

EXAM PREPARATION MATERIAL

Practice files designed to help you understand exam patterns, question styles and the overall assessment format.

IMPORTANT NOTE



Please note that these files may not be up to date.
However, the questions will help you understand the exam format and typical question patterns.



Stay connected with our support team
for the latest file updates and exam-related information.

Question: 1

Your team needs to make sure that a Compute Engine instance does not have access to the internet or to any Google APIs or services.

Which two settings must remain disabled to meet these requirements? (Choose two.)

- A. Public IP
- B. IP Forwarding
- C. Private Google Access
- D. Static routes
- E. IAM Network User Role

Answer: AC

Explanation:

To ensure that a Compute Engine instance does not have access to the internet or to any Google APIs or services, you need to disable the following settings:

Public IP: Disabling the public IP address ensures that the instance does not have a direct connection to the internet.

Without a public IP address, the instance cannot be accessed from or communicate with the internet directly.

Private Google Access: Disabling Private Google Access ensures that the instance does not have access to Google APIs and services through the internal Google network. Private Google Access allows instances without a public IP to reach Google APIs and services using private IP addresses, but disabling it will block this path.

Disabling these settings will effectively isolate the instance from both the public internet and Google's internal API services.

Reference:

Google Cloud VPC Documentation - Overview

Configuring Private Google Access

Compute Engine Network Overview

Question: 2

Which two implied firewall rules are defined on a VPC network? (Choose two.)

- A. A rule that allows all outbound connections
- B. A rule that denies all inbound connections
- C. A rule that blocks all inbound port 25 connections
- D. A rule that blocks all outbound connections
- E. A rule that allows all inbound port 80 connections

Answer: AB

Explanation:

Implied IPv4 allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination

Implied IPv4 deny ingress rule. An ingress rule whose action is deny, source is 0.0.0.0/0, and priority is the lowest possible (65535) protects all instances by blocking incoming connections to them.

https://cloud.google.com/vpc/docs/firewalls?hl=en#default_firewall_rules

Question: 3

A customer needs an alternative to storing their plain text secrets in their source-code management (SCM) system.

How should the customer achieve this using Google Cloud Platform?

- A. Use Cloud Source Repositories, and store secrets in Cloud SQL.
- B. Encrypt the secrets with a Customer-Managed Encryption Key (CMEK), and store them in Cloud Storage.
- C. Run the Cloud Data Loss Prevention API to scan the secrets, and store them in Cloud SQL.
- D. Deploy the SCM to a Compute Engine VM with local SSDs, and enable preemptible VMs.

Answer: B

Explanation:

Storing secrets securely is crucial for maintaining the integrity and confidentiality of your applications. Here is how you can achieve this using Google Cloud Platform:

Encrypt the Secrets: Use Customer-Managed Encryption Keys (CMEK) to encrypt your secrets. CMEK

allows you to have greater control over the encryption keys used to protect your data. This ensures that even if the storage medium is compromised, the secrets remain protected by strong encryption. Store in Cloud Storage: Store the encrypted secrets in Google Cloud Storage. Cloud Storage is a secure and scalable object storage service. By using encrypted storage, you can ensure that the secrets are securely stored and can only be accessed by authorized entities.

This method provides a secure and managed way to store secrets, ensuring that they are not exposed in plain text within your source code management system.

Reference:

Customer-Managed Encryption Keys (CMEK)

Google Cloud Storage Security

Question: 4

Your team wants to centrally manage GCP IAM permissions from their on-premises Active Directory Service. Your team

wants to manage permissions by AD group membership.

What should your team do to meet these requirements?

- A. Set up Cloud Directory Sync to sync groups, and set IAM permissions on the groups.
- B. Set up SAML 2.0 Single Sign-On (SSO), and assign IAM permissions to the groups.
- C. Use the Cloud Identity and Access Management API to create groups and IAM permissions from Active Directory.
- D. Use the Admin SDK to create groups and assign IAM permissions from Active Directory.

Answer: A

Explanation:

"In order to be able to keep using the existing identity management system, identities need to be synchronized between AD and GCP IAM. To do so Google provides a tool called Cloud Directory Sync. This tool will read all identities in AD and replicate those within GCP. Once the identities have been replicated then it's possible to apply IAM permissions on the groups. After that you will configure SAML so Google can act as a service provider and either you ADFS or other third party tools like Ping or Okta will act as the identity provider. This way you effectively delegate the authentication from Google to something that is under your control."

Question: 5

When creating a secure container image, which two items should you incorporate into the build if possible? (Choose two.)

- A. Ensure that the app does not run as PID 1.
- B. Package a single app as a container.
- C. Remove any unnecessary tools not needed by the app.
- D. Use public container images as a base image for the app.
- E. Use many container image layers to hide sensitive information.

Answer: BC

Explanation:

When creating a secure container image, it is essential to follow best practices to minimize vulnerabilities and ensure the container operates as intended. Here are the two key practices: Package a Single App as a Container: By packaging only a single application within a container, you reduce complexity and potential attack surfaces. This practice aligns with the principle of single responsibility, ensuring each container has a clear and focused purpose.

Remove Any Unnecessary Tools: Any additional tools or software that are not required by the application should be removed from the container image. This minimizes the number of potential vulnerabilities and reduces the attack surface. A minimal container image also leads to smaller image sizes and faster deployment times.

These practices contribute to creating a more secure and efficient container image. Reference:

Question: 6

A customer needs to launch a 3-tier internal web application on Google Cloud Platform (GCP). The customer's internal compliance requirements dictate that end-user access may only be allowed if the traffic seems to originate from a specific known good CIDR. The customer accepts the risk that their application will only have SYN flood DDoS protection. They want to use GCP's native SYN flood protection.

Which product should be used to meet these requirements?

