GIAC
GASAE
GIAC AI Security Automation Engineer

For More Information – Visit link below:
https://www.examsempire.com/

Product Version

1. **Up to Date products, reliable and verified.**
2. **Questions and Answers in PDF Format.**



Latest Version: 6.0

Question: 1

A deployIfNotExists policy that enables Defender for Storage settings is assigned at a management group. Newly created storage accounts in scope show as noncompliant, and the policy's deployments fail with an authorization error. The policy definition lists the roles its deployment requires.
What is the most likely cause?

A. No remediation task has been created, so the policy only evaluates and never deploys
B. The assignment's managed identity lacks the required roles at the scope it deploys into
C. The existence condition is evaluated before the account finishes provisioning
D. Storage accounts must be onboarded to Sentinel before Defender settings apply

Answer: B

Question: 2

A triage script truncates each alert payload to 6,000 characters before sending it to a hosted LLM API. Alerts written as English prose of that length are accepted, but alerts of the same character length that consist mostly of base64-encoded PowerShell and hex dumps are rejected with a context-length error.
What explains the difference?

A. The API meters UTF-8 bytes, and base64 text takes several bytes for each character
B. The model decodes the base64 before counting, which expands the script back to its full size
C. The limit counts words, and each unbroken blob is padded out to a fixed word size
D. Encoded strings split into far more tokens per character, and the limit counts tokens

Answer: D

Question: 3

An email-triage agent reads messages from a shared phishing-report mailbox and summarizes them for analysts. It has two tools: search_mailbox and send_email. In a red-team test, a reported message contains hidden text telling the agent to forward the last ten reports to an external address, and the agent does so.
Which two changes most reduce this risk?
(Choose two.)

A. Require analyst approval before any email the agent composes is sent
B. Add a system-prompt rule telling the agent to ignore instructions in emails

C. Remove send_email and give the triage agent read-only mailbox access
D. Use a model with a larger context window so whole threads fit in one call
E. Lower the temperature so the agent adheres more closely to its system prompt

Answer: A,C

Question: 4

A MITRE Caldera operation report for a lab exercise shows that the discovery abilities completed and the agent kept running later abilities afterwards. One ability that copies a test file to a remote share, which is in the adversary profile, was not run. Its command references a variable for the remote host that should be filled from discovery output, and the operation's fact list holds no value of that type.
What best explains why the ability was skipped?

A. The adversary profile omitted the ability, so the planner had nothing to schedule
B. A required fact was absent
C. The agent lost its connection and stopped beaconing before that ability
D. The ability's ATT&CK technique ID was missing from its definition

Answer: B

Question: 5

A pipeline deploys a Bicep template that defines a baseline network security group into a resource group that also holds a jump-box VM, its NIC and its disk, none of which are in the template. After the latest run the VM, NIC and disk were deleted. The template content was unchanged; only a deployment setting in the pipeline was edited.
What most likely caused the deletion?

A. The NSG baseline blocked management traffic, so the VM was deprovisioned
B. A what-if run applied its deletions
C. An Azure Policy deny effect removed resources violating the baseline
D. The deployment mode was switched to complete

Answer: D

Question: 6

An enrichment script makes two consecutive calls to a hosted LLM API. Call 1 sends a system prompt and the full alert JSON and asks for a one-paragraph summary, which comes back accurate. Call 2 sends the same system prompt and the single user message Now list every IOC from that alert as JSON. The reply contains plausible-looking IP addresses and hashes that do not appear in the alert.

What is the most likely cause?

A. Each API call is independent, so call 2 never contained the alert; prior turns must be resent with every request
B. The alert JSON from call 1 overflowed the context window; it was dropped before call 2 ran
C. The cached copy of call 1 expired on the provider side before the second request reached it
D. Call 2 used a higher default temperature than call 1, which made the IOC output random

Answer: A

Question: 7

An attacker used a stolen SSH private key to log in to a Linux bastion host as svc-deploy and added a public key of their own. The automated remediation playbook currently only resets the account password.
Which two additional actions are needed to evict the attacker?
(Choose two.)

A. Regenerate the bastion's SSH host keys so the stolen client key no longer matches the server
B. Move the SSH daemon to a non-standard port and update the firewall to allow only the new port
C. Remove unauthorized entries from the account's authorized_keys file and revoke the stolen key
D. Lock the password for svc-deploy so password authentication is refused
E. Terminate the account's existing SSH sessions after the credentials are cleaned up

Answer: C,E

Question: 8

A scheduled analytics rule in Microsoft Sentinel is set to run every 60 minutes and to query the previous 10 minutes of data. During a purple-team exercise, a burst of password-spray sign-ins that the query logic matches produced no alert. The rule had not fired earlier that day.
What explains the miss?

A. Scheduled rules create incidents only when an automation rule is attached to them
B. Each run looks back only 10 minutes, so 50 minutes of each hour go unqueried and the burst fell in that gap
C. Alert suppression on the rule discarded the alerts raised after the first trigger
D. The query must join a watchlist before a rule can match sign-in events

Answer: B

Question: 9

A Step Functions workflow contains a compromised EC2 instance by replacing its security groups with an isolation group, then queues EBS snapshots and a memory capture. During the last incident the instance was terminated a few minutes after isolation, before memory was captured. It belongs to an Auto Scaling group that uses load balancer health checks, and those checks began failing once the isolation group was applied.
Which change to the workflow prevents a repeat?

A. Detach the instance from its Auto Scaling group as the first step
B. Enable termination protection on the instance before replacing its security groups
C. Lengthen the group's health check grace period
D. Take the EBS snapshots first, so disk contents are kept before the swap

Answer: A

Question: 10

An EDR alert names source IP 10.20.4.17. The enrichment step queries the CMDB by IP and labels the asset as a low-criticality lab printer owned by Facilities. The analyst later finds the alert came from a finance manager's laptop on DHCP.
What is the most likely enrichment design flaw?

A. The enrichment ran before deduplication; it attached data from an earlier alert or case
B. The CMDB does not record criticality for laptops, so the step fell back to a default rating
C. IP addresses were used as stable keys, so resolve the IP to a host identity from DHCP or EDR data for the alert time first
D. Geolocation was not run on the source IP, so the step could not place the device in the finance office

Answer: C

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

Check our Customer Testimonials and ratings available on every product page.

Visit our website.

https://examsempire.com/