PassQueen
a.k.a dumpsqueen.com

Implementing End-to-End Security Controls for Cloud and AI Workloads

Microsoft SC-500

Version Demo

Total Demo Questions: 8

Total Premium Questions: 85

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Implement security solutions for cloud and AI workloads	69
Topic 2, Manage security operations for cloud and AI workloads	15
Topic 3, Mix Questions	1
Total	85

QUESTION NO: 1 - (SIMULATION)

You have a Microsoft Entra tenant.

You need to implement password less authentication. The solution must meet the following requirements:

•Users can sign in without a password by using a mobile device.

•New users that sign in for the first time must use a helpdesk issued sign in method that expires.

Which authentication method should you enable for each requirement? To answer, drag the appropriate methods to the correct requirements. Each method 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.

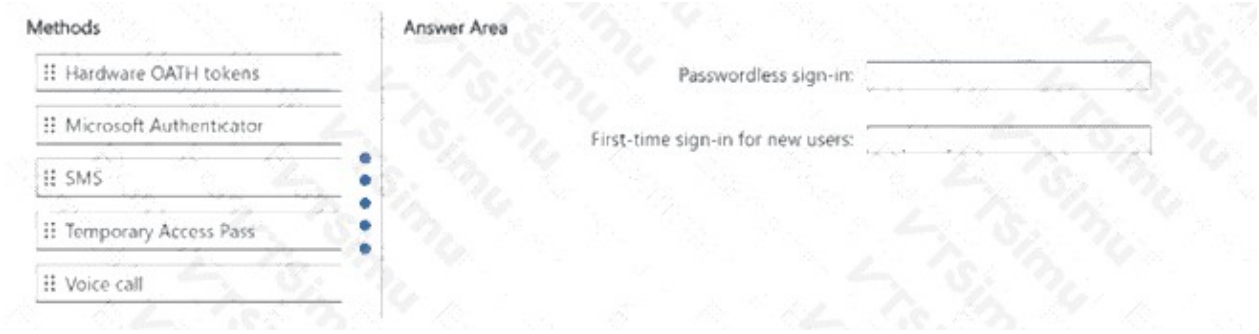

ANSWER: See the explanation for the answer

Explanation:

Configure the methods as follows:

Passwordless sign-in: `Microsoft Authenticator`
First-time sign-in for new users: `Temporary Access Pass`

Microsoft Authenticator supports passwordless phone sign-in from a mobile device. A Temporary Access Pass (TAP) is a time-limited code that can be issued by the helpdesk to let a new user complete initial sign-in and register passwordless methods.

QUESTION NO: 2

You have an Azure key vault named KV1 that currently uses the vault access policy permission model.

An Azure App Service web app named App1 uses a managed identity to retrieve secrets from KV1. You must migrate KV1 to Azure role-based access control (Azure RBAC) authorization. App1 must continue to retrieve secrets, but it must not be able to create, update, or delete secrets.

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

NOTE: Each correct selection is worth one point.

A. Change KV1 to use the Azure RBAC permission model.

B. Assign the Key Vault Secrets User role to the managed identity of App1.

C. Assign the Key Vault Secrets Officer role to the managed identity of App1.

D. Create a vault access policy that grants Get secret permission to App1.

E. Assign the Key Vault Contributor role to the managed identity of App1.

ANSWER: A B

Explanation:

Change KV1 to use the Azure RBAC permission model and assign the **Key Vault Secrets User** role to the managed identity of App1. The role grants permission to read secret contents without granting secret-management permissions, which satisfies the least-privilege requirement.

After the vault is switched to Azure RBAC authorization, vault access policies are no longer used for data-plane authorization. Key Vault Secrets Officer grants permissions to manage secrets and is broader than required. Key Vault Contributor is a management-plane role and does not grant permission to retrieve secret values.

QUESTION NO: 3 - (HOTSPOT)

You have an Azure subscription named Sub1 that contains 50 virtual machines. Sub1 has Microsoft Defender for Cloud enabled.

Sub1 contains an Azure key vault named KV1 and an Azure policy that enforces storing all secrets in KV1.

Occasionally, the developers at your company store plaintext tokens and SSH private keys on the virtual machines.

You need to configure Defender for Cloud to detect plaintext secrets on the virtual machines. The solution must minimize administrative changes to the virtual machines.

How should you configure Defender for Cloud? To answer, select the appropriate options in the answer area

NOTE: Each correct selection is worth one point.

