PassQueen
a.k.a dumpsqueen.com

Google Cloud Certified - Professional Cloud Security Engineer

Google Professional-Cloud-Security-Engineer

Version Demo

Total Demo Questions: 38

Total Premium Questions: 386

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Configuring access within a cloud solution environment	127
Topic 2, Configuring network security	97
Topic 3, Ensuring data protection	75
Topic 4, Managing operations within a cloud solution environment	52
Topic 5, Supporting compliance requirements	34
Topic 6, Mix Questions	1
Total	386

QUESTION NO: 1

Your organization requires that production administrators use stronger authentication than a password when accessing Google Cloud. The security team also wants to reduce the risk from stolen session cookies and phishing-resistant authentication is required where possible.

Which two controls should the organization implement? (Choose two.)

A. Require 2-Step Verification for production administrator accounts.

B. Require production administrators to use FIDO security keys.

C. Require administrators to change passwords every seven days.

D. Assign the Owner role to a shared administrator account.

E. Allow administrator access only from a VM with an external IP address.

ANSWER: A B

Explanation:

Require 2-Step Verification for administrator accounts. Two-step verification adds an additional authentication factor and reduces the impact of a compromised password. Cloud Identity and Google Workspace administrators can enforce 2-Step Verification for organizational users or groups according to the organization's policy.

Use security keys for administrator accounts. FIDO security keys provide phishing-resistant authentication because the authentication is bound to the legitimate relying party, making a captured password or deceptive sign-in page insufficient for an attacker to complete authentication. Security keys are an appropriate stronger method for highly privileged accounts. See Enforce 2-Step Verification and Sign in with a security key.

QUESTION NO: 2

You are working with a network engineer at your company who is extending a large BigQuery-based data analytics application. Currently, all of the data for that application is ingested from on-premises applications over a Dedicated Interconnect connection with a 20Gbps capacity. You need to onboard a data source on Microsoft Azure that requires a daily ingestion of approximately 250 TB of data. You need to ensure that the data gets transferred securely and efficiently. What should you do?

A. Establish a Cross-Cloud Interconnect connection between Microsoft Azure and Google Cloud. Configure a network route over this connection to transfer the data.

B. Establish a VPN connection with the Microsoft Azure subscription where the source application is running. Transfer the data through the VPN connection.

C. Use the existing Dedicated Interconnect connection through the on-premises network and establish connectivity to Microsoft Azure.

D. Set up a SFTP server with a public IP address that runs on a VM in your Google Cloud project. Connect from Microsoft Azure to this server.

ANSWER: A

Explanation:

Establish a Cross-Cloud Interconnect connection between Microsoft Azure and Google Cloud. Configure a network route over this connection to transfer the data. Cross-Cloud Interconnect is designed specifically for private, dedicated connectivity between Google Cloud and supported external cloud service providers, including Microsoft Azure. It provides a direct physical path between the provider networks rather than sending traffic through the public internet or through the company's on-premises environment.

The required sustained throughput is substantial: transferring 250 TB in 24 hours requires roughly 23.1 Gbps before accounting for protocol overhead, operational headroom, or variations in the transfer rate. A Cross-Cloud Interconnect

deployment can be sized with sufficient capacity and redundancy for this workload, while keeping the data path private and predictable. The implementation includes provisioning the Cross-Cloud Interconnect, creating VLAN attachments, configuring Cloud Router and BGP, and advertising the appropriate Azure and Google Cloud prefixes so that the ingestion traffic uses the dedicated route.

This architecture also keeps the new Azure ingestion independent of the existing 20 Gbps Dedicated Interconnect used for on-premises traffic, avoiding contention and preserving capacity for the current analytics feeds. Google documents Cross-Cloud Interconnect as high-bandwidth dedicated connectivity for connecting Google Cloud with another cloud provider: Cross-Cloud Interconnect overview. Configuration details are available in the Cross-Cloud Interconnect setup guide.

QUESTION NO: 3

You manage a fleet of virtual machines (VMs) in your organization. You have encountered issues with lack of patching in many VMs. You need to automate regular patching in your VMs and view the patch management data across multiple projects.

What should you do?

Choose 2 answers

A. Deploy patches with VM Manager by using OS patch management

B. View patch management data in VM Manager by using OS patch management.

C. Deploy patches with Security Command Center by using Rapid Vulnerability Detection.

D. View patch management data in a Security Command Center dashboard.

E. View patch management data in Artifact Registry.

ANSWER: A B

Explanation:

“Deploy patches with VM Manager by using OS patch management” is correct because VM Manager’s OS patch management capability is designed to automate operating-system updates for Compute Engine VM instances. It lets administrators create patch deployment schedules that run recurring patch jobs, define which instances are targeted through instance filters or labels, and control patch execution behavior. This provides a managed, repeatable approach to addressing missing OS patches across a VM fleet rather than relying on manual updates.

“View patch management data in VM Manager by using OS patch management.” is also correct because VM Manager provides patch compliance and patch-job reporting. The OS patch management dashboard shows patch status and compliance information for managed VM instances, and it supports viewing data at organization, folder, and project scopes when the required permissions and scope selection are available. This centralized visibility enables security and operations teams to review patching posture across multiple projects and identify instances requiring attention.

These capabilities rely on the OS Config service and its agent on supported VM instances. For implementation details, see Google Cloud’s OS patch management documentation and the VM Manager patch dashboard documentation.

QUESTION NO: 4

Your team needs to configure their Google Cloud Platform (GCP) environment so they can centralize the control over networking resources like firewall rules, subnets, and routes. They also have an onpremises environment where resources need access back to the GCP resources through a private VPN connection. The networking resources will need to be controlled by the network security team.

Which type of networking design should your team use to meet these requirements?

A. Shared VPC Network with a host project and service projects

B. Grant Compute Admin role to the networking team for each engineering project

C. VPC peering between all engineering projects using a hub and spoke model

D. Cloud VPN Gateway between all engineering projects using a hub and spoke model

ANSWER: A

Explanation:

Shared VPC Network with a host project and service projects is the appropriate design because it separates centralized network administration from application and workload administration. The network security team can own the host project and create and manage the shared VPC network, including its subnets, routes, firewall policies and rules, and Cloud VPN connectivity to the on-premises environment. Engineering teams can deploy resources in distinct service projects while attaching those resources to subnets defined in the centrally managed host-project network.

This model gives the organization a single, governed connectivity layer for workloads across projects. A private Cloud VPN connection can terminate in the shared VPC, allowing eligible service-project workloads to reach on-premises resources using the host network's routing and security controls. IAM can then delegate narrowly scoped permissions to service-project teams without granting them authority to alter the core network architecture. This is Google Cloud's recommended organizational pattern when central network teams must enforce consistent controls while separate teams operate workloads in their own projects.

See Shared VPC overview and Google Cloud enterprise networking guidance.

QUESTION NO: 5

In order to meet PCI DSS requirements, a customer wants to ensure that all outbound traffic is authorized.

Which two cloud offerings meet this requirement without additional compensating controls? (Choose two.)

A. App Engine

B. Cloud Functions

C. Compute Engine

D. Google Kubernetes Engine

E. Cloud Storage

ANSWER: C D

Explanation:

Compute Engine and **Google Kubernetes Engine** can meet this requirement through Google Cloud VPC firewall rules. VPC firewall rules are stateful and support direction-specific egress policies. An organization can replace reliance on the default implied egress allow behavior with explicit, priority-ordered deny and allow rules, restricting outbound connections to only approved destination ranges, protocols, and ports. Rules can also be targeted to particular VM instances through network tags or service accounts, allowing controls to be applied precisely to workloads in the cardholder data environment.

