az group delete

[▼] `--name $webappname - -resource-group myResourceGroup`

- az webapp create
- az appservice plan create
- az webapp deployment
- az group delete

[▼]

- --repo-url $gitrepo - -branch master - -manual-integration
- git clone $gitrepo
- --plan $webappname

[▼] `source config --name $webappname`

- az webapp create
- az appservice plan create
- az webapp deployment
- az group delete

`--resource-group myResourceGroup` [▼]

- --repo-url $gitrepo - -branch master - -manual-integration
- git clone $gitrepo
- --plan $webappname

Answer:

Answer Area

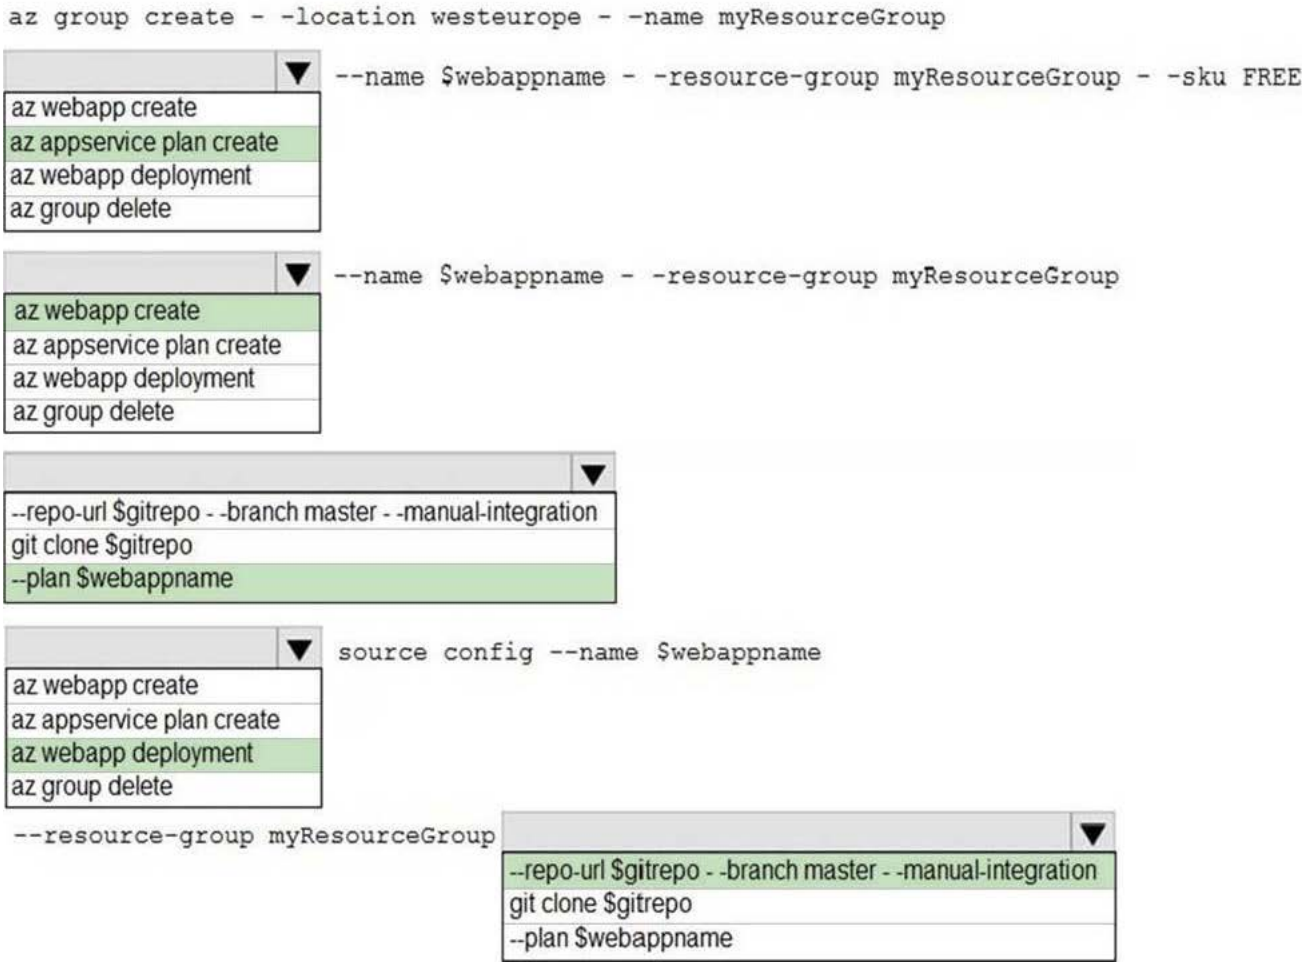

Explanation:

