PassQueen
a.k.a dumpsqueen.com

Google Cloud CertifiedProfessional Security Operations Engineer (PSOE) Exam

Google Security-Operations-Engineer

Version Demo

Total Demo Questions: 8

Total Premium Questions: 80

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Planning and configuring a security operations strategy	6
Topic 2, Managing and implementing security operations	33
Topic 3, Detecting, investigating, and responding to threats	35
Topic 4, Optimizing security operations	6
Total	80

QUESTION NO: 1

You are helping a new Google Security Operations (SecOps) customer configure access for their SOC team. The customer ' s Google SecOps administrators currently have access to the Google SecOps instance. The customer is reporting that the SOC team members are not getting authorized to access the instance, but they are able to authenticate to the third-party identity provider (IdP). How should you fix the issue?

Choose 2 answers

A. Link Google SecOps to a Google Cloud project with the Chronicle API.

B. Connect Google SecOps with the third-party IdP using Workforce Identity Federation.

C. Grant the appropriate data access scope to the SOC team ' s IdP group in IAM.

D. Grant the roles/chronicle.viewer role to the SOC team ' s IdP group in IAM.

E. Grant the Basic permission to the appropriate IdP groups in the Google SecOps SOAR Advanced Settings.

ANSWER: D E

Explanation:

Successful authentication at the third-party identity provider establishes that the users can prove their identities, but it does not grant them permissions in Google Security Operations. The SOC team's federated IdP group must be authorized in Google Cloud IAM for the Google Cloud project associated with the SecOps instance. Granting `roles/chronicle.viewer` to that group supplies the baseline Google Security Operations viewer access required for users to enter and view the SecOps experience. For federated identities, the group is typically represented in IAM through the Workforce Identity Federation principal or principal set configured for the IdP.

Access to Google Security Operations SOAR is managed separately from the Google Cloud IAM permission that provides access to the SecOps instance. The same appropriate IdP group must therefore also be assigned a SOAR permission in Google SecOps SOAR Advanced Settings. The Basic permission establishes SOAR access for members of the group, allowing the organization to apply the appropriate SOAR role model after authentication and instance authorization are in place. Together, the IAM role assignment and the SOAR group permission address both authorization layers required for SOC users. See the Google Security Operations IAM roles documentation and Google SecOps SOAR user management documentation.

QUESTION NO: 2

Your organization has onboarded a third-party endpoint detection and response tool to Google Security Operations (SecOps). The SOC wants to automate endpoint containment for high-confidence malware alerts but retain forensic evidence and avoid automatically isolating every low-confidence detection.

Choose 2 answers

A. Use conditional playbook logic to require an approved high-confidence threshold before containment runs.

B. Use the EDR integration to quarantine the affected endpoint when the containment criteria are met.

C. Reboot the endpoint automatically after every malware alert to remove persistence.

D. Delete the endpoint from endpoint management after the EDR alert is received.

E. Block only the destination IP address observed in the alert, and leave the endpoint fully connected.

ANSWER: A B

Explanation:

Use a conditional playbook branch based on alert severity, confidence, or other approved risk criteria before invoking the containment action. This limits automatic isolation to alerts that meet the organization's defined threshold and avoids broadly disrupting endpoints based on low-confidence detections.

Use the EDR integration to quarantine the affected endpoint when the containment criteria are met. Endpoint quarantine limits network communication while generally preserving the running system for evidence collection and investigation. Rebooting or deleting the endpoint can destroy volatile evidence, and blocking only an observed IP address does not reliably contain a compromised host that might communicate with other infrastructure.

QUESTION NO: 3

You are conducting proactive threat hunting in your company ' s Google Cloud environment. You suspect that an attacker compromised a developer ' s credentials and is attempting to move laterally from a development Google Kubernetes Engine (GKE) cluster to critical production systems. You need to identify IoCs and prioritize investigative actions by using Google Cloud ' s security tools before analyzing raw logs in detail. What should you do next?

A. In the Security Command Center (SCC) console, apply filters for the cluster and analyze the resulting aggregated findings ' timeline and details for IoCs. Examine the attack path simulations associated with attack exposure scores to prioritize subsequent actions.

B. Review threat intelligence feeds within Google Security Operations (SecOps), and enrich any anomalies with context on known IoCs, attacker tactics, techniques, and procedures (TTPs), and campaigns.

C. Investigate Virtual Machine (VM) Threat Detection findings in Security Command Center (SCC). Filter for VM Threat Detection findings to target the Compute Engine instances that serve as the nodes for the cluster, and look for malware or rootkits on the nodes.

D. Create a Google SecOps SOAR playbook that automatically isolates any GKE resources exhibiting unusual network connections to production environments and triggers an alert to the incident response team.