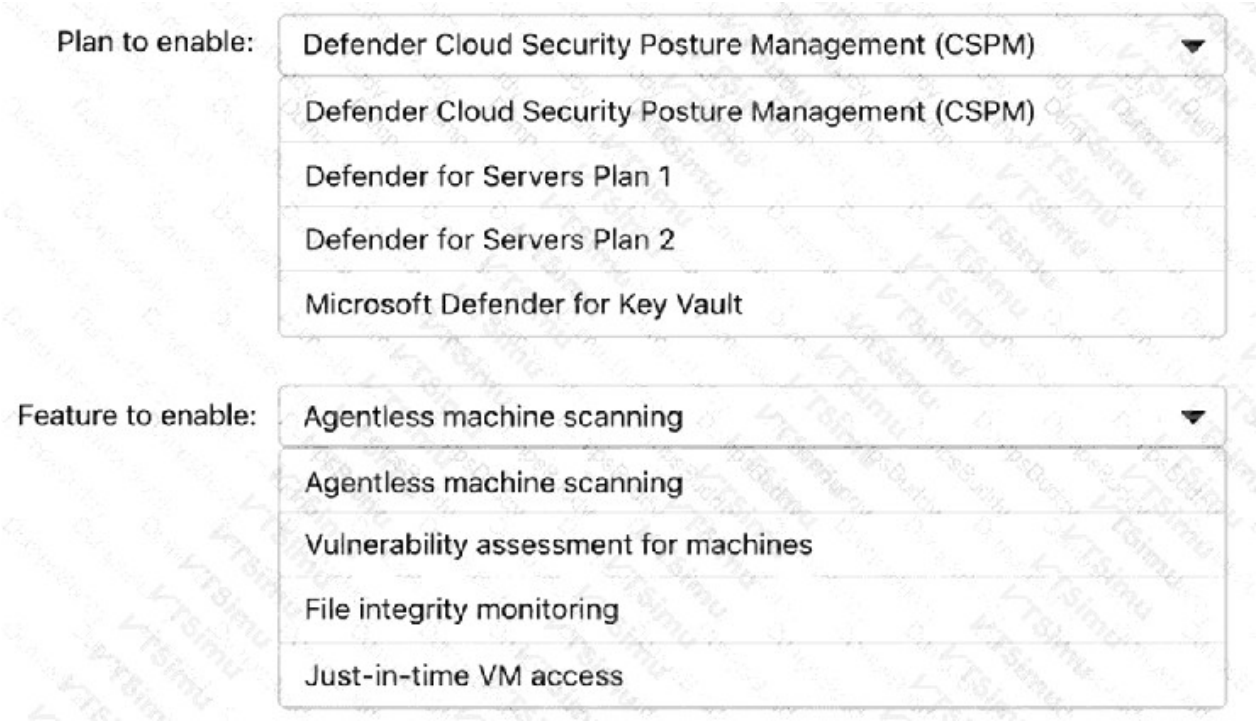

ANSWER:

Explanation:

Enable **Defender Cloud Security Posture Management (CSPM)** and then enable **Agentless machine scanning**. This configuration provides the agentless secret-scanning capability needed to identify potentially exposed credentials, such as plaintext tokens, passwords, connection strings, and SSH private keys, in virtual-machine file systems. Defender for Cloud performs the assessment by using an agentless approach to inspect supported Azure VM disks, rather than requiring developers or administrators to install, configure, update, and maintain a security agent on every machine.

This directly satisfies the requirement to minimize administrative changes across the 50 VMs. The organization can centrally enable the relevant Defender for Cloud capability at the subscription scope and receive security findings for discovered secrets. Those findings can then be investigated and remediated by moving the secret material into the existing Key Vault

and removing it from the VM. The existing policy requiring secrets to be stored in KV1 supports the desired remediation process, while agentless secret scanning supplies the detection mechanism for credentials that were nevertheless written to machine disks.

Microsoft documents secret scanning as part of the agentless machine-scanning capabilities available through Defender CSPM. It is designed to improve cloud posture by discovering sensitive credentials in workloads without adding software to the assessed virtual machines. See Agentless machine scanning in Microsoft Defender for Cloud and Defender Cloud Security Posture Management.

QUESTION NO: 4 - (HOTSPOT)

You have an Azure virtual network named VNet1 that contains three subnets named Subnet1, Subnet2 and Subnet3. A single network security group (NSG) named NSG1 is associated with all the subnets. You have the following virtual machines:

â€¢VM1 on Subnet1

â€¢VM2 on Subnet2

VM3 on Subnet3

You create two application security groups named ASG1 and ASG2. VM2 is a member of ASG1, and VM3 is a member of ASG2.

You need to ensure that only VM2 can connect to VM3. The solution must continue to work if the private IP address of VM2 changes.

How should you configure the inbound rule on NSG1 ? To answer, drag the settings to the correct configurations. Each setting 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.