- A. Cloud Armor
- B. VPC Firewall Rules
- C. Cloud Identity and Access Management
- D. Cloud CDN

Answer: B

Explanation:

To ensure end-user access is only allowed if the traffic originates from a specific known good CIDR and to utilize GCP's native SYN flood protection, you can use the following product:

VPC Firewall Rules: By configuring VPC firewall rules, you can control traffic to and from your

instances based on IP address, protocol, and port. You can set rules to only allow traffic from a specific CIDR block, ensuring that only authorized traffic can reach your application.

Additionally, Google Cloud Platform provides built-in protections against SYN flood attacks, which are a type of DDoS attack.

These protections are part of the underlying infrastructure and do not require additional configuration.

Using VPC firewall rules will help you comply with the internal requirement of allowing access only from a specific CIDR and provide the necessary SYN flood DDoS protection.

Reference:

Google Cloud VPC Firewall Rules

Google Cloud DDoS Protection

Question: 7

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: AC

Explanation:

To connect Compute Engine instances within a Google Cloud Platform project to workloads running in a dedicated server room that can only be accessed from within the private company network, you can use the following approaches:

Cloud VPN: Cloud VPN securely connects your on-premises network to your Google Cloud Virtual Private Cloud (VPC) network through an IPsec VPN connection. This enables secure communication between your GCP instances and your on-premises workloads over the internet.

Cloud Interconnect: Cloud Interconnect provides direct physical connections between your on-premises network and Google's network. It offers higher bandwidth and lower latency compared to Cloud VPN, making it suitable for workloads that require fast and reliable connectivity.

Both Cloud VPN and Cloud Interconnect allow you to securely connect your on-premises environments to Google Cloud, ensuring that the workloads remain within the private company network.

Reference:

Cloud VPN Overview

Cloud Interconnect Overview

Question: 8

A customer implements Cloud Identity-Aware Proxy for their ERP system hosted on Compute Engine. Their security team wants to add a security layer so that the ERP systems only accept traffic from Cloud Identity-Aware Proxy.

What should the customer do to meet these requirements?

- A. Make sure that the ERP system can validate the JWT assertion in the HTTP requests.
- B. Make sure that the ERP system can validate the identity headers in the HTTP requests.
- C. Make sure that the ERP system can validate the x-forwarded-for headers in the HTTP requests.
- D. Make sure that the ERP system can validate the user's unique identifier headers in the HTTP requests.

Answer: A

Explanation:

Use Cryptographic Verification If there is a risk of IAP being turned off or bypassed, your app can check to make sure the identity information it receives is valid. This uses a third web request header added by IAP, called X-Goog-IAP-JWT-Assertion.

The value of the header is a cryptographically signed object that also contains the user identity data. Your application can verify the digital signature and use the data provided in this object to be certain that it was provided by IAP without alteration.

Question: 9

A company has been running their application on Compute Engine. A bug in the application allowed a malicious user to repeatedly execute a script that results in the Compute Engine instance crashing. Although the bug has been fixed, you want to get notified in case this hack re-occurs.

What should you do?

- A. Create an Alerting Policy in Stackdriver using a Process Health condition, checking that the number of executions of the script remains below the desired threshold. Enable notifications.
- B. Create an Alerting Policy in Stackdriver using the CPU usage metric. Set the threshold to 80% to be notified when the CPU usage goes above this 80%.
- C. Log every execution of the script to Stackdriver Logging. Create a User-defined metric in Stackdriver Logging on the logs, and create a Stackdriver Dashboard displaying the metric.
- D. Log every execution of the script to Stackdriver Logging. Configure BigQuery as a log sink, and create a BigQuery scheduled query to count the number of executions in a specific timeframe.

Answer: A

Explanation:

To monitor and get notified in case the script causing the Compute Engine instance to crash is executed again, you should create an Alerting Policy in Stackdriver (now known as Google Cloud Monitoring). The Process Health condition can be set to monitor the number of executions of the script and ensure it remains below the desired threshold. By enabling notifications, you will be alerted if this threshold is exceeded.

Step-by-Step:

Log Script Executions: Ensure that the script execution is logged.

Create a User-Defined Metric: Go to Google Cloud Console > Logging > Logs-based Metrics, and create a new user-defined metric that counts the number of times the script executes.

Set Up Alerting Policy:

Navigate to Google Cloud Console > Monitoring > Alerting.

Click on "Create Policy".

Add a condition and select "Logs-based Metric".

Configure the condition to trigger when the number of script executions exceeds the threshold.

Configure Notifications: Add notification channels (email, SMS, etc.) to the alerting policy.

Save and Test: Save the policy and test to ensure notifications are received when the script is executed beyond the threshold.

Reference::

Google Cloud Logging User-defined Metrics

Google Cloud Monitoring Alerting Policies

Question: 10

Your team needs to obtain a unified log view of all development cloud projects in your SIEM. The development projects are under the NONPROD organization folder with the test and pre-production projects. The development projects share the ABC-BILLING billing account with the rest of the organization.

Which logging export strategy should you use to meet the requirements?

- A. 1. Export logs to a Cloud Pub/Sub topic with folders/NONPROD parent and includeChildren property set to True in a dedicated SIEM project.
2. Subscribe SIEM to the topic.
- B. 1. Create a Cloud Storage sink with billingAccounts/ABC-BILLING parent and includeChildren property set to False in a dedicated SIEM project.
2. Process Cloud Storage objects in SIEM.
- C. 1. Export logs in each dev project to a Cloud Pub/Sub topic in a dedicated SIEM project.
2. Subscribe SIEM to the topic.
- D. 1. Create a Cloud Storage sink with a publicly shared Cloud Storage bucket in each project.
2. Process Cloud Storage objects in SIEM.

Answer: C

Explanation:

"Your team needs to obtain a unified log view of all development cloud projects in your SIEM" - This means we are ONLY interested in development projects. "The development projects are under the NONPROD organization folder with the test and pre-production projects" - We will need to filter out development from others i.e test and pre-prod. "The development projects share the ABC-BILLING billing account with the rest of the organization." - This is unnecessary information.

Question: 11

A customer needs to prevent attackers from hijacking their domain/IP and redirecting users to a malicious site through a man-in-the-middle attack.

Which solution should this customer use?