Box 1: az appservice plan create

The azure group creates command successfully returns JSON result. Now we can use resource group to create a azure app service plan

Box 2: az webapp create

Create a new web app..

Box 3: –plan $webappname

with the serviceplan we created in step 1.

Box 4: az webapp deployment

Continuous Delivery with GitHub. Example:

az webapp deployment source config –name firstsamplewebsite1 –resource-group websites–repo-url $gitrepo –branch master –git-token $token

Box 5: –repo-url $gitrepo –branch master –manual-integration

Question: 120

DRAG DROP

You have a web app named MainApp. You are developing a triggered App Service background task by using the WebJobs SDK. This task automatically invokes a function code whenever any new data is received in a queue.

You need to configure the services.

Which service should you use for each scenario? To answer, drag the appropriate services to the correct scenarios. Each service 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.

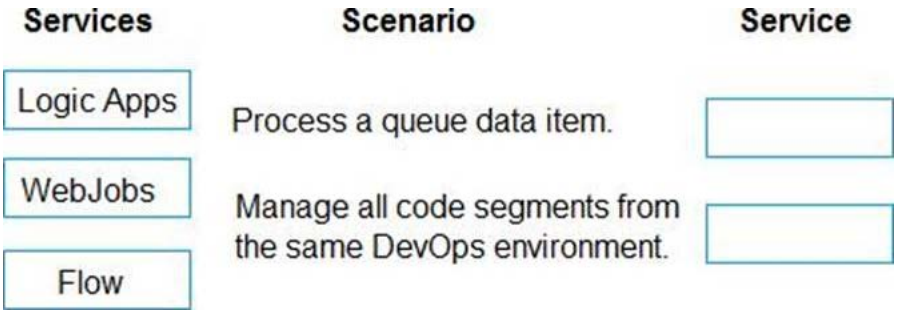

Answer:

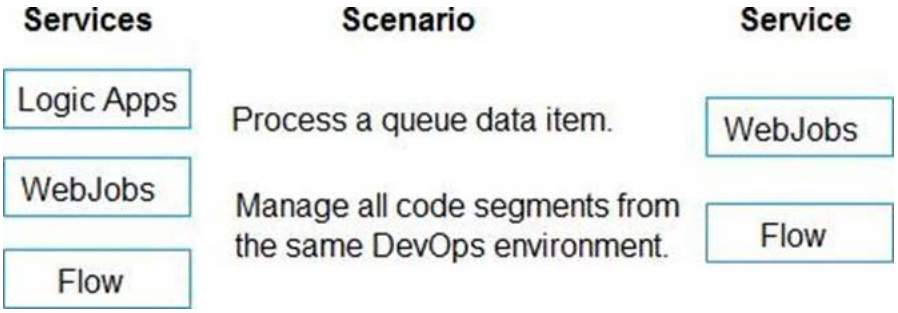

Explanation:

Box 1: WebJobs

A WebJob is a simple way to set up a background job, which can process continuously or on a schedule. WebJobs differ from a cloud service as it gives you get less fine-grained control over your processing environment, making it a more true PaaS service.

Box 2: Flow

Question: 121

You need to ensure that the solution can meet the scaling requirements for Policy Service.

Which Azure Application Insights data model should you use?
A. an Application Insights dependency
B. an Application Insights event
C. an Application Insights trace
D. an Application Insights metric

Answer: D

Explanation:

Application Insights provides three additional data types for custom telemetry:

Trace – used either directly, or through an adapter to implement diagnostics logging using an instrumentation framework that is familiar to you, such as Log4Net or System.Diagnostics.

Event – typically used to capture user interaction with your service, to analyze usage patterns.

Metric – used to report periodic scalar measurements.