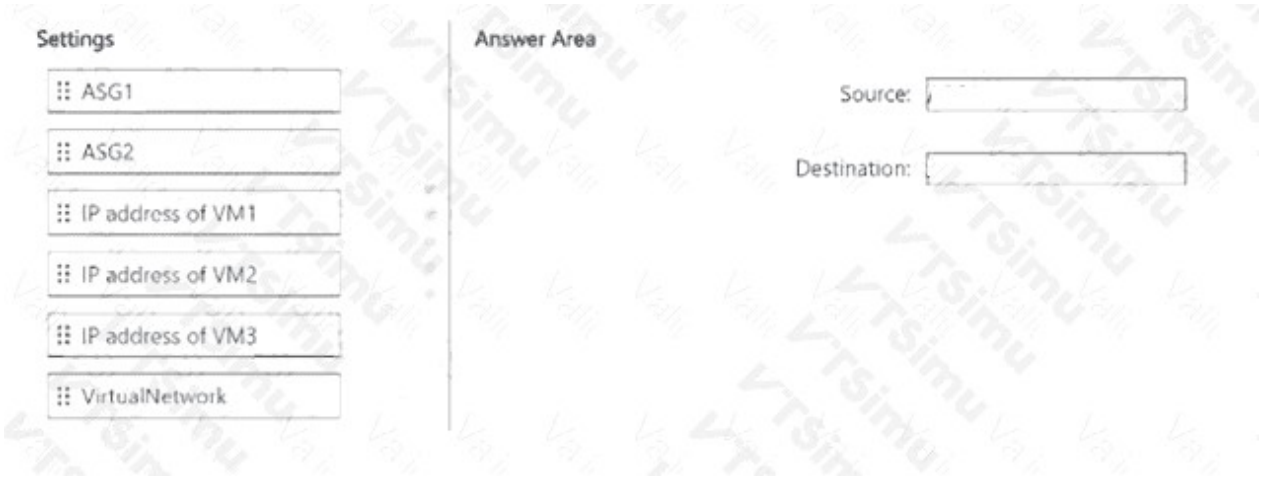

ANSWER:

Answer:

Explanation:

Source: ASG1;

Destination: ASG2 The requirement is identity-stable network filtering between virtual machines even if VM2 receives a different private IP address. Application security groups solve exactly that problem: rules refer to VM membership instead of a fixed IP. Because VM2 is a member of ASG1 and VM3 is a member of ASG2, the inbound allow rule on the shared NSG must use ASG1 as source and ASG2 as destination. Choosing IP addresses would fail the change-resilience requirement. The important exam skill is separating data-plane access, management-plane administration, and network reachability. A storage, database, or firewall setting must be selected because it enforces the exact path requested in the scenario. Distractors often look plausible because they improve security generally, but they do not satisfy the protocol, scope, or automation requirement stated in the question. The result is a direct exam-style implementation choice: it changes the required security behavior without relying on unrelated monitoring, manual cleanup, or excessive privilege. Official Microsoft source/topic: SC-500 Study Guide > NSGs and ASGs; Microsoft Learn > Application security groups in network security rules.==========

Explanation:

The inbound rule should use **ASG1** as its source and **ASG2** as its destination. Application security groups let an NSG rule target groups of network interfaces based on their assigned application security group membership. In this scenario, VM2's network interface is a member of ASG1 and VM3's network interface is a member of ASG2. The rule therefore expresses the intended workload-to-workload path directly: traffic originating from the VM2 workload group is permitted to reach the VM3 workload group.

This is the appropriate configuration when the source VM's private address must be allowed to change without requiring NSG rule maintenance. An ASG reference is not a hard-coded IP address. Azure evaluates the membership of the application security groups when it processes traffic, so the association remains effective when VM2 receives a different private IP address, provided its network interface remains assigned to ASG1. Likewise, the destination remains represented by VM3's ASG2 membership rather than its individual address.

Because NSG1 is associated with the subnets, an inbound rule that uses these ASG references supplies a reusable, workload-oriented policy across the virtual network. This approach makes the rule easier to understand and maintain: the security intent is based on the role of each machine rather than on temporary network addressing. Microsoft documents that application security groups can be used as the source and destination in NSG security rules, allowing security policy to be defined around applications and workloads. See Application security groups and Network security groups for the supported rule behavior and configuration model.

QUESTION NO: 5

You have Microsoft Security Copilot agents that authenticate by using Microsoft Entra service principals.

You receive a Microsoft Defender alert triggered by the anomalous OAuth authentication of an agent ' s Microsoft Entra service principal.

You need to assess the impact of the agent identity and identify which resources are affected if the identity is abused for lateral movement The solution must minimize administrative effort.

What should you do?

A. From Advanced hunting, create a query against the IdentityLogonEvents table to list all the sign-ins performed by the identity.