- A. VPC Flow Logs
- B. Cloud Armor
- C. DNS Security Extensions
- D. Cloud Identity-Aware Proxy

Answer: C

Explanation:

DNSSEC — use a DNS registrar that supports DNSSEC, and enable it. DNSSEC digitally signs DNS communication, making it more difficult (but not impossible) for hackers to intercept and spoof. Domain Name System Security Extensions (DNSSEC) adds security to the Domain Name System (DNS) protocol by enabling DNS responses to be validated. Having a trustworthy Domain Name System (DNS) that translates a domain name like www.example.com into its associated IP address is an increasingly important building block of today's web-based applications. Attackers can hijack this process of domain/IP lookup and redirect users to a malicious site through DNS hijacking and man-in-the-middle attacks. DNSSEC helps mitigate the risk of such attacks by cryptographically signing DNS records. As a result, it prevents attackers from issuing fake DNS

responses that may misdirect browsers to nefarious websites. <https://cloud.google.com/blog/products/gcp/dnssec-now-available-in-cloud-dns>

Question: 12

A customer deploys an application to App Engine and needs to check for Open Web Application Security Project (OWASP) vulnerabilities.

Which service should be used to accomplish this?

- A. Cloud Armor
- B. Google Cloud Audit Logs
- C. Cloud Security Scanner
- D. Forseti Security

Answer: C

Explanation:

Reference:: <https://cloud.google.com/security-scanner/>

Web Security Scanner supports categories in the OWASP Top Ten, a document that ranks and provides remediation guidance for the top 10 most critical web application security risks, as determined by the Open Web Application Security Project (OWASP).

https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview#detectors_and_compliance

Question: 13

A customer's data science group wants to use Google Cloud Platform (GCP) for their analytics workloads. Company policy dictates that all data must be company-owned and all user authentications must go through their own Security Assertion Markup Language (SAML) 2.0 Identity Provider (IdP). The Infrastructure Operations Systems Engineer was trying to set up Cloud Identity for the customer and realized that their domain was already being used by G Suite.

How should you best advise the Systems Engineer to proceed with the least disruption?

- A. Contact Google Support and initiate the Domain Contestation Process to use the domain name in your new Cloud Identity domain.
- B. Register a new domain name, and use that for the new Cloud Identity domain.
- C. Ask Google to provision the data science manager's account as a Super Administrator in the existing domain.
- D. Ask customer's management to discover any other uses of Google managed services, and work with the existing Super Administrator.

Answer: D

Explanation:

Since the domain is already being used by G Suite, the best course of action is to minimize disruption by discovering any existing uses of Google-managed services. Collaborate with the existing Super Administrator to align the setup with the company's requirements.

Step-by-Step:

Identify Existing Usage: Have the customer's management identify all current uses of the domain within Google-managed services.

Collaboration: Work closely with the existing Super Administrator of the domain.

Provision Required Accounts: Ask the Super Administrator to provision necessary accounts and permissions for the data science manager or other relevant personnel.

Integrate SAML IdP: Ensure that the existing domain integrates with the company's SAML 2.0 IdP for user authentication.

Set Up Cloud Identity: Configure Cloud Identity under the guidance of the Super Administrator without disrupting current services.

Reference::

Google Cloud Identity Administration

Google Support for Domain Issues

Question: 14

A business unit at a multinational corporation signs up for GCP and starts moving workloads into GCP. The business unit creates a Cloud Identity domain with an organizational resource that has hundreds of projects.

Your team becomes aware of this and wants to take over managing permissions and auditing the domain resources.

Which type of access should your team grant to meet this requirement?

- A. Organization Administrator
- B. Security Reviewer
- C. Organization Role Administrator
- D. Organization Policy Administrator

Answer: C

Explanation:

Here are the permissions available to organizationRoleAdmin

iam.roles.create

iam.roles.delete

iam.roles undelete

iam.roles.get

iam.roles.list
iam.roles.update
resourcemanager.projects.get
resourcemanager.projects.getIamPolicy
resourcemanager.projects.list
resourcemanager.organizations.get
resourcemanager.organizations.getIamPolicy

There are sufficient as per least privilege policy. You can do user management as well as auditing.

<https://cloud.google.com/iam/docs/understanding-custom-roles>

Question: 15

An application running on a Compute Engine instance needs to read data from a Cloud Storage bucket. Your team does not allow Cloud Storage buckets to be globally readable and wants to ensure the principle of least privilege.

Which option meets the requirement of your team?

- A. Create a Cloud Storage ACL that allows read-only access from the Compute Engine instance's IP address and allows the application to read from the bucket without credentials.
- B. Use a service account with read-only access to the Cloud Storage bucket, and store the credentials to the service account in the config of the application on the Compute Engine instance.
- C. Use a service account with read-only access to the Cloud Storage bucket to retrieve the credentials from the instance metadata.
- D. Encrypt the data in the Cloud Storage bucket using Cloud KMS, and allow the application to decrypt the data with the KMS key.

Answer: C

Explanation:

If the environment variable `GOOGLE_APPLICATION_CREDENTIALS` is set, ADC uses the service account key or configuration file that the variable points to. If the environment variable `GOOGLE_APPLICATION_CREDENTIALS` isn't set, ADC uses the service account that is attached to the resource that is running your code.

https://cloud.google.com/docs/authentication/production#passing_the_path_to_the_service_account_key_in_code

Question: 16

An organization's typical network and security review consists of analyzing application transit routes, request handling, and firewall rules. They want to enable their developer teams to deploy new applications without the overhead of this full review.

How should you advise this organization?

- A. Use Forseti with Firewall filters to catch any unwanted configurations in production.
- B. Mandate use of infrastructure as code and provide static analysis in the CI/CD pipelines to enforce policies.

- C. Route all VPC traffic through customer-managed routers to detect malicious patterns in production.
- D. All production applications will run on-premises. Allow developers free rein in GCP as their dev and QA platforms.