Google Kubernetes Engine runs its nodes in a VPC network, so VPC firewall policy can govern traffic leaving those nodes. For more granular workload-level controls, GKE can additionally apply Kubernetes NetworkPolicy egress restrictions to pods. Together, these native networking capabilities enable an organization to establish and enforce an authorized outbound-traffic policy without requiring a separate compensating control solely because of the compute platform. The implementation must include appropriately scoped explicit egress rules and ongoing rule governance, consistent with PCI DSS expectations for restricting traffic to necessary and authorized flows. See VPC firewall rules documentation and GKE network overview.

QUESTION NO: 6

You run applications on Cloud Run. You already enabled container analysis for vulnerability scanning. However, you are concerned about the lack of control on the applications that are deployed. You must ensure that only trusted container images are deployed on Cloud Run.

What should you do?

Choose 2 answers

A. Enable Binary Authorization on the existing Kubernetes cluster.

B. Set the organization policy constraint constraints/run. allowedBinaryAuthorizationPolicie tothe list of allowed Binary Authorization policy names.

C. Set the organization policy constraint constraints/compute.trustedimageProjects to the list ofprotects that contain the trusted container images.

D. Enable Binary Authorization on the existing Cloud Run service.

E. Use Cloud Run breakglass to deploy an image that meets the Binary Authorization policy by default.

ANSWER: B D

Explanation:

Enable Binary Authorization on the existing Cloud Run service to apply a deploy-time admission control to that service. Binary Authorization evaluates the image presented for deployment against the configured policy before Cloud Run accepts it. A policy can require attestations from trusted attestors, allowing the organization to define the evidence required for an image to be considered trusted, such as successful review, testing, or vulnerability-scanning gates. This complements Container Analysis: scanning produces security metadata, whereas Binary Authorization enforces whether an image that meets the organization's trust requirements may be deployed.

Set the organization policy constraint constraints/run. allowedBinaryAuthorizationPolicie tothe list of allowed Binary Authorization policy names to limit the Binary Authorization policies that Cloud Run services in the applicable organization-policy scope can use. This governance control prevents service deployments from selecting an arbitrary or weaker policy, helping ensure that Cloud Run services use only organization-approved Binary Authorization policies. Configure the policy hierarchy at the appropriate organization, folder, or project scope, and grant policy-administration permissions separately from routine Cloud Run deployment permissions. Google documents Cloud Run Binary Authorization configuration at Binary Authorization for Cloud Run and the associated Cloud Run organization-policy constraint at Organization Policy constraints.

QUESTION NO: 7

As adoption of the Cloud Data Loss Prevention (DLP) API grows within the company, you need to optimize usage to reduce cost. DLP target data is stored in Cloud Storage and BigQuery. The location and region are identified as a suffix in the resource name.

Which cost reduction options should you recommend?

A. Set appropriate rowsLimit value on BigQuery data hosted outside the US and set appropriate bytesLimitPerFile value on multiregional Cloud Storage buckets.

B. Set appropriate rowsLimit value on BigQuery data hosted outside the US, and minimize transformation units on multiregional Cloud Storage buckets.

C. Use rowsLimit and bytesLimitPerFile to sample data and use CloudStorageRegexFileSet to limit scans.

D. Use FindingLimits and TimespanContfig to sample data and minimize transformation units.

ANSWER: C

Explanation:

Use rowsLimit and bytesLimitPerFile to sample data and use CloudStorageRegexFileSet to limit scans. is correct because Sensitive Data Protection (formerly Cloud DLP) charges inspection based on the volume of content scanned. Reducing the amount of data submitted for inspection is therefore a direct, controlled way to reduce cost while retaining the ability to identify sensitive-data patterns in representative or targeted content.

For BigQuery inspection jobs, `rowsLimit` limits how many rows are inspected per table. For Cloud Storage, `bytesLimitPerFile` limits the amount of each file that is inspected. These settings support sampling, allowing a security team to tune scans to the level of assurance required rather than always processing every row and every byte.

`CloudStorageRegexFileSet` further reduces scanned volume by selecting only Cloud Storage objects whose names match an intended include pattern, while optional exclusion patterns can omit known irrelevant objects. This is especially useful when a bucket contains mixed datasets, temporary exports, or historical files that do not need to be included in a particular inspection job. Together, data sampling and explicit file selection reduce inspection volume at the source of DLP API cost. See the InspectJobConfig reference and the Cloud Storage inspection documentation.

QUESTION NO: 8

A customer wants to make it convenient for their mobile workforce to access a CRM web interface that is hosted on Google Cloud Platform (GCP). The CRM can only be accessed by someone on the corporate network. The customer wants to make it available over the internet. Your team requires an authentication layer in front of the application that supports two-factor authentication Which GCP product should the customer implement to meet these requirements?

A. Cloud Identity-Aware Proxy

B. Cloud Armor

C. Cloud Endpoints

D. Cloud VPN

ANSWER: A

Explanation:

Cloud Identity-Aware Proxy is the appropriate service because it provides an identity-aware authentication and authorization layer in front of web applications hosted on Google Cloud. Rather than requiring users to connect to the corporate network before reaching the CRM, Identity-Aware Proxy can protect the application when it is exposed through an external HTTPS load balancer. Before a request reaches the CRM, Identity-Aware Proxy verifies the user's identity and evaluates whether that identity is authorized by the configured access policy.

For the required two-factor authentication, Identity-Aware Proxy integrates with Google Cloud identity systems, including Cloud Identity and Google Workspace, whose user accounts can be governed by two-step verification policies. It can also be used with workforce identity configurations where the organization's identity provider enforces multifactor authentication. This permits mobile employees to authenticate securely over the internet using their managed corporate identities, while retaining centralized, identity-based access control to the CRM. Google documents Identity-Aware Proxy as a service for controlling access to applications based on user identity and context: Identity-Aware Proxy overview. Guidance for requiring multifactor authentication through Cloud Identity or Google Workspace is available at Turn on 2-Step Verification.

QUESTION NO: 9

You are tasked with exporting and auditing security logs for login activity events for Google Cloud console and API calls that modify configurations to Google Cloud resources. Your export must meet the following requirements:

Export related logs for all projects in the Google Cloud organization.

Export logs in near real-time to an external SIEM.

What should you do? (Choose two.)

A. Create a Log Sink at the organization level with a Pub/Sub destination.

B. Create a Log Sink at the organization level with the includeChildren parameter, and set the destination to a Pub/Sub topic.

C. Enable Data Access audit logs at the organization level to apply to all projects.

D. Enable Google Workspace audit logs to be shared with Google Cloud in the Admin Console.

E. Ensure that the SIEM processes the AuthenticationInfo field in the audit log entry to gather identity information.

ANSWER: B D

Explanation:

Create a Log Sink at the organization level with the includeChildren parameter, and set the destination to a Pub/Sub topic. An aggregated sink at the organization scope with `includeChildren` captures matching log entries generated by descendant projects, rather than requiring individual sinks in every project. Pub/Sub is an appropriate destination for near-real-time integration because a subscriber or forwarding service can consume messages and deliver them to the external SIEM. The sink should be configured with filters that retain the relevant audit-log categories, including Admin Activity logs for configuration-changing API and console actions.

Enable Google Workspace audit logs to be shared with Google Cloud in the Admin Console. Google Cloud Audit Logs identify actions performed in Google Cloud resources, while Google Workspace login audit data supplies the user-login activity needed for Google Cloud console access when the organization uses Google Workspace or Cloud Identity. Sharing those audit logs with Google Cloud makes them available in Cloud Logging, where they can be routed and retained alongside the Google Cloud logs for centralized SIEM analysis. This combination provides both organization-wide collection and the required login-event source. See the Cloud Logging routing overview and the Cloud Audit Logs documentation.