ANSWER: A

Explanation:

In the Security Command Center (SCC) console, apply filters for the cluster and analyze the resulting aggregated findings' timeline and details for IoCs. Examine the attack path simulations associated with attack exposure scores to prioritize subsequent actions. This approach directly supports early-stage threat hunting because Security Command Center centralizes security findings and asset context from Google Cloud security services. Filtering findings to the development GKE cluster gives the analyst a focused view of relevant signals, including potentially suspicious activity, vulnerable configurations, exposed resources, and identities or service accounts associated with the cluster. The finding details and event timelines provide immediately actionable context that can be used to identify potential indicators of compromise before undertaking broad, raw-log investigation.

Attack path simulation is especially applicable to the suspected lateral-movement scenario. It models reachable attack paths through an organization’s cloud resources, identities, permissions, configurations, and network exposure. The associated attack exposure scores help rank paths by risk, enabling the analyst to investigate the routes most likely to enable a compromise of production resources. This combines evidence collection with risk-based prioritization, so investigation can concentrate first on the developer identity, GKE workload identities or service accounts, and production assets connected by the highest-risk paths. See Google’s Attack path simulation overview and Security Command Center findings documentation.

QUESTION NO: 4

You are responsible for monitoring the ingestion of critical Windows server logs to Google Security Operations (SecOps) by using the Bindplane agent. You want to receive an immediate notification when no logs have been ingested for over 30 minutes. You want to use the most efficient notification solution. What should you do?

A. Configure the Windows server to send an email notification if there is an error in the Bindplane process.

B. Create a new YARA-L rule in Google SecOps SIEM to detect the absence of logs from the server within a 30-minute window.

C. Configure a Bindplane agent to send a heartbeat signal to Google SecOps every 15 minutes, and create an alert if two heartbeats are missed.

D. Create a new alert policy in Cloud Monitoring that triggers a notification based on the absence of logs from the server's hostname.

ANSWER: D

Explanation:

Creating an alert policy in Cloud Monitoring that triggers a notification based on the absence of logs from the server's hostname is the most efficient solution because it centrally monitors the ingestion signal and automatically opens an incident when the expected data stops arriving. Configure an absence condition for the ingestion metric or log-based metric scoped to the Windows server hostname, set the absence duration to 30 minutes, and attach an appropriate notification channel, such as email, PagerDuty, or a webhook. This approach detects the operational condition that matters—missing security telemetry—rather than relying on the agent process to identify and report its own failure. It also supports standardized alerting, incident management, escalation, and auditability across all monitored log sources. Cloud Monitoring metric-absence conditions are designed specifically for alerting when a time series stops reporting during a configured interval. Google SecOps ingestion monitoring uses Cloud Monitoring metrics to provide visibility into ingestion activity and health, making the alert policy a managed and scalable way to identify interruptions promptly. See Monitor data ingestion for Google SecOps and Metric-absence alerting conditions in Cloud Monitoring.

QUESTION NO: 5

You are developing a new detection rule in Google Security Operations (SecOps). You are defining the YARA-L logic that includes complex event, match, and condition sections. You need to develop and test the rule to ensure that the detections are accurate before the rule is migrated to production. You want to minimize impact to production processes. What should you do?

A. Develop the rule logic in the UDM search, review the search output to inform changes to filters and logic, and copy the rule into the Rules Editor.

B. Use Gemini in Google SecOps to develop the rule by providing a description of the parameters and conditions, and transfer the rule into the Rules Editor.

C. Develop the rule in the Rules Editor, define the sections of the rule logic, and test the rule using the test rule feature.

D. Develop the rule in the Rules Editor, define the sections of the rule logic, and test the rule by setting it to live but not alerting. Run a YARA-L retrohunt from the rules dashboard.

ANSWER: C

Explanation:

Develop the rule in the Rules Editor, define the sections of the rule logic, and test the rule using the test rule feature. This is the appropriate low-impact workflow because the Rules Editor is Google Security Operations' native environment for authoring YARA-L detections and validating the complete rule before deployment. A YARA-L rule's behavior depends on the interaction of its event, match, outcome, and condition logic; testing the rule as a whole is therefore essential for confirming that it identifies the intended activity and produces useful results.

The Rules Editor's test capability lets an analyst execute the draft rule against available historical security-event data and inspect the detections that the logic would generate. This enables iterative tuning of event filters, variable joins, match windows, aggregations, and conditions without enabling the rule for ongoing production alerting. The engineer can use the returned matches to validate expected detections and refine the rule before it is deployed as a live rule. This directly meets the requirement to test accuracy while minimizing operational impact, because the draft-validation activity does not require introducing a production alerting workflow.

Google documents rule creation and testing workflows in the Create detection rules guide and the Test a rule guide.

QUESTION NO: 6

You received an alert from Container Threat Detection that an added binary has been executed in a business critical workload. You need to investigate and respond to this incident. What should you do?