Answer: B

Explanation:

To enable developer teams to deploy new applications without the extensive overhead of network and security reviews, it's recommended to mandate the use of infrastructure as code (IaC) and enforce policies through static analysis in CI/CD pipelines. This approach ensures that security and compliance policies are checked automatically during the development process.

Step-by-Step:

Adopt IaC: Use tools like Terraform or Google Cloud Deployment Manager to manage infrastructure as code.

CI/CD Pipeline Integration: Integrate static analysis tools such as TFLint or Checkov in the CI/CD pipeline to enforce security policies.

Policy Definition: Define security policies and best practices that need to be adhered to in the code. **Automated Checks:**

Configure automated checks in the CI/CD pipeline to review code against these policies before deployment.

Monitor and Audit: Continuously monitor and audit deployed applications to ensure ongoing compliance.

Reference::

Infrastructure as Code on Google Cloud

[Static Analysis for Terraform](#)

[Checkov for IaC](#)

Question: 17

An employer wants to track how bonus compensations have changed over time to identify employee outliers and correct earning disparities. This task must be performed without exposing the sensitive compensation data for any individual and must be reversible to identify the outlier.

Which Cloud Data Loss Prevention API technique should you use to accomplish this?

- A. Generalization
- B. Redaction
- C. CryptoHashConfig
- D. CryptoReplaceFfxFpeConfig

Answer: D

Explanation:

De-identifying sensitive data Cloud Data Loss Prevention (DLP) can de-identify sensitive data in text content, including text stored in container structures such as tables. De-identification is the process of removing identifying information from data.

The API detects sensitive data such as personally

identifiable information (PII), and then uses a de-identification transformation to mask, delete, or otherwise obscure the data. For example, de-identification techniques can include any of the following: Masking sensitive data by partially or fully

replacing characters with a symbol, such as an asterisk (*) or hash (#). Replacing each instance of sensitive data with a token, or surrogate, string. Encrypting and replacing sensitive data using a randomly generated or pre-determined key. When you de-identify data using the CryptoReplaceFfxFpeConfig or CryptoDeterministicConfig infoType transformations, you can re-identify that data, as long as you have the CryptoKey used to originally de-identify the data.

<https://cloud.google.com/dlp/docs/deidentify-sensitive-data>

Question: 18

An organization adopts Google Cloud Platform (GCP) for application hosting services and needs guidance on setting up password requirements for their Cloud Identity account. The organization has a password policy requirement that corporate employee passwords must have a minimum number of characters.

Which Cloud Identity password guidelines can the organization use to inform their new requirements?

- A. Set the minimum length for passwords to be 8 characters.
- B. Set the minimum length for passwords to be 10 characters.
- C. Set the minimum length for passwords to be 12 characters.
- D. Set the minimum length for passwords to be 6 characters.

Answer: A

Explanation:

The minimum length for passwords in Cloud Identity can be set to 8 characters. This aligns with common security best practices for password policies, ensuring a basic level of complexity and security.

Step-by-Step:

Access Admin Console: Log in to the Google Admin console.

Navigate to Security Settings: Go to Security > Password Management.

Set Minimum Length: Set the minimum length for passwords to 8 characters.

Save Changes: Save the settings and ensure that all user accounts adhere to the new policy. Reference::

Google Cloud Identity Security Settings

Password Policy Best Practices

Question: 19

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:

Reference:: <https://cloud.google.com/kms/docs/envelope-encryption>

Envelope Encryption: <https://cloud.google.com/kms/docs/envelope-encryption>

Here are best practices for managing DEKs:

- Generate DEKs locally.
- When stored, always ensure DEKs are encrypted at rest.
- For easy access, store the DEK near the data that it encrypts.

The DEK is encrypted (also known as wrapped) by a key encryption key (KEK). The process of encrypting a key with another key is known as envelope encryption.

Here are best practices for managing KEKs:

- Store KEKs centrally. (KMS)
- Set the granularity of the DEKs they encrypt based on their use case. For example, consider a workload that requires multiple DEKs to encrypt the workload's data chunks. You could use a single KEK to wrap all DEKs that are responsible for that workload's encryption.
- Rotate keys regularly, and also after a suspected incident.

Question: 20

How should a customer reliably deliver Stackdriver logs from GCP to their on-premises SIEM system?

- A. Send all logs to the SIEM system via an existing protocol such as syslog.
- B. Configure every project to export all their logs to a common BigQuery DataSet, which will be queried by the SIEM system.
- C. Configure Organizational Log Sinks to export logs to a Cloud Pub/Sub Topic, which will be sent to the SIEM via Dataflow.
- D. Build a connector for the SIEM to query for all logs in real time from the GCP RESTful JSON APIs.

Answer: C

Explanation:

Scenarios for exporting Cloud Logging data: Splunk This scenario shows how to export selected logs from Cloud Logging to Pub/Sub for ingestion into Splunk. Splunk is a security information and event management (SIEM) solution that supports several ways of ingesting data, such as receiving streaming data out of Google Cloud through Splunk HTTP Event Collector (HEC) or by fetching data from Google Cloud APIs through Splunk Add-on for Google Cloud. Using the Pub/Sub to Splunk Dataflow template, you can natively forward logs and events from a Pub/Sub topic into Splunk HEC. If Splunk HEC is not available in your Splunk deployment, you can use the Add-on to collect the logs and events from the Pub/Sub topic.

<https://cloud.google.com/solutions/exporting-stackdriver-logging-for-splunk>

Question: 21

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: CD

Explanation:

App Engine ingress firewall rules are available, but egress rules are not currently available. Per requirements 1.2.1 and 1.3.4, you must ensure that all outbound traffic is authorized. SAQ A-EP and SAQ D-type merchants must provide compensating controls or use a different Google Cloud product. Compute Engine and GKE are the preferred alternatives.

<https://cloud.google.com/solutions/pci-dss-compliance-in-gcp>

Question: 22

A website design company recently migrated all customer sites to App Engine. Some sites are still in progress and should only be visible to customers and company employees from any location.

Which solution will restrict access to the in-progress sites?