QUESTION NO: 10

Your company uses GitHub Actions to deploy infrastructure into Google Cloud. Corporate policy prohibits downloading, creating, or storing user-managed service account keys in GitHub or any other external system.

What should you do?

A. Configure Workload Identity Federation for GitHub Actions and allow the federated identity to impersonate a dedicated service account.

B. Create a user-managed service account key and store it as an encrypted GitHub Actions secret.

C. Grant the GitHub Actions users the Owner role on the deployment projects.

D. Create a shared Google user account and use its password in the GitHub Actions workflow.

ANSWER: A

Explanation:

Configure **Workload Identity Federation** between GitHub Actions and Google Cloud, and allow the external workload to impersonate a suitably scoped service account. Workload Identity Federation exchanges an identity token issued by the external identity provider for short-lived Google Cloud credentials. This avoids distributing a long-lived service account private key to GitHub.

The organization can configure attribute conditions and attribute mappings so that only approved GitHub organizations, repositories, branches, or workflows are eligible to authenticate. The federated principal can then be granted permission to impersonate only the deployment service account, and that service account can receive only the IAM roles required for the deployment. This follows least privilege and provides auditable, short-lived access.

See Workload Identity Federation with deployment pipelines and Workload Identity Federation overview.

QUESTION NO: 11

You are using Security Command Center (SCC) to protect your workloads and receive alerts for suspected security breaches at your company. You need to detect cryptocurrency mining software. Which SCC service should you use?

A. Web Security Scanner

B. Container Threat Detection

C. Rapid Vulnerability Detection

D. Virtual Machine Threat Detection

ANSWER: D

Explanation:

Virtual Machine Threat Detection is the appropriate Security Command Center service for detecting cryptocurrency mining software running on Compute Engine virtual machine instances. It monitors VM workloads for suspicious activity and produces Security Command Center findings when it identifies known indicators of compromise, including malware and cryptomining behavior. Cryptomining detection is particularly important because an attacker who obtains access to a VM can install mining software to consume CPU and other resources, often causing unexpected cost and performance degradation. Virtual Machine Threat Detection uses managed detection capabilities designed for VM runtime threats, so it can surface these findings centrally in SCC for investigation and response. Google documents cryptomining as a supported detection category and provides findings for recognized mining software and related activity. For implementation details and supported threat detections, see Virtual Machine Threat Detection overview and Use Virtual Machine Threat Detection.

QUESTION NO: 12

You are migrating an on-premises data warehouse to BigQuery Cloud SQL, and Cloud Storage. You need to configure security services in the data warehouse. Your company compliance policies mandate that the data warehouse must:

• Protect data at rest with full lifecycle management on cryptographic keys

• Implement a separate key management provider from data management

• Provide visibility into all encryption key requests

What services should be included in the data warehouse implementation?

Choose 2 answers

A. Customer-managed encryption keys

B. Customer-Supplied Encryption Keys

C. Key Access Justifications

D. Access Transparency and Approval

E. Cloud External Key Manager

ANSWER: C E

Explanation:

Cloud External Key Manager is appropriate because it lets Google Cloud services use keys held and managed by a supported external key-management partner or an external key-management system. The external provider retains control of the cryptographic key lifecycle, including key creation, rotation, disabling, and destruction, while the data remains managed by BigQuery, Cloud SQL, and Cloud Storage. This creates the required separation between the data-management environment and the key-management provider. Cloud EKM is used through Cloud KMS external keys, allowing supported Google Cloud services to encrypt data at rest with externally managed key material.

Key Access Justifications provides the required visibility into encryption-key requests when external keys are used. For eligible requests, Google Cloud includes a justification describing the reason that Google infrastructure needs to access the external key. The external key-management system can log these requests and can use the justification in its authorization decision, giving the organization an auditable record and control point for access to its cryptographic keys. Together, these services satisfy external lifecycle ownership and key-request visibility requirements. See Cloud External Key Manager documentation and Key Access Justifications documentation.

QUESTION NO: 13

Your organization needs to reduce the risk of data exfiltration from a project that stores regulated data in Cloud Storage and BigQuery. Access to those services must be allowed only from the organization's approved VPC networks and from approved Google Cloud projects.

Which two actions should you take? (Choose two.)

A. Create a VPC Service Controls service perimeter that includes the regulated-data project.

B. Configure ingress rules and access levels for the approved networks, projects, and identities.

C. Create VPC peering connections between the regulated-data project and every approved project.

D. Enable Cloud NAT on the subnets used by the regulated-data project.

E. Apply Cloud Storage object lifecycle rules to delete objects after they are accessed.

ANSWER: A B

Explanation:

Create a service perimeter around the regulated-data project by using VPC Service Controls. A service perimeter helps mitigate data exfiltration risks by restricting access to supported Google Cloud services and by controlling access across protected projects. Cloud Storage and BigQuery are supported services that can be protected by a perimeter.

Configure ingress rules to allow only the required access from approved projects and identities, and configure access levels when network-based context is required. Access levels can evaluate request attributes such as source IP address and device or identity context. Together, the perimeter and explicit ingress policy restrict access to protected services while allowing the approved usage paths. See VPC Service Controls overview and Ingress and egress rules.

QUESTION NO: 14

You are creating a new infrastructure CI/CD pipeline to deploy hundreds of ephemeral projects in your Google Cloud organization to enable your users to interact with Google Cloud. You want to restrict the use of the default networks in your organization while following Google-recommended best practices. What should you do?

A. Enable the constraints/compute.skipDefaultNetworkCreation organization policy constraint at the organization level.

B. Create a cron job to trigger a daily Cloud Function to automatically delete all default networks for each project.

C. Grant your users the IAM Owner role at the organization level. Create a VPC Service Controls perimeter around the project that restricts the compute.googleapis.com API.

D. Only allow your users to use your CI/CD pipeline with a predefined set of infrastructure templates they can deploy to skip the creation of the default networks.

ANSWER: A

Explanation:

Enable the `constraints/compute.skipDefaultNetworkCreation` organization policy constraint at the organization level. This is the Google-recommended preventive control for an environment that creates many new projects, because it stops Google Cloud from automatically creating a default VPC network in newly created projects beneath the policy's scope. Applying it at the organization level ensures that the policy is inherited consistently by the ephemeral projects created through the CI/CD pipeline, rather than relying on project-by-project cleanup or user behavior.

Default networks contain preconfigured networking resources and broadly permissive firewall rules intended to simplify initial experimentation. In a governed production organization, infrastructure should instead create only explicitly designed VPC networks, subnetworks, routes, and firewall policies through approved infrastructure-as-code templates. Preventing default-network creation is therefore more reliable than detecting and deleting default networks after project creation, and it supports a least-privilege, secure-by-default project provisioning process. Google Cloud documents this Boolean organization policy

constraint as the control that skips creation of default networks for new projects. See Organization Policy constraints: skip default network creation and Default network documentation.

QUESTION NO: 15

You are setting up a new Cloud Storage bucket in your environment that is encrypted with a customer managed encryption key (CMEK). The CMEK is stored in Cloud Key Management Service (KMS). in project "pr j -a", and the Cloud Storage bucket will use project "prj-b". The key is backed by a Cloud Hardware Security Module (HSM) and resides in the region europe-west3. Your storage bucket will be located in the region europe-west1. When you create the bucket, you cannot access the key. and you need to troubleshoot why.

What has caused the access issue?

A. A firewall rule prevents the key from being accessible.

B. Cloud HSM does not support Cloud Storage

C. The CMEK is in a different project than the Cloud Storage bucket

D. The CMEK is in a different region than the Cloud Storage bucket.

ANSWER: D

Explanation:

The CMEK is in a different region than the Cloud Storage bucket. For a Cloud Storage bucket that uses customer-managed encryption keys, the Cloud KMS key must have a location compatible with the bucket's location. In particular, a bucket in the single region `europe-west1` requires a key ring and key in `europe-west1`. A key in `europe-west3` is a distinct regional resource and cannot be selected for encryption of that bucket, so Cloud Storage cannot access it during bucket configuration.