B. From Attack paths, select the identity and view the blast radius.

C. From AI Observability in Microsoft Purview Data Security Posture Management (DSPM), review the agent activity.

D. From Microsoft Purview Audit, query the audit logs for all the role assignments granted to the identity.

E. From Incidents, review incidents related to OAuth events reported by Microsoft Defender for Cloud Apps.

ANSWER: B

Explanation:

From Attack paths, select the identity and view the blast radius. is correct because Microsoft Security Exposure Management in the Defender portal is designed to model how an attacker could traverse relationships between identities, permissions, resources, and exposures. Selecting the affected Microsoft Entra service principal lets the security team assess its blast radius in the context of reachable resources and potential lateral-movement paths, rather than merely reviewing individual authentication records. This provides an environment-level view of the identity's effective exposure and helps prioritize response based on the resources that could be impacted if the service principal were compromised or abused.

Attack paths correlate security signals and configuration relationships into actionable paths, reducing the need for administrators to manually collect sign-in events, enumerate role assignments, and map resource access across separate tools. That directly meets the requirement to identify affected resources while minimizing administrative effort. The investigation can then guide containment actions for the service principal, such as disabling its credentials, restricting permissions, or remediating the exposures identified along the path. Microsoft documents Attack paths as a capability for identifying and investigating routes that attackers can use to reach critical assets: Attack paths in Microsoft Defender XDR.

QUESTION NO: 6

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 a Microsoft Sentinel workspace

You have a multi-tier Security Operations Center (SOC) team.

You need to ensure that all new security incidents are assigned immediately to the Tier 1 analysts group and flagged for triage.

Solution You create a hunting query.

Does this meet the goal'

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. A Microsoft Sentinel hunting query is designed for proactive, analyst-driven investigation of data already collected in the workspace. It uses Kusto Query Language (KQL) to search for indicators, patterns, and potentially suspicious activity that may not be covered by an existing analytics rule. Running or saving a hunting query does not create an incident-management action when a new incident is generated. The stated goal requires immediate operational handling of incidents: assigning each new incident to the Tier 1 analysts group and applying a triage flag. In Microsoft Sentinel, this type of incident automation is implemented by an automation rule, optionally invoking a playbook where further workflow actions are required. Automation rules can run when an incident is created and can change incident properties, including status, owner, severity, and tags. A tag can identify the incident as requiring triage, while assignment ensures it is directed to the appropriate SOC team without waiting for manual analyst intervention. Hunting remains useful for investigating threats, but it is not the control that enforces automatic incident routing and tagging. See Microsoft Sentinel hunting and automate incident handling with automation rules.

QUESTION NO: 7

You have Microsoft Entra service principals that are used by Microsoft Security Copilot agents.

A Microsoft Defender alert indicates anomalous OAuth authentication by the service principal used by Agent1. You need to contain the identity immediately and determine which cloud resources could be affected if the identity is used for lateral movement.

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

NOTE: Each correct selection is worth one point.

A. Disable the service principal used by Agent1.

B. From Attack paths, select the service principal and review the blast radius.

C. Rotate the access keys for every Azure Storage account in the subscription.

D. Delete the Microsoft Security Copilot workspace used by Agent1.

E. Query Microsoft Purview Audit logs only for historical sign-in activity.

ANSWER: A B

Explanation:

Disable the service principal to immediately prevent new token acquisition and use Attack paths to select the identity and review its blast radius. Disabling the service principal is an effective containment action when the identity is suspected of compromise. Attack paths provides contextual analysis of how the identity's permissions, resource relationships, and exposures could enable access to other cloud resources.

Reviewing audit or sign-in logs can provide useful investigation evidence, but it does not provide the same relationship-based blast-radius analysis. Rotating an unrelated Azure Storage account key does not contain a Microsoft Entra service principal. Removing the Security Copilot workspace does not revoke the service principal's permissions in Azure or Microsoft Entra ID.

QUESTION NO: 8 - (SIMULATION)

You have three internet-facing Azure App Service web apps named App1, App2, and App1 Each app uses built-in authentication.

App2 hosts a backend API.

Some corporate users can sign in to App2, even though they should NOT be able to use the API.

You need to restrict App2 access to assigned Microsoft Entra users and groups.

What should you configure for App2? To answer, drag the appropriate configurations to the correct methods. 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.

ANSWER: See the explanation for the answer

Explanation:

Configure App2 as follows:

Authentication method: The Microsoft identity provider
Access control method: Enterprise app assignment required

The Microsoft identity provider enables Microsoft Entra ID authentication for the App Service built-in authentication feature. Setting the related enterprise application to require assignment ensures that only explicitly assigned Microsoft Entra users and groups can obtain access. Users who are not assigned are denied, even if they otherwise belong to the tenant.

`IP access restrictions` filters by source network address, not user/group assignment. RBAC, local authentication, tenant-wide admin consent, and the other listed options do not enforce this application sign-in restriction.