- A. Upload an .htaccess file containing the customer and employee user accounts to App Engine.
- B. Create an App Engine firewall rule that allows access from the customer and employee networks and denies all other traffic.
- C. Enable Cloud Identity-Aware Proxy (IAP), and allow access to a Google Group that contains the customer and employee user accounts.
- D. Use Cloud VPN to create a VPN connection between the relevant on-premises networks and the company's GCP Virtual Private Cloud (VPC) network.

Answer: C

Explanation:

Cloud Identity-Aware Proxy (IAP) allows you to control access to your web applications running on Google Cloud. It ensures that only authenticated users who are part of a specified Google Group can access the application. Here's how you can restrict access to in-progress sites using IAP: Enable IAP: First, you need to enable Cloud IAP for your App Engine application.

This will require configuring OAuth consent and setting up necessary permissions.

Create a Google Group: Create a Google Group that includes all the customers and company employees who should have access to the in-progress sites.

Configure Access: Configure IAP to allow access only to members of the created Google Group. This involves setting up the necessary IAP policies and ensuring that only authenticated users in the group can access the application.

By using IAP, you ensure that the access control is centrally managed and only authorized users can view the in-progress sites from any location.

Reference:

Cloud Identity-Aware Proxy Documentation

Setting up IAP

Question: 23

When working with agents in a support center via online chat, an organization's customers often share pictures of their documents with personally identifiable information (PII). The organization that owns the support center is concerned that the PII is being stored in their databases as part of the regular chat logs they retain for review by internal or external analysts for customer service trend analysis.

Which Google Cloud solution should the organization use to help resolve this concern for the customer while still maintaining data utility?

- A. Use Cloud Key Management Service (KMS) to encrypt the PII data shared by customers before storing it for analysis.
- B. Use Object Lifecycle Management to make sure that all chat records with PII in them are discarded and not saved for analysis.
- C. Use the image inspection and redaction actions of the DLP API to redact PII from the images before storing them for analysis.
- D. Use the generalization and bucketing actions of the DLP API solution to redact PII from the texts before storing them for analysis.

Answer: C

Explanation:

To handle images containing personally identifiable information (PII) in chat logs while maintaining data utility, you can use Google Cloud's Data Loss Prevention (DLP) API. The DLP API provides capabilities to inspect and redact sensitive information from images. Here's how you can use it: Inspect Images: Use the DLP API to inspect images shared by customers for PII. This involves configuring the API to detect various types of sensitive information, such as names, social security numbers, and other PII.

Redact PII: Apply the redaction actions provided by the DLP API to remove or mask the PII in the images. The redaction can blur, mask, or replace sensitive information with placeholders, ensuring that the PII is not stored in the databases.

Store Redacted Images: Store the redacted images in your database for further analysis. This ensures that the sensitive information is not retained, addressing privacy concerns while still preserving the utility of the data for analysis.

By using the DLP API, the organization can effectively manage PII in customer-provided images, ensuring compliance with privacy regulations.

Reference:

Cloud DLP Documentation

Redacting Sensitive Data with DLP API

Question: 24

A company'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:

Rotating a user-managed Service Account key involves creating a new key, updating your application to use the new key, and then deleting the old key to maintain security. Here's the step-by-step process:

Create a New Key: Use the Google Cloud Console or `gcloud` command-line tool to create a new key for the service account. This generates a new key pair and provides you with the private key. `gcloud iam service-accounts keys create new-key-file.json --iam-account=YOUR_SERVICE_ACCOUNT_EMAIL`

Update Application: Update your application configuration to use the new key. This might involve replacing the old key file with the new one or updating the environment variables or configurations that point to the key file.

Delete the Old Key: Once you have confirmed that the application is working correctly with the new key, delete the old key from the service account to ensure it cannot be used for unauthorized access. `gcloud iam service-accounts keys delete OLD_KEY_ID --iam-account=YOUR_SERVICE_ACCOUNT_EMAIL`

This process ensures that your service account keys are regularly rotated, reducing the risk of key compromise.

Reference:

Managing Service Account Keys

Service Account Key Rotation

Question: 25

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:

Reference:: https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#centralize_network_control

Use Shared VPC to connect to a common VPC network. Resources in those projects can communicate with each other securely and efficiently across project boundaries using internal IPs. You can manage shared network resources, such as subnets, routes, and firewalls, from a central host project,

enabling you to apply and enforce consistent network policies across the projects.

Question: 26

An organization is migrating from their current on-premises productivity software systems to G Suite. Some network security controls were in place that were mandated by a regulatory body in their region for their previous on-premises system. The organization's risk team wants to ensure that network security controls are maintained and effective in G Suite. A security architect supporting this migration has been asked to ensure that network security controls are in place as part of the new shared responsibility model between the organization and Google Cloud.

What solution would help meet the requirements?

- A. Ensure that firewall rules are in place to meet the required controls.
- B. Set up Cloud Armor to ensure that network security controls can be managed for G Suite.
- C. Network security is a built-in solution and Google's Cloud responsibility for SaaS products like G Suite.
- D. Set up an array of Virtual Private Cloud (VPC) networks to control network security as mandated by the relevant regulation.

Answer: C

Explanation:

<https://gsuite.google.com/learn-more/security/security-whitepaper/page-1.html>

Shared responsibility "Security of the Cloud" - GCP is responsible for protecting the infrastructure that runs all of the services offered in the GCP Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run GCP Cloud services.

Question: 27

A customer's company has multiple business units. Each business unit operates independently, and each has their own engineering group. Your team wants visibility into all projects created within the company and wants to organize their Google Cloud Platform (GCP) projects based on different business units. Each business unit also requires separate sets of IAM permissions.

Which strategy should you use to meet these needs?

- A. Create an organization node, and assign folders for each business unit.
- B. Establish standalone projects for each business unit, using gmail.com accounts.
- C. Assign GCP resources in a project, with a label identifying which business unit owns the resource.
- D. Assign GCP resources in a VPC for each business unit to separate network access.