The fact that the KMS key is held in a separate project is supported for CMEK configurations, provided the relevant Cloud Storage service agent receives the required permission to encrypt and decrypt with that key. Likewise, a Cloud HSM-backed key can be used with Cloud Storage; the protection level does not remove the location requirement. The corrective action is to create or use an appropriate KMS key in the bucket-compatible location, then grant the Cloud Storage service agent access to that key. Google documents the Cloud Storage CMEK location requirements at Customer-managed encryption keys and the applicable KMS resource-location concepts at Cloud KMS locations.

QUESTION NO: 16

Your organization must retain security audit logs for five years. Project administrators must not be able to shorten the retention period or delete the retained logs from their own projects.

Which two actions should you take? (Choose two.)

A. Create an organization-level aggregated sink that routes the required logs to a log bucket in a dedicated security project.

B. Lock a five-year retention policy on the destination log bucket.

C. Grant the Project Owner role to the security team in every workload project.

D. Export the logs to a Cloud Storage bucket that has public access prevention disabled.

E. Create a logs-based metric in each project and retain the metric data for five years.

ANSWER: A B

Explanation:

Create an **organization-level aggregated sink** that routes the required audit logs to a log bucket in a dedicated security project. An aggregated sink can collect matching logs from projects throughout the organization, while the destination

security project can be administered separately from the source workload projects. This prevents individual project administrators from controlling the central log destination.

Configure the destination log bucket with a five-year retention period and **lock the retention policy**. Locking a Cloud Logging bucket retention policy makes the configured retention period immutable. Even principals with permissions to update the bucket cannot reduce the locked retention period, which provides the required protection against premature deletion.

See aggregated sinks documentation and Cloud Logging retention policy documentation.

QUESTION NO: 17

You have created an OS image that is hardened per your organization's security standards and is being stored in a project managed by the security team. As a Google Cloud administrator, you need to make sure all VMs in your Google Cloud organization can only use that specific OS image while minimizing operational overhead. What should you do? (Choose two.)

A. Grant users the compute.imageUser role in their own projects.

B. Grant users the compute.imageUser role in the OS image project.

C. Store the image in every project that is spun up in your organization.

D. Set up an image access organization policy constraint, and list the security team managed project in the projects allow list.

E. Remove VM instance creation permission from users of the projects, and only allow you and your team to create VM instances.

ANSWER: B D

Explanation:

Granting users the `roles/compute.imageUser` role on the security team's image project authorizes them to use the shared hardened image when they create VM instances in their own projects. This is the least-operational-overhead permission model because the image is maintained centrally: the security team can publish updated hardened image versions in one project, while appropriately authorized users across the organization can consume them without copying the image into each workload project.

An organization policy using the `constraints/compute.trustedImageProjects` list constraint enforces the corresponding organization-wide control plane restriction. Adding the security-managed image project to the allowed-project list ensures that boot disks for newly created Compute Engine VMs can use images only from that trusted project. When the hardened image is the approved image maintained there, this central policy prevents VM creation from unapproved image projects throughout the policy's scope. Apply the policy at the organization level to cover all current and future projects, while allowing inheritance where appropriate. Google documents both the trusted-image-projects constraint and its use for restricting boot-image sources at Organization Policy constraints. The required cross-project image-use permission is documented in Managing access to custom images.

QUESTION NO: 18

A company is running workloads in a dedicated server room. They must only be accessed from within the private company network. You need to connect to these workloads from Compute Engine instances within a Google Cloud Platform project.

Which two approaches can you take to meet the requirements? (Choose two.)

A. Configure the project with Cloud VPN.

B. Configure the project with Shared VPC.

C. Configure the project with Cloud Interconnect.

D. Configure the project with VPC peering.

E. Configure all Compute Engine instances with Private Access.

ANSWER: A C

Explanation:

Configuring the project with Cloud VPN provides private network connectivity between the Google Cloud VPC network used by the Compute Engine instances and the company's on-premises network. Cloud VPN uses IPsec tunnels, allowing workloads to communicate through private RFC 1918 addressing and controlled VPC/on-premises routes rather than exposing the server-room workloads to public access. For production connectivity, Google recommends HA VPN, which uses redundant tunnels to support high availability.

Configuring the project with Cloud Interconnect also meets the requirement by extending the company network into Google Cloud through private connectivity. Dedicated Interconnect provides direct physical connections to Google's network, while Partner Interconnect is delivered through a supported service provider. In either case, VLAN attachments and Cloud Router exchange routes between the on-premises network and the VPC, enabling Compute Engine instances to reach the dedicated-server-room workloads using private IP addresses. These approaches preserve the requirement that the workloads remain reachable only from the private company network. See the Cloud VPN overview and the Cloud Interconnect overview.

QUESTION NO: 19

You are tasked with exporting and auditing security logs for login activity events for Google Cloud console and API calls that modify configurations to Google Cloud resources. Your export must meet the following requirements:

- Export related logs for all projects in the Google Cloud organization.
- Export logs in near real-time to an external SIEM.

What should you do? (Choose two.)

A. Create a Log Sink at the organization level with a Pub/Sub destination.

B. Create a Log Sink at the organization level with the includeChildren parameter, and set the destination to a Pub/Sub topic.

C. Enable Data Access audit logs at the organization level to apply to all projects.

D. Enable Google Workspace audit logs to be shared with Google Cloud in the Admin Console.

E. Ensure that the SIEM processes the AuthenticationInfo field in the audit log entry to gather identity information.

ANSWER: B E

Explanation:

Creating a Log Sink at the organization level with the includeChildren parameter, and setting the destination to a Pub/Sub topic, provides centralized export coverage for logs generated by projects and other child resources in the organization. An organization-level aggregated sink with `includeChildren` enabled evaluates log entries from descendant resources, avoiding the operational burden and coverage risk of creating and maintaining separate sinks in every project. Pub/Sub is an appropriate destination for near-real-time integration because an external SIEM can consume exported log messages through a subscriber or a supported ingestion integration. The sink's filter can be scoped to the relevant Cloud Audit Logs, including administrative activity associated with configuration-changing API calls.

Ensuring that the SIEM processes the `authenticationInfo` field is also essential for auditability. Cloud Audit Log entries record the identity that made an authenticated request in this field, including identity attributes such as the principal email or principal subject when available. Parsing this information lets the SIEM associate console and API activity with the acting user, service account, or federated principal, enabling investigations, correlation, and alerting based on identity. See Google Cloud documentation for aggregated sinks and the AuditLog authenticationInfo field.

QUESTION NO: 20

A companyâ€™s application is deployed with a user-managed Service Account key. You want to use Google- recommended practices to rotate the key.

What should you do?

A. Open Cloud Shell and run gcloud iam service-accounts enable-auto-rotate --iam- account=IAM_ACCOUNT.

B. Open Cloud Shell and run gcloud iam service-accounts keys rotate --iam- account=IAM_ACCOUNT --key=NEW_KEY.

C. Create a new key, and use the new key in the application. Delete the old key from the Service Account.

D. Create a new key, and use the new key in the application. Store the old key on the system as a backup key.

ANSWER: C

Explanation:

Create a new key, and use the new key in the application. Delete the old key from the Service Account. This follows Google Cloud's recommended manual rotation process for a user-managed service account key. First, create a replacement user-managed key for the same service account and securely deliver it to the application. Then update the application's configuration or secret-management reference so that authentication uses the new key. Before removing the existing key, validate that the application is operating successfully with the replacement, avoiding an authentication outage. Once the new key is confirmed in use, delete the old key from the service account; deletion prevents the old private key from being used if it was copied, exposed, or otherwise compromised. Google recommends avoiding service account keys where possible and using more secure alternatives such as attached service accounts, Workload Identity Federation, or the metadata server. However, when a user-managed key is required, maintaining a limited key lifetime and performing this create, deploy, verify, and delete sequence minimizes the period in which an older credential remains valid. See Creating and managing service account keys and Best practices for managing service account keys.