Choose 2 answers

A. Review the finding, quarantine the cluster containing the running pod. and delete the running pod to prevent further compromise.

B. Keep the cluster and pod running, and investigate the behavior to determine whether the activity is malicious.

C. Notify the workload owner. Follow the response playbook. and ask the threat hunting team to identify the root cause of the incident.

D. Review the finding, investigate the pod and related resources, and research the related attack and response methods.

E. Silence the alert in the Security Command Center (SCC) console, as the alert is a low severity finding.

ANSWER: C D

Explanation:

“Review the finding, investigate the pod and related resources, and research the related attack and response methods.” is correct because an added-binary-executed finding can indicate that an unexpected executable was introduced into a running container. The finding details establish the affected workload, container, and detection context. Investigation should then expand to relevant Kubernetes resources and runtime evidence, such as pod specifications, image provenance, service accounts, node context, processes, and network activity. Researching the applicable attack technique and recommended response helps accurately scope the event and select evidence-preserving containment actions.

“Notify the workload owner. Follow the response playbook. and ask the threat hunting team to identify the root cause of the incident.” is also correct because this alert involves a business-critical workload and requires coordinated incident handling in addition to technical triage. The workload owner can validate expected application behavior and assess operational impact, while the established response playbook ensures that escalation, evidence collection, containment approval, and communications follow the organization’s defined process. Threat hunting support is appropriate for determining initial access, persistence, lateral movement, and the broader scope of a potential compromise. Google recommends reviewing Security Command Center findings and investigating affected resources as part of finding management and threat response. See Container Threat Detection documentation and Security Command Center finding management documentation.

QUESTION NO: 7

Your organization uses Security Command Center (SCC) with Event Threat Detection enabled. The SOC needs detections for possible data exfiltration from a limited set of sensitive Cloud Storage and BigQuery workloads. The sensitive resources are located in dedicated projects, and the organization wants to minimize Cloud Logging costs.

Choose 2 answers

A. Enable DATA_READ Data Access audit logs for Cloud Storage in the projects containing the sensitive buckets.

B. Enable DATA_READ Data Access audit logs for BigQuery in the projects containing the sensitive datasets.

C. Enable DATA_WRITE Data Access audit logs for all Cloud Storage and BigQuery projects in the organization.

D. Enable Data Access audit logs only on the individual sensitive buckets and datasets.

E. Enable VPC Flow Logs for the relevant VPC networks instead of Cloud Data Access audit logs.

ANSWER: A B

Explanation:

Enable DATA_READ Data Access audit logs for Cloud Storage in the projects containing the sensitive buckets. Read operations provide the audit evidence required to detect potentially suspicious extraction of stored objects.

Enable DATA_READ Data Access audit logs for BigQuery in the projects containing the sensitive datasets. BigQuery read activity is the relevant source for identifying data-access behavior that could indicate exfiltration. Applying the configuration to the dedicated projects is the narrowest practical hierarchy scope for limiting log volume while retaining coverage for the sensitive resources.

DATA_WRITE logs are not required for the stated read-and-exfiltration detection objective. Audit-log configuration cannot be applied directly to individual buckets or datasets, and VPC Flow Logs do not provide the API-level Data Access audit evidence used by these detections.

QUESTION NO: 8

Your organization uses the curated detection rule set in Google Security Operations (SecOps) for high priority network indicators. You are finding a vast number of false positives coming from your on-premises proxy servers. You need to reduce the number of alerts. What should you do?

A. Configure a rule exclusion for the target.ip field.

B. Configure a rule exclusion for the principal.ip field.

C. Configure a rule exclusion for the network.asset.ip field.

D. Configure a rule exclusion for the target.domain field.

ANSWER: B

Explanation:

Configure a rule exclusion for the principal.ip field. In the Google Security Operations Unified Data Model (UDM), `principal` represents the entity that initiates or performs an action. For network telemetry traversing an on-premises proxy, the proxy commonly appears as the source-side, initiating network entity recorded in `principal.ip`. Consequently, a high-volume proxy can repeatedly match network-indicator logic and generate alerts that do not provide useful security signal for the organization.

A rule exclusion scoped to `principal.ip` lets the curated detection continue to evaluate other traffic while suppressing matches associated with the known proxy IP address or addresses. This is the appropriate targeted tuning approach: it reduces alert noise without disabling the curated rule set or broadly excluding unrelated indicators. Before applying the exclusion, validate the proxy address inventory and document the business justification, then periodically review the exclusion as proxy infrastructure and threat-detection requirements change. Google SecOps supports rule exclusions for curated detections to tune noisy, known-safe activity, and its UDM field model defines the principal as the actor or source associated with an event. See the Google SecOps curated detections documentation and the UDM field list.