Answer: A

Explanation:

To organize GCP projects based on different business units and manage IAM permissions, you should create an organization node and assign folders for each business unit. This approach allows you to logically separate projects under folders and apply IAM policies at the folder level.

Step-by-Step:

Create Organization Node: Ensure that your GCP account is linked to an organization.

Create Folders for Business Units:

Navigate to the GCP Console > IAM & Admin > Resource Manager.

Create a folder for each business unit under the organization node.

Move Projects to Folders:

Move existing projects into the respective folders according to the business unit.

Set IAM Policies:

Assign IAM roles and permissions at the folder level to manage access for each business unit independently.

Monitor and Manage: Use Cloud Audit Logs and other GCP tools to monitor the activities and ensure compliance with the organization's policies.

Reference::

Creating and Managing Folders

Managing IAM Policies

Question: 28

A company has redundant mail servers in different Google Cloud Platform regions and wants to route customers to the nearest mail server based on location.

How should the company accomplish this?

- A. Configure TCP Proxy Load Balancing as a global load balancing service listening on port 995.
- B. Create a Network Load Balancer to listen on TCP port 995 with a forwarding rule to forward traffic based

on location.

- C. Use Cross-Region Load Balancing with an HTTP(S) load balancer to route traffic to the nearest region.
- D. Use Cloud CDN to route the mail traffic to the closest origin mail server based on client IP address.

Answer: A

Explanation:

<https://cloud.google.com/load-balancing/docs/tcp>

TCP Proxy Load Balancing is implemented on GFEs that are distributed globally. If you choose the Premium Tier of Network Service Tiers, a TCP proxy load balancer is global. In Premium Tier, you can deploy backends in multiple regions, and the load balancer automatically directs user traffic to the closest region that has capacity. If you choose the Standard Tier, a TCP proxy load balancer can only direct traffic among backends in a single region.

<https://cloud.google.com/load-balancing/docs/load-balancing-overview#tcp-proxy-load-balancing>

Question: 29

Your team sets up a Shared VPC Network where project co-vpc-prod is the host project. Your team has configured the firewall rules, subnets, and VPN gateway on the host project. They need to enable Engineering Group A to attach a Compute Engine instance to only the 10.1.1.0/24 subnet.

What should your team grant to Engineering Group A to meet this requirement?

- A. Compute Network User Role at the host project level.
- B. Compute Network User Role at the subnet level.
- C. Compute Shared VPC Admin Role at the host project level.
- D. Compute Shared VPC Admin Role at the service project level.

Answer: B

Explanation:

To enable Engineering Group A to attach a Compute Engine instance to a specific subnet (10.1.1.0/24) in a Shared VPC, you should grant the Compute Network User Role at the subnet level. This role allows users to use the subnetwork for their instances without giving them broader permissions at the project level.

Step-by-Step:

Identify the Subnet: Locate the subnet (10.1.1.0/24) in the host project.

Grant Role:

Navigate to the GCP Console > VPC network > VPC networks.

Select the Shared VPC host project and locate the specific subnet.

Click on "Edit" and go to the "IAM & Admin" section.

Assign the "Compute Network User" role to Engineering Group A at the subnet level.

Verification: Ensure that Engineering Group A can now attach Compute Engine instances to the specified subnet.

Reference::

Shared VPC Overview

Compute Network User Role

Question: 30

A company migrated their entire data/center to Google Cloud Platform. It is running thousands of instances across multiple projects managed by different departments. You want to have a historical record of what was running in Google Cloud Platform at any point in time.

What should you do?

- A. Use Resource Manager on the organization level.
- B. Use Forseti Security to automate inventory snapshots.
- C. Use Stackdriver to create a dashboard across all projects.
- D. Use Security Command Center to view all assets across the organization.

Answer: B

Explanation:

To maintain a historical record of what was running in Google Cloud Platform at any point in time, you should use Forseti Security to automate inventory snapshots. Forseti Security is an open-source toolkit that helps to automate security and compliance in GCP by taking inventory snapshots of GCP resources.

Step-by-Step:

Install Forseti Security:

Follow the installation guide to deploy Forseti Security on your GCP environment.

Configure Inventory:

Set up the inventory module in Forseti to capture and store snapshots of GCP resources.

Schedule Snapshots:

Use Forseti's configuration to schedule regular inventory snapshots.

Access Historical Data:

Review and access historical records through Forseti's dashboard or by querying the Forseti database.

Compliance and Monitoring: Use Forseti to ensure compliance and monitor changes over time. Reference::

[Forseti Security Overview](#)

Inventory Module

Question: 31

An organization is starting to move its infrastructure from its on-premises environment to Google Cloud Platform (GCP). The first step the organization wants to take is to migrate its current data backup and disaster recovery solutions to GCP for later analysis. The organization's production environment will remain on- premises for an indefinite time. The organization wants a scalable and cost-efficient solution.

Which GCP solution should the organization use?

- A. BigQuery using a data pipeline job with continuous updates
- B. Cloud Storage using a scheduled task and gsutil
- C. Compute Engine Virtual Machines using Persistent Disk

D. Cloud Datastore using regularly scheduled batch upload jobs

Answer: B

Explanation:

To migrate the current data backup and disaster recovery solutions to GCP while keeping the production environment on-premises, the most scalable and cost-efficient solution is using Google Cloud Storage with scheduled tasks and the `gsutil` command.

Setup Cloud Storage: Create a Cloud Storage bucket to store the backups.

Go to the Cloud Console and navigate to Cloud Storage.

Click "Create bucket" and follow the prompts to configure the storage bucket.

Install `gsutil`: Ensure `gsutil` is installed on the on-premises servers.

`gsutil` is a command-line tool for interacting with Cloud Storage.

Follow the installation guide here.

Create Backup Script: Write a script to upload data to Cloud Storage using `gsutil`.

```
#!/bin/bash gsutil -m cp -r /path/to/local/backup gs://your-bucket-name
```

Schedule Backup Task: Use a scheduling tool like `cron` on Linux to run the backup script at regular intervals.