QUESTION NO: 21

An office manager at your small startup company is responsible for matching payments to invoices and creating billing alerts. For compliance reasons, the office manager is only permitted to have the

Identity and Access Management (IAM) permissions necessary for these tasks. Which two IAM roles should the office manager have? (Choose two.)

A. Organization Administrator

B. Project Creator

C. Billing Account Viewer

D. Billing Account Costs Manager

E. Billing Account User

ANSWER: C D

Explanation:

Billing Account Viewer supplies the read-only billing-account access needed to review financial information, including invoices and transactions. This supports matching payments with invoices without granting administrative control over the organization, projects, or billing-account associations. Google describes this role as appropriate for users who need to view billing-account cost information and financial documents.

Billing Account Costs Manager supplies the cost-management permissions needed to create and manage Cloud Billing budgets and their associated alerting configuration. Budget alerts notify configured recipients when actual or forecasted costs reach specified thresholds. The role follows least privilege because it is scoped to billing cost-management tasks rather than broad organization administration or project-creation authority.

These two roles can be granted on the relevant Cloud Billing account, limiting the office manager's visibility and ability to act to the billing account for which they perform invoice reconciliation and budget-alert work. This meets the stated compliance requirement to provide only permissions necessary for the assigned financial operations.

See Google Cloud's Cloud Billing IAM roles and permissions and budget and alert documentation.

QUESTION NO: 22

Your security team uses encryption keys to ensure confidentiality of user data. You want to establish a process to reduce the impact of a potentially compromised symmetric encryption key in Cloud Key Management Service (Cloud KMS).

Which steps should your team take before an incident occurs? (Choose two.)

A. Disable and revoke access to compromised keys.

B. Enable automatic key version rotation on a regular schedule.

C. Manually rotate key versions on an ad hoc schedule.

D. Limit the number of messages encrypted with each key version.

E. Disable the Cloud KMS API.

ANSWER: B D

Explanation:

Enable automatic key version rotation on a regular schedule because a Cloud KMS key rotation schedule automatically creates a new key version and makes it the primary version for future encryption operations. Establishing this process in advance limits the time during which any one key version is used for newly encrypted data. If a version is later suspected of compromise, the organization has a defined rotation cadence and a smaller set of recently encrypted data to assess and, where necessary, re-encrypt. Cloud KMS does not automatically re-encrypt existing ciphertext during rotation, so teams should also plan the associated data re-encryption and incident-response procedures.

Limit the number of messages encrypted with each key version because reducing the amount of data protected by a single cryptographic key version constrains the potential exposure if that version is compromised. In practice, this can be achieved through a rotation interval appropriate to the data volume and sensitivity, as well as through key-separation designs that use distinct keys for separate applications, tenants, or data domains. Together, regular automated rotation and bounded key-version usage provide a proactive cryptographic risk-reduction process. See Cloud KMS key rotation documentation and Cloud KMS keys and key versions documentation.

QUESTION NO: 23

You are setting up a CI/CD pipeline to deploy containerized applications to your production clusters on Google Kubernetes Engine (GKE). You need to prevent containers with known vulnerabilities from being deployed. You have the following requirements for your solution:

- Must be cloud-native
- Must be cost-efficient
- Minimize operational overhead

How should you accomplish this? (Choose two.)

A. Create a Cloud Build pipeline that will monitor changes to your container templates in a Cloud Source Repositories repository. Add a step to analyze Container Analysis results before allowing the build to continue.

B. Use a Cloud Function triggered by log events in Google Cloud's operations suite to automatically scan your container images in Container Registry.

C. Use a cron job on a Compute Engine instance to scan your existing repositories for known vulnerabilities and raise an alert if a non-compliant container image is found.

D. Deploy Jenkins on GKE and configure a CI/CD pipeline to deploy your containers to Container Registry. Add a step to validate your container images before deploying your container to the cluster.

E. In your CI/CD pipeline, add an attestation on your container image when no vulnerabilities have been found. Use a Binary Authorization policy to block deployments of containers with no attestation in your cluster.

ANSWER: A E

Explanation:

Creating a Cloud Build pipeline that analyzes Container Analysis results before the build continues provides a managed, cloud-native CI control. Container Analysis can scan supported container images for known vulnerabilities and expose vulnerability occurrences that the pipeline can evaluate. A Cloud Build gate can therefore stop an artifact from progressing when its vulnerability findings exceed the organization's defined acceptance criteria. This uses managed Google Cloud services rather than requiring a self-managed scanner or continuously maintained compute infrastructure.

Adding an attestation only after the image has passed vulnerability validation, then enforcing that attestation through Binary Authorization, provides the deployment-time control needed for GKE. The attestation serves as signed evidence that the image completed the required validation process. A Binary Authorization policy can require that evidence before admitting an image to the cluster, preventing an image that bypassed the CI pipeline or lacks the required approval from being deployed. Together, CI validation and admission enforcement create defense in depth: findings are evaluated before release, and GKE independently enforces that only approved artifacts run in production. This approach is managed, cost-efficient, and has low operational overhead. See Artifact Analysis container scanning overview and Binary Authorization overview.

QUESTION NO: 24

A batch job running on Compute Engine needs temporary write access to a Cloud Storage bucket. You want the batch job to use the minimum permissions necessary to complete the task. What should you do?

A. Create a service account with full Cloud Storage administrator permissions. Assign the service account to the Compute Engine instance.

B. Grant the predefined storage.objectcreator role to the Compute Engine instances default service account.

C. Create a service account and embed a long-lived service account key file that has write permissions specified directly in the batch jobscript.

D. Create a service account with the storage .objectcreator role. Use service account impersonation in the batch job's code.

ANSWER: B

Explanation:

Grant the predefined `storage.objectCreator` role to the Compute Engine instance's default service account. Compute Engine workloads obtain short-lived credentials automatically from their attached service account through the metadata server, so the batch job can authenticate to Cloud Storage without distributing or managing a service account key. This is appropriate for temporary workload access because the credentials are automatically rotated and are valid only while the workload can use that service account identity.

The `roles/storage.objectCreator` role provides the narrowly scoped permission required to create objects in Cloud Storage: `storage.objects.create`. It does not include general bucket administration or object read permissions, which supports least privilege when the job only needs to upload output. For further scope reduction, grant this role directly on the specific target bucket rather than at the project level. If the job must overwrite existing objects, the required role must be evaluated separately because overwriting can require additional permissions; for a create-only batch upload, `storage.objectCreator` is the appropriate predefined role.

See the Cloud Storage IAM roles documentation and Google Cloud guidance on using attached service accounts.

QUESTION NO: 25

Your company has been creating users manually in Cloud Identity to provide access to Google Cloud resources. Due to continued growth of the environment, you want to authorize the Google Cloud Directory Sync (GCDS) instance and integrate it with your on-premises LDAP server to onboard hundreds of users. You are required to:

Replicate user and group lifecycle changes from the on-premises LDAP server in Cloud Identity.

Disable any manually created users in Cloud Identity.

You have already configured the LDAP search attributes to include the users and security groups in scope for Google Cloud. What should you do next to complete this solution?