Scenario:

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Reference: https://docs.microsoft.com/en-us/azure/azure-monitor/app/data-model

Question: 122

Topic 7, VanArsdel. Ltd

Case study

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

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

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

To start the case study

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

Background

VanArsdel. Ltd. is a global office supply company. The company fs based in Canada and has retail store locations

across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment

Requirements

The application components must meet the following requirements:

Corporate website

• Secure the website by using SSL

• Minimize costs tor data storage and hosting.

• Implement native GitHub workflows for continuous integration and continuous deployment (Cl/CO).

• Distribute the website content globally for local use.

• Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.

• The website must have 99.95 percent uptime.

Corporate website

The company provides a public website located at htlp://www. vanaisdelttd.com. The website consists of a React JavaScript user interface, HTML,CSS, image assets, and several APIs hosted in Azure functions.

Retail store locations

• Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos D3 by using native SQL language queries.

• Audit store sale transaction information nightly to validate data, process sates financials, and reconcile inventory.

Delivery services

• Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id. the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.

• Store delivery driver profile information in Azure Active Directory Azure AD) by using an Azure Function called from the corporate website.

Inventory services

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security

• All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.

• Authentication and authorization must use Azure AD and services must use managed identities where possible.

Retail Store Locations

• You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.

• Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

You need to secure the Azure Functions to meet the security requirements.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.
A. Store the RSA-HSM key in Azure Cosmos D
B. Apery the built-in policies for customer-managed keys and allowed locations.
C. Create a free tier Azure App Configuration instance with a new Azure AD service principal.
D. Store the RSA-HSM key in Azure Key Vault with soft-delete and purge-protection features enabled.
E. Store the RSA-HSM key in Azure Blob storage with an Immutability policy applied to the container.
F. Create a standard tier Azure App Configuration instance with an assigned Azure AD managed identity.

Answer: C,E

Explanation:

Scenario: All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.

Microsoft Azure Key Vault is a cloud-hosted management service that allows users to encrypt keys and small secrets by using keys that are protected by hardware security modules (HSMs).

You need to create a managed identity for your application.

Reference: https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

Question: 123

HOTSPOT

You are developing an application that uses Azure Storage Queues.

You have the following code:

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse
(CloudConfigurationManager.GetSetting("StorageConnectionString"));
CloudQueueClient queueClient = storageAccount.CreateCloudQueueClient()

CloudQueue queue = queueClient.GetQueueReference("appqueue") ;
await queu.CreateIfNotExistsAsync() ;