Edit the crontab file with `crontab -e` and add an entry like:

Reference::

Cloud Storage Documentation

`gsutil` Documentation

Question: 32

You are creating an internal App Engine application that needs to access a user's Google Drive on the user's behalf. Your company does not want to rely on the current user's credentials. It also wants to follow Google-recommended practices.

What should you do?

A. Create a new Service account, and give all application users the role of Service Account User.

B. Create a new Service account, and add all application users to a Google Group. Give this group the role of Service Account User.

C. Use a dedicated G Suite Admin account, and authenticate the application's operations with these G Suite credentials.

D. Create a new service account, and grant it G Suite domain-wide delegation. Have the application use it to impersonate the user.

Answer: D

Explanation:

To access a user's Google Drive on their behalf without relying on the user's credentials and following Google-recommended practices, you should use a service account with domain-wide delegation.

Create a Service Account:

Go to the Cloud Console, navigate to IAM & Admin > Service Accounts.
Click "Create Service Account" and provide necessary details.
Grant Domain-Wide Delegation:
Edit the service account to enable "G Suite Domain-wide Delegation".

Download the JSON key file.

Configure API Access in G Suite:

Go to the Google Admin Console.

Navigate to Security > API Controls > Domain-wide Delegation.

Add a new API client and use the client ID from the service account.

Authorize the necessary API scopes (e.g., <https://www.googleapis.com/auth/drive>).

Implement in Application:

Use the Google API Client Library for the desired language.

Load the service account credentials and perform user impersonation to access Google Drive. Reference::

Domain-wide Delegation of Authority

Using OAuth 2.0 for Server to Server Applications

Question: 33

A customer wants to move their sensitive workloads to a Compute Engine-based cluster using Managed Instance Groups (MIGs). The jobs are bursty and must be completed quickly. They have a requirement to be able to manage and rotate the encryption keys.

Which boot disk encryption solution should you use on the cluster to meet this customer's requirements?

- A. Customer-supplied encryption keys (CSEK)
- B. Customer-managed encryption keys (CMEK) using Cloud Key Management Service (KMS)
- C. Encryption by default
- D. Pre-encrypting files before transferring to Google Cloud Platform (GCP) for analysis

Answer: B

Explanation:

For managing and rotating encryption keys in a Compute Engine-based cluster using Managed Instance Groups (MIGs), Customer-Managed Encryption Keys (CMEK) with Cloud KMS is the appropriate solution.

Set Up Cloud KMS:

Go to the Cloud Console and navigate to Security > Cryptographic Keys.

Create a keyring and a key.

Create and Use CMEK:

While creating or updating a Compute Engine instance, specify the CMEK key.

Example command:

```
gcloud compute instances create example-instance \
  --image-family=debian-9 \
  --image-project=debian-cloud \
  --boot-disk-kms-key=projects/[PROJECT_ID]/locations/global/keyRings/[KEY_RING]/cryptoKeys/[KEY]
```

Rotate Keys:

Rotate keys periodically using Cloud KMS by creating new key versions and updating the instances to use the new key

versions.

Reference::

Customer-Managed Encryption Keys (CMEK)

Using Customer-Managed Encryption Keys

Question: 34

Your company is using Cloud Dataproc for its Spark and Hadoop jobs. You want to be able to create, rotate, and destroy symmetric encryption keys used for the persistent disks used by Cloud Dataproc. Keys can be stored in the cloud.

What should you do?

- A. Use the Cloud Key Management Service to manage the data encryption key (DEK).
- B. Use the Cloud Key Management Service to manage the key encryption key (KEK).
- C. Use customer-supplied encryption keys to manage the data encryption key (DEK).
- D. Use customer-supplied encryption keys to manage the key encryption key (KEK).

Answer: B

Explanation:

This PD and bucket data is encrypted using a Google-generated data encryption key (DEK) and key encryption key (KEK). The CMEK feature allows you to create, use, and revoke the key encryption key (KEK). Google still controls the data encryption key (DEK). For more information on Google data encryption keys, see Encryption at Rest.

<https://cloud.google.com/dataproc/docs/concepts/configuring-clusters/customer-managed-encryption>

<https://codelabs.developers.google.com/codelabs/encrypt-and-decrypt-data-with-cloud-kms#0>

Question: 35

You are a member of the security team at an organization. Your team has a single GCP project with credit card payment processing systems alongside web applications and data processing systems. You want to reduce the scope of systems subject to PCI audit standards.

What should you do?

- A. Use multi-factor authentication for admin access to the web application.
- B. Use only applications certified compliant with PA-DSS.
- C. Move the cardholder data environment into a separate GCP project.
- D. Use VPN for all connections between your office and cloud environments.

Answer: C

Explanation:

To reduce the scope of systems subject to PCI audit standards, segregate the cardholder data environment (CDE) into a separate GCP project. This ensures that only the project containing the CDE will be subject to PCI DSS compliance, reducing the audit scope for other projects.

Create Separate GCP Project:

Go to the Cloud Console, navigate to IAM & Admin > Manage Resources.

Click "Create Project" and set up a new project for the CDE.

Migrate CDE:

Transfer the systems processing, storing, or transmitting cardholder data to the new project.

Apply PCI DSS Controls:

Implement PCI DSS required controls on the new project.

Use appropriate security measures such as firewalls, access controls, and encryption.

Reference::

Google Cloud and PCI DSS

Creating and Managing Projects

Question: 36

A retail customer allows users to upload comments and product reviews. The customer needs to make sure the text does not include sensitive data before the comments or reviews are published.

Which Google Cloud Service should be used to achieve this?

- A. Cloud Key Management Service
- B. Cloud Data Loss Prevention API
- C. BigQuery
- D. Cloud Security Scanner

Answer: B

Explanation:

To ensure user-uploaded comments and product reviews do not include sensitive data before publication, use the Cloud Data Loss Prevention (DLP) API.

Enable DLP API:

Go to the Cloud Console and navigate to APIs & Services > Library.

Search for "Data Loss Prevention API" and enable it.