A. 2. Set up a recurring GCDS task.

B. 2. Run GCDS after user and group lifecycle changes.

C. 2. Set up a recurring GCDS task.

D. 2. Run GCDS after user and group lifecycle changes.

E. Configure GCDS to suspend Google users not found in LDAP, and schedule recurring synchronization.

ANSWER: E

Explanation:

Configure GCDS to suspend Google users that are not found in the LDAP search results, then schedule the synchronization to run recurrently. GCDS is designed to synchronize users, groups, and group memberships from an LDAP directory into Cloud Identity or Google Workspace. A recurring scheduled sync ensures that onboarding, offboarding, and group-membership changes in LDAP are propagated without an administrator having to manually launch each synchronization.

The user deletion/suspension policy is essential for the requirement concerning manually created Cloud Identity users. When GCDS evaluates the configured LDAP scope, users that exist in Cloud Identity but are absent from that scope can be suspended. Suspension disables sign-in while preserving the account and its data, which is the appropriate lifecycle state when access must be removed rather than permanently deleting the identity. The LDAP search configuration defines the authoritative in-scope population; the suspension policy applies the required treatment to accounts outside it. Google recommends testing synchronization and reviewing the proposed changes before enabling a production schedule. See Google's GCDS documentation and its guidance for synchronizing users and groups.

QUESTION NO: 26

You discovered that sensitive personally identifiable information (PII) is being ingested to your Google Cloud environment in the daily ETL process from an on-premises environment to your BigQuery datasets. You need to redact this data to obfuscate the PII, but need to re-identify it for data analytics purposes. Which components should you use in your solution? (Choose two.)

A. Secret Manager

B. Cloud Key Management Service

C. Cloud Data Loss Prevention with cryptographic hashing

D. Cloud Data Loss Prevention with automatic text redaction

E. Cloud Data Loss Prevention with deterministic encryption using AES-SIV

ANSWER: B E

Explanation:

Cloud Data Loss Prevention with deterministic encryption using AES-SIV is appropriate because it performs reversible pseudonymization of sensitive values. With deterministic encryption, identical plaintext values produce identical encrypted

values when the same key and context are used. This allows the ETL pipeline to replace PII with stable surrogate values that remain useful for joins, grouping, and other BigQuery analytics. Authorized processes can subsequently use the same cryptographic configuration to re-identify the original values. Cloud Data Loss Prevention supports deterministic encryption through its de-identification and re-identification transformations, including AES-SIV-based encryption.

Cloud Key Management Service should be used to protect the cryptographic material used by the transformation. In particular, Sensitive Data Protection can use a data-encryption key wrapped by a Cloud KMS key; its service agent requires permission to decrypt the wrapped key. This avoids embedding or directly managing sensitive encryption keys in the ETL application, while enabling controlled key access, rotation, auditing, and recovery for authorized re-identification workflows. Together, these services provide reversible protection rather than irreversible masking. See Sensitive Data Protection transformation reference and encryption support documentation.

QUESTION NO: 27

You are the security admin of your company. Your development team creates multiple GCP projects under the "implementation" folder for several dev, staging, and production workloads. You want to prevent data exfiltration by malicious insiders or compromised code by setting up a security perimeter. However, you do not want to restrict communication between the projects.

What should you do?

A. Use a Shared VPC to enable communication between all projects, and use firewall rules to prevent data exfiltration.

B. Create access levels in Access Context Manager to prevent data exfiltration, and use a shared VPC for communication between projects.

C. Use an infrastructure-as-code software tool to set up a single service perimeter and to deploy a Cloud Function that monitors the "implementation" folder via Stackdriver and Cloud Pub/Sub. When the function notices that a new project is added to the folder, it executes Terraform to add the new project to the associated perimeter.

D. Use an infrastructure-as-code software tool to set up three different service perimeters for dev, staging, and prod and to deploy a Cloud Function that monitors the "implementation" folder via Stackdriver and Cloud Pub/Sub. When the function notices that a new project is added to the folder, it executes Terraform to add the new project to the respective perimeter.

ANSWER: C

Explanation:

Use an infrastructure-as-code software tool to set up a single service perimeter and to deploy a Cloud Function that monitors the "implementation" folder via Stackdriver and Cloud Pub/Sub. When the function notices that a new project is added to the folder, it executes Terraform to add the new project to the associated perimeter. A VPC Service Controls service perimeter is designed to mitigate data exfiltration risks for supported Google Cloud services, including risks from compromised workloads and malicious insiders with valid credentials. A perimeter can contain multiple projects, so placing all development, staging, and production projects in one perimeter permits protected-service requests between resources in those projects without requiring an ingress or egress boundary crossing. The protected boundary is instead enforced for requests entering or leaving the perimeter. Service perimeters are explicitly configured with projects; placing projects beneath a Resource Manager folder does not automatically make new projects members of the perimeter. Therefore, automating perimeter membership when a project is created or moved into the implementation folder keeps the intended security boundary current while avoiding manual administrative gaps. Managing the perimeter as infrastructure as code also provides repeatable, reviewable configuration. See the VPC Service Controls service perimeter documentation and the guidance for creating service perimeters.

QUESTION NO: 28

You are a consultant for an organization that is considering migrating their data from its private cloud to Google Cloud. The organization’s compliance team is not familiar with Google Cloud and needs guidance on how compliance requirements will be met on Google Cloud. One specific compliance requirement is for customer data at rest to reside within specific geographic boundaries. Which option should you recommend for the organization to meet their data residency requirements on Google Cloud?

A. Organization Policy Service constraints

B. Shielded VM instances

C. Access control lists

D. Geolocation access controls

E. Google Cloud Armor

ANSWER: A

Explanation:

Organization Policy Service constraints are the appropriate recommendation because they provide centralized, enforceable governance over where supported Google Cloud resources can be created. For a data-residency requirement, an administrator configures the `constraints/gcp.resourceLocations` list constraint at the organization, folder, or project level. The policy can allow only the required regions or multi-regions, such as a specified country-level multi-region or individual regional locations. Requests to create or configure governed resources outside those permitted locations are denied, preventing accidental placement of customer data in an unapproved geography.

This approach complements selecting the appropriate resource location during deployment. The organization should first identify which Google Cloud products store its customer data and confirm that each product supports the required location and the resource-locations constraint. It can then apply the policy at the highest practical level of the resource hierarchy, with carefully controlled exceptions where necessary. This creates an auditable preventive control that helps teams consistently meet geographic data-location commitments as projects and workloads grow.

Google Cloud documents the resource-locations organization policy constraint and its supported values in the Defining resource locations guide. General information on selecting and understanding Google Cloud locations is available in the Cloud locations documentation.

QUESTION NO: 29

You need to follow Google-recommended practices to leverage envelope encryption and encrypt data at the application layer.

What should you do?

A. Generate a data encryption key (DEK) locally to encrypt the data, and generate a new key encryption key (KEK) in Cloud KMS to encrypt the DEK. Store both the encrypted data and the encrypted DEK.

B. Generate a data encryption key (DEK) locally to encrypt the data, and generate a new key encryption key (KEK) in Cloud KMS to encrypt the DEK. Store both the encrypted data and the KEK.

C. Generate a new data encryption key (DEK) in Cloud KMS to encrypt the data, and generate a key encryption key (KEK) locally to encrypt the key. Store both the encrypted data and the encrypted DEK.

D. Generate a new data encryption key (DEK) in Cloud KMS to encrypt the data, and generate a key encryption key (KEK) locally to encrypt the key. Store both the encrypted data and the KEK.

ANSWER: A

Explanation:

Generate a data encryption key (DEK) locally to encrypt the data, and generate a new key encryption key (KEK) in Cloud KMS to encrypt the DEK. Store both the encrypted data and the encrypted DEK.