CloudQueueMessage peekedMessage = await queue.PeekMessageAsync() ;
if (peekedMessage != null)
{
     Console.WriteLine("The peeked message is: {0}", peekedMessage.AsString);
}
CloudQueueMessage message = await queue.GetMessageAsync() ;
```

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

Statement	Yes	No
The code configures the lock duration for the queue.	○	○
The last message read remains in the queue after the code runs.	○	○
The storage queue remains in the storage account after the code runs.	○	○

Answer:

Answer Area

Statement	Yes	No
The code configures the lock duration for the queue.	○	○
The last message read remains in the queue after the code runs.	○	○
The storage queue remains in the storage account after the code runs.	○	○

Explanation:

Box 1: No

The QueueDescription.LockDuration property gets or sets the duration of a peek lock; that is, the amount of time that the message is locked for other receivers. The maximum value for LockDuration is 5 minutes; the default value is 1 minute.

Box 2: Yes

You can peek at the message in the front of a queue without removing it from the queue by calling the PeekMessage method.

Box 3: Yes

Question: 124

DRAG DROP

You are developing an application to securely transfer data between on-premises file systems and Azure Blob storage. The application stores keys, secrets, and certificates in Azure Key Vault. The application uses the Azure Key Vault APIs.

The application must allow recovery of an accidental deletion of the key vault or key vault objects. Key vault objects must be retained for 90 days after deletion.

You need to protect the key vault and key vault objects.

Which Azure Key Vault feature should you use? To answer, drag the appropriate features to the correct actions. 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.

Features

- Access policy
- Purge protection
- Soft delete
- Shared access signature

Answer Area

Action	Feature
Enable retention period and accidental deletion.	Feature
Enforce retention period and accidental deletion.	Feature

Answer:

Features

- Access policy
- Purge protection
- Soft delete
- Shared access signature

Answer Area

Action	Feature
Enable retention period and accidental deletion.	Soft delete