Configure DLP API:

Create an inspection template specifying the types of sensitive data to detect.

Set up de-identification templates if you want to redact or mask sensitive data.

Implement DLP in Application:

Use the Google Cloud DLP Client Library for the desired programming language.

Send the text data to the DLP API for inspection before saving or publishing.

```
from google.cloud import dlp_v2
dlp_client = dlp_v2.DlpServiceClient()
parent =
```

```
f"projects/{project_id}"
item = {"value": "User comment text here"}
inspect_config = {"info_types": [{"name": "PERSON_NAME"}, {"name": "CREDIT_CARD_NUMBER"}]}
response = dlp_client.inspect_content(parent=parent,
inspect_config=inspect_config, item=item)
```

Reference::

Cloud Data Loss Prevention API Documentation

DLP API Client Libraries

Question: 37

A company allows every employee to use Google Cloud Platform. Each department has a Google Group, with all department members as group members. If a department member creates a new project, all members of that department should automatically have read-only access to all new project resources. Members of any other department should not have access to the project. You need to configure this behavior.

What should you do to meet these requirements?

- A. Create a Folder per department under the Organization. For each department's Folder, assign the Project Viewer role to the Google Group related to that department.
- B. Create a Folder per department under the Organization. For each department's Folder, assign the Project Browser role to the Google Group related to that department.
- C. Create a Project per department under the Organization. For each department's Project, assign the Project Viewer role to the Google Group related to that department.
- D. Create a Project per department under the Organization. For each department's Project, assign the Project Browser role to the Google Group related to that department.

Answer: A

Explanation:

To configure the behavior where each department member automatically has read-only access to all new project resources created by any department member, you should use Google Cloud's folder structure and IAM roles effectively. Here are the steps:

Create Folders for Departments: Create a folder under your Organization for each department. Folders help organize resources and provide a hierarchy for applying policies and permissions.

Assign IAM Roles to Google Groups: Assign the Project Viewer role to the Google Group associated with each department at the folder level. This ensures that all members of the group have the necessary permissions.

Inherited Permissions: When a department member creates a new project under their department's folder, the permissions assigned to the folder are inherited by the new project. Thus, all department members will automatically have read-only access to the project's resources.

Navigate to IAM & Admin in the GCP Console.

Select "Folders" from the left-hand menu.

For each department, create a new folder under the organization.

Select the newly created folder, and then go to the "Permissions" tab.

Click on "Add" to assign a new role.

Enter the email address of the Google Group for the department.

Assign the "Project Viewer" role to the group.

Access Restrictions: Since the permissions are applied at the folder level, only the members of the specific department's Google Group will have read-only access to the projects created in that folder. Other departments will not have access unless explicitly granted.

By following these steps, you ensure that department members have the required access to their respective projects without manual configuration for each new project.

Reference::

Google Cloud IAM Documentation

Google Cloud Resource Manager Documentation

Question: 38

A customer's internal security team must manage its own encryption keys for encrypting data on Cloud Storage and decides to use customer-supplied encryption keys (CSEK).

How should the team complete this task?

- A. Upload the encryption key to a Cloud Storage bucket, and then upload the object to the same bucket.
- B. Use the gsutil command line tool to upload the object to Cloud Storage, and specify the location of the encryption key.
- C. Generate an encryption key in the Google Cloud Platform Console, and upload an object to Cloud Storage using the specified key.
- D. Encrypt the object, then use the gsutil command line tool or the Google Cloud Platform Console to upload the object to Cloud Storage.

Answer: B

Explanation:

To use customer-supplied encryption keys (CSEK) for encrypting data on Cloud Storage, follow these steps:

Generate an Encryption Key: Generate a 256-bit AES encryption key. This key should be base64- encoded.

sh

Copy code

```
openssl rand -base64 32
```

Upload Object with CSEK: Use the gsutil command-line tool to upload the object to Cloud Storage, specifying the location of the encryption key using the -o option.

```
gsutil -o "GSUtil:encryption_key=<base64-encoded-key>" cp [LOCAL_OBJECT_PATH] gs://[BUCKET_NAME]/
```

Verify Encryption: After uploading the object, you can verify that it is encrypted using the provided CSEK by checking the object's metadata.

```
gsutil stat gs://[BUCKET_NAME]/[OBJECT_NAME]
```

Key Management: Ensure that the encryption key is securely stored and managed. It should not be hard-coded in scripts or applications.

By using the gsutil tool and specifying the encryption key, you ensure that the object is encrypted using the customer-

supplied encryption key during the upload process.

Reference::

Customer-Supplied Encryption Keys (CSEK) Documentation

gsutil Command Line Tool Documentation

Question: 39

A customer has 300 engineers. The company wants to grant different levels of access and efficiently manage IAM permissions between users in the development and production environment projects.

Which two steps should the company take to meet these requirements? (Choose two.)

- A. Create a project with multiple VPC networks for each environment.
- B. Create a folder for each development and production environment.
- C. Create a Google Group for the Engineering team, and assign permissions at the folder level.
- D. Create an Organizational Policy constraint for each folder environment.
- E. Create projects for each environment, and grant IAM rights to each engineering user.

Answer: BC

Explanation:

To manage IAM permissions efficiently for a large engineering team with different levels of access in development and production environments, follow these steps:

Create Separate Folders:

Create a folder for the development environment.

Create a folder for the production environment.

This allows you to organize projects and apply different policies and permissions to each environment.

Navigate to IAM & Admin in the GCP Console.

Select "Folders" from the left-hand menu.

Create a new folder named "Development".

Create a new folder named "Production".

Create Google Groups:

Create Google Groups for different teams within the engineering department (e.g., Development Team, Production Team).

This helps in managing permissions centrally.

Use the Google Admin Console to create groups.

Add relevant engineers to each group.

Assign Permissions at the Folder Level:

Assign appropriate IAM roles to the Google Groups at the folder level.

For example, grant Viewer role to the Development Team group for the development folder.

Grant Editor or more restrictive roles as