This implements envelope encryption, the recommended pattern for application-layer encryption when data must be protected before it is sent to, or stored in, another service. The application generates a unique, high-entropy DEK locally and uses that key for the actual data encryption. This approach is efficient because encryption of potentially large data objects occurs locally rather than requiring the data itself to be sent to Cloud KMS.

The application then sends only the DEK to Cloud KMS for encryption by a Cloud KMS key, which serves as the KEK. The resulting wrapped, encrypted DEK is stored alongside the ciphertext. To decrypt later, the application retrieves the encrypted DEK, asks Cloud KMS to decrypt it subject to IAM permissions and key state, and uses the recovered DEK locally to decrypt

the data. Cloud KMS retains and protects the KEK; applications store the ciphertext and wrapped DEK, not the KEK itself. This design also supports centralized access control, auditing, rotation, and lifecycle management for the key that protects the data keys.

See the Cloud KMS envelope encryption documentation and the Cloud KMS encryption algorithms documentation.

QUESTION NO: 30

You're developing the incident response plan for your company. You need to define the access strategy that your DevOps team will use when reviewing and investigating a deployment issue in your Google Cloud environment. There are two main requirements:

Least-privilege access must be enforced at all times.

The DevOps team must be able to access the required resources only during the deployment issue.

How should you grant access while following Google-recommended best practices?

A. Assign the Project Viewer Identity and Access Management (1AM) role to the DevOps team.

B. Create a custom 1AM role with limited list/view permissions, and assign it to the DevOps team.

C. Create a service account, and grant it the Project Owner 1AM role. Give the Service Account User Role on this service account to the DevOps team.

D. Create a service account, and grant it limited list/view permissions. Give the Service Account User Role on this service account to the DevOps team.

E. Create a Privileged Access Manager entitlement that grants the DevOps team the least-privilege role for a limited duration during deployment incidents.

ANSWER: E

Explanation:

Use Privileged Access Manager (PAM) to provide just-in-time, time-bound access during an incident. Create a PAM entitlement that makes the DevOps team eligible to receive only the narrowly scoped IAM role required for deployment investigation, ideally on the specific project, folder, or resource involved. Configure the entitlement with an access duration appropriate for incident response and, where required by the organization, an approval workflow. When a deployment issue occurs, an eligible DevOps engineer requests and activates the entitlement; PAM then creates a temporary IAM grant. The grant expires automatically when its configured duration ends, ensuring that access is not retained after the incident.

This approach directly supports least privilege because the role granted through the entitlement should contain only the permissions needed to inspect or remediate the deployment issue. It also provides auditable elevation requests and grants, helping incident responders and security teams review who obtained access, when it was activated, and for how long. Google documents PAM as its just-in-time privileged-access capability and recommends using entitlements to define eligible principals, roles, scope, duration, and optional approvals. See the Privileged Access Manager overview and guidance for creating PAM entitlements.

QUESTION NO: 31

A customer deployed an application on Compute Engine that takes advantage of the elastic nature of cloud computing.

How can you work with Infrastructure Operations Engineers to best ensure that Windows Compute Engine VMs are up to date with all the latest OS patches?

A. Build new base images when patches are available, and use a CI/CD pipeline to rebuild VMs, deploying incrementally.

B. Federate a Domain Controller into Compute Engine, and roll out weekly patches via Group Policy Object.

C. Use Deployment Manager to provision updated VMs into new serving Instance Groups (IGs).

D. Reboot all VMs during the weekly maintenance window and allow the StartUp Script to download the latest patches from the internet.

ANSWER: A

Explanation:

Build new base images when patches are available, and use a CI/CD pipeline to rebuild VMs, deploying incrementally. is the best approach because it applies an immutable-infrastructure model that aligns with elastic Compute Engine deployments. Infrastructure Operations Engineers can maintain a hardened, version-controlled Windows base image, validate newly released operating-system patches in a build and test pipeline, and publish an approved image version. The CI/CD process can then create replacement VM instances from that image and roll them out incrementally through managed instance groups or the application deployment process.

This provides consistent patch levels across the fleet rather than relying on each running VM to independently download, install, and successfully reboot after updates. It also supports repeatable testing, traceability of the exact image version deployed, controlled rollback to a known-good image if needed, and safe gradual rollout that preserves application availability. Replacing instances is particularly appropriate for an elastic workload, because instances are designed to be created and removed as capacity and releases change. Google recommends creating and managing custom images for consistent VM configuration, and managed instance groups support controlled rolling updates of VM instances. See Create custom images and Roll out updates to managed instance groups.

QUESTION NO: 32

Your DevOps team uses Packer to build Compute Engine images by using this process:

1 Create an ephemeral Compute Engine VM.

2 Copy a binary from a Cloud Storage bucket to the VM's file system.

3 Update the VM's package manager.

4 Install external packages from the internet onto the VM.

Your security team just enabled the organizational policy. consrraints/compure.vnExtemallpAccess. to restrict the usage of public IP Addresses on VMs. In response your DevOps team updated their scripts to remove public IP addresses on the Compute Engine VMs however the build pipeline is failing due to connectivity issues.

What should you do?

Choose 2 answers

A. Provision a Cloud NAT instance in the same VPC and region as the Compute Engine VM

B. Provision an HTTP load balancer with the VM in an unmanaged instance group to allow inbound connections from the internet to your VM.

C. Update the VPC routes to allow traffic to and from the internet.

D. Provision a Cloud VPN tunnel in the same VPC and region as the Compute Engine VM.

E. Enable Private Google Access on the subnet that the Compute Engine VM is deployed within.

ANSWER: A E

Explanation:

Provision a Cloud NAT instance in the same VPC and region as the Compute Engine VM and enable Private Google Access on the subnet that the Compute Engine VM is deployed within. Together, these controls restore the distinct outbound paths required by the image build while preserving the organizational restriction against external IP addresses on the ephemeral VM.

Cloud NAT provides source network address translation for eligible VM instances that have only internal IP addresses. It enables the VM to initiate outbound connections to public internet package repositories for package-manager updates and installation of external packages. Cloud NAT is a regional service and must be configured for the relevant VPC network and region, typically through a Cloud Router.

Private Google Access enables instances without external IP addresses to reach Google APIs and services, including Cloud Storage, using Google's network. This supports the build step that retrieves the binary from the Cloud Storage bucket without assigning an external address to the VM. The service is enabled at the subnet level, so the Packer-created VM must use that configured subnet. Google documents Cloud NAT at Cloud NAT overview and describes access to Google APIs without external IPs at Private Google Access.

QUESTION NO: 33

Your organization previously stored files in Cloud Storage by using Google Managed Encryption Keys (GMEK). but has recently updated the internal policy to require Customer Managed Encryption Keys (CMEK). You need to re-encrypt the files quickly and efficiently with minimal cost.

What should you do?

A. Encrypt the files locally, and then use gsutil to upload the files to a new bucket.

B. Copy the files to a new bucket with CMEK enabled in a secondary region

C. Reupload the files to the same Cloud Storage bucket specifying a key file by using gsutil.

D. Configure a CMEK as the bucket's default encryption key, then rewrite the objects in place.

ANSWER: D

Explanation:

Configure a CMEK as the bucket's default encryption key, then rewrite the objects in place is correct because Cloud Storage applies encryption to an object when the object is written. Setting a bucket's default Cloud KMS key ensures that subsequent writes to that bucket use the required CMEK. Rewriting each existing object then creates a new encrypted version of the object within Cloud Storage, using that configured CMEK, without requiring the data to be downloaded to and uploaded from an external client.

This approach is efficient because the rewrite occurs service-side and can retain the same bucket and object names. It avoids unnecessary client-side transfer time, local storage requirements, and network egress or ingestion workflow overhead. The process should be performed only after ensuring that the Cloud Storage service agent has permission to encrypt and decrypt with the selected Cloud KMS key, and that the key location is compatible with the bucket's location. Google documents that changing the encryption key for existing Cloud Storage objects requires rewriting them, while a bucket default key applies to new object writes. See Use customer-managed encryption keys and Rewrite objects.