Enforce retention period and accidental deletion.	Purge protection

Explanation:

Box 1: Soft delete

When soft-delete is enabled, resources marked as deleted resources are retained for a specified period (90 days by default). The service further provides a mechanism for recovering the deleted object, essentially undoing the deletion.

Box 2: Purge protection

Purge protection is an optional Key Vault behavior and is not enabled by default. Purge protection can only be enabled once soft-delete is enabled.

When purge protection is on, a vault or an object in the deleted state cannot be purged until the retention period has passed. Soft-deleted vaults and objects can still be recovered, ensuring that the retention policy will be followed.

Question: 125

HOTSPOT

You need to add code at line AM09 to ensure that users can review content using ContentAnalysisService.

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

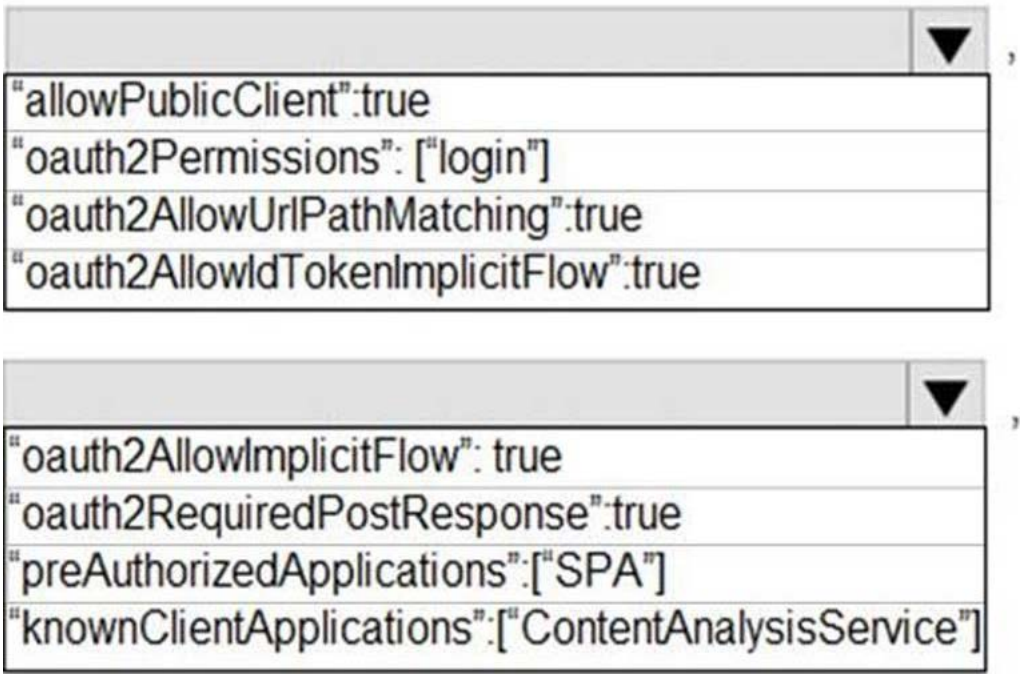

Answer:

Answer Area

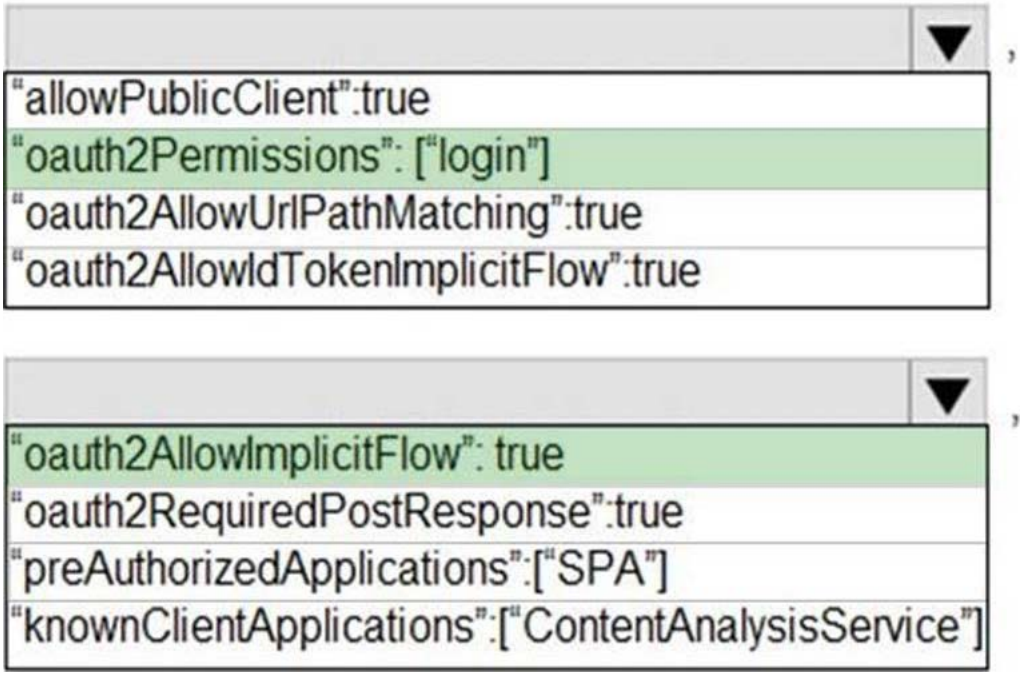

Explanation:

Box 1: "oauth2Permissions": ["login"]

oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

Box 2: "oauth2AllowImplicitFlow":true

For applications (Angular, Ember.js, React.js, and so on), Microsoft identity platform supports the OAuth 2.0 Implicit Grant flow.

Question: 126

HOTSPOT

You need to retrieve the database connection string.

Which values should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

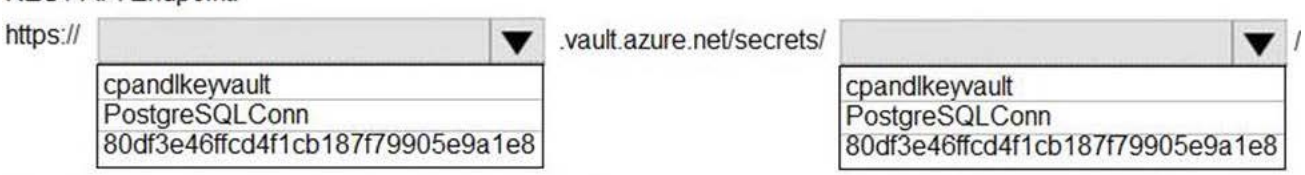

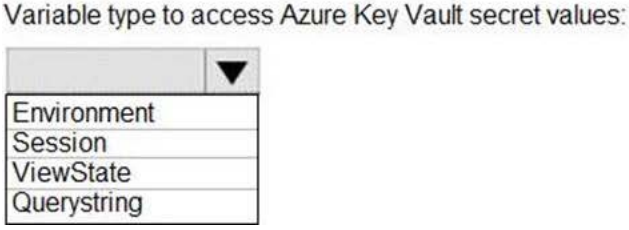

Answer:

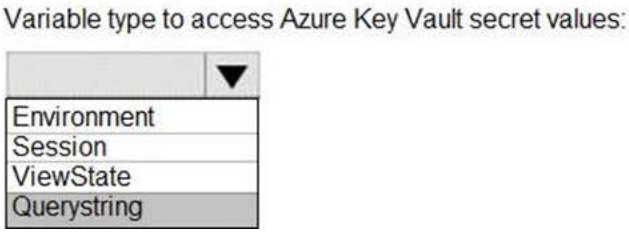

Explanation:

Azure database connection string retrieve REST API vault.azure.net/secrets/

Box 1: cpandlkeyvault

We specify the key vault, cpandlkeyvault.

Scenario: The database connection string is stored in Azure Key Vault with the following attributes:

Azure Key Vault name: cpandlkeyvault

Secret name: PostgreSQLConn

Id: 80df3e46ffcd4f1cb187f79905e9a1e8

Box 2: PostgreSQLConn

We specify the secret, PostgreSQLConn

Example, sample request:

https://myvault.vault.azure.net//secrets/mysecretname/4387e9f3d6e14c459867679a90fd0f79?api-version=7.1

Box 3: Querystring