QUESTION NO: 34

You have been tasked with configuring Security Command Center for your organization's Google Cloud environment. Your security team needs to receive alerts of potential crypto mining in the organization's compute environment and alerts for common Google Cloud misconfigurations that impact security. Which Security Command Center features should you use to configure these alerts? (Choose two.)

A. Event Threat Detection

B. Container Threat Detection

C. Security Health Analytics

D. Cloud Data Loss Prevention

E. Google Cloud Armor

ANSWER: A C

Explanation:

Event Threat Detection is the Security Command Center capability for detecting suspicious activity by continuously analyzing supported Google Cloud logs and producing findings. Its built-in detectors include detections for cryptocurrency mining activity, allowing the security team to receive findings when log events indicate that compute resources may be used for cryptomining. These findings can be routed through Security Command Center notification mechanisms and integrated with response workflows.

Security Health Analytics is the appropriate feature for identifying common configuration weaknesses that affect an organization's cloud security posture. It runs managed detectors against Google Cloud resource configurations and generates findings for issues such as overly permissive access, insecure network settings, insufficient logging configurations, and other security-relevant misconfigurations. This gives the team an ongoing view of configuration risks across supported projects and resources, rather than relying on ad hoc manual reviews.

Together, these capabilities address both requested alert types: Event Threat Detection provides threat-oriented monitoring for potential cryptomining, while Security Health Analytics provides posture-oriented findings for common Google Cloud security misconfigurations. See the Event Threat Detection overview and the Security Health Analytics overview.

QUESTION NO: 35

A BigQuery table contains employee data, including salary and government identification columns. Human resources administrators must be able to view the original values, while analysts must be able to query the table but see masked values in the sensitive columns.

Which two actions should you take? (Choose two.)

A. Assign policy tags to the sensitive columns.

B. Create and apply data masking policies to the policy tags.

C. Enable BigQuery table partitioning by employee department.

D. Grant the BigQuery Data Owner role to all analysts.

E. Export the table to Cloud Storage and encrypt it with a customer-managed encryption key.

ANSWER: A B

Explanation:

Assign **policy tags** to the sensitive BigQuery columns. Policy tags provide column-level access control and let the organization classify sensitive fields independently of table-level permissions. IAM can then grant access to the protected data only to the appropriate human resources users and administrators.

Create and apply a **data masking policy** to the policy tags. BigQuery data masking policies allow users who do not have fine-grained access to view a masked value instead of the original sensitive data. This allows analysts to run permitted queries while reducing unnecessary exposure of salary and identifier values.

See BigQuery column-level security and BigQuery dynamic data masking.

QUESTION NO: 36

You need to connect your organization's on-premises network with an existing Google Cloud environment that includes one Shared VPC with two subnets named Production and Non-Production. You are required to:

Use a private transport link.

Configure access to Google Cloud APIs through private API endpoints originating from on-premises environments.

Ensure that Google Cloud APIs are only consumed via VPC Service Controls.

What should you do?

A. 1. Set up a Cloud VPN link between the on-premises environment and Google Cloud.2. Configure private access using the restricted googleapis.com domains in on-premises DNS configurations.

B. 1. Set up a Partner Interconnect link between the on-premises environment and Google Cloud.2. Configure private access using the private.googleapis.com domains in on-premises DNS configurations.

C. 1. Set up a Direct Peering link between the on-premises environment and Google Cloud.2. Configure private access for both VPC subnets.

D. 1. Set up a Dedicated Interconnect link between the on-premises environment and Google Cloud.2. Configure private access using the restricted.googleapis.com domains in on-premises DNS configurations.

ANSWER: D

Explanation:

Set up a Dedicated Interconnect link between the on-premises environment and Google Cloud, then configure on-premises DNS to resolve applicable Google API names to `restricted.googleapis.com`. Dedicated Interconnect supplies private, physical connectivity from the organization's network to Google's network. A VLAN attachment connects that interconnect capacity to the Shared VPC, allowing on-premises systems to reach private Google API virtual IP addresses through the VPC rather than through public internet paths.

The `restricted.googleapis.com` endpoint is specifically intended for workloads that must access only Google APIs supported by VPC Service Controls. DNS resolution to its restricted virtual IP range, together with routing of that range across the VLAN attachment, directs API calls from the on-premises environment over the private hybrid connection. This endpoint therefore supports the requirement to make API access private while enforcing the intended VPC Service Controls-compatible API boundary. The DNS configuration must cover the required Google API hostnames, and the hybrid routing configuration must advertise or otherwise provide a route to the restricted API address range.

Google documents private access to Google APIs from on-premises hosts, including use of restricted VIPs, at Private access options for Google APIs. For the endpoint behavior and its VPC Service Controls use case, see Configure private Google access for on-premises hosts.

QUESTION NO: 37

Your company runs workloads on AWS that need to access Google Cloud resources. The company does not allow the creation or storage of long-lived Google Cloud service account keys outside Google Cloud.

Which two actions should you take to meet these requirements? (Choose two.)

A. Create a Workload Identity Pool and configure an AWS workload identity provider.

B. Grant the federated principal the Workload Identity User role on a dedicated service account.

C. Create a user-managed service account key and store it in AWS Secrets Manager.

D. Grant the AWS IAM role the Google Cloud Owner role on the target project.

E. Create a Cloud VPN tunnel so that the AWS workload can use the Compute Engine metadata server.

ANSWER: A B

Explanation:

Create a Workload Identity Pool and configure an AWS workload identity provider. Workload Identity Federation lets external workloads exchange credentials issued by a supported external identity provider, including AWS, for short-lived Google Cloud credentials. The AWS workload can authenticate using its AWS identity rather than a downloaded Google Cloud service account key.

Grant the federated principal permission to impersonate a dedicated Google Cloud service account by granting the Workload Identity User role. The workload can then obtain short-lived credentials for that service account, and the service account can

be granted only the resource permissions required by the AWS application. This supports least privilege and removes the need to manage long-lived service account private keys. See Workload Identity Federation and Workload Identity Federation with AWS.

QUESTION NO: 38

You want to limit the images that can be used as the source for boot disks. These images will be stored in a dedicated project.

What should you do?

A. Use the Organization Policy Service to create a compute.trustedimageProjects constraint on the organization level. List the trusted project as the whitelist in an allow operation.

B. Use the Organization Policy Service to create a compute.trustedimageProjects constraint on the organization level. List the trusted projects as the exceptions in a deny operation.

C. In Resource Manager, edit the project permissions for the trusted project. Add the organization as member with the role: Compute Image User.

D. In Resource Manager, edit the organization permissions. Add the project ID as member with the role: Compute Image User.

ANSWER: A

Explanation:

Use the Organization Policy Service to create a compute.trustedimageProjects constraint on the organization level. List the trusted project as the whitelist in an allow operation. This organization policy constraint is specifically designed to control which Google Cloud projects can supply Compute Engine images for boot disks. Applying the policy at the organization level establishes a consistent guardrail for all current and future projects in the organization, unless a more specific supported policy configuration applies. The allowed-values list should contain the dedicated image project, using its project resource format, so that VM boot disks may be created only from images hosted in that approved project. This helps enforce a centralized, controlled image pipeline: security teams can harden, validate, and maintain approved images in one location while preventing use of unapproved image sources elsewhere in the organization. The policy restricts image use; appropriate IAM permissions are still required for principals that need to access the approved images. Google documents this control as the trusted image projects organization policy constraint and shows configuring the trusted projects through an allowed-values policy. See Restrict access to images and Using organization policy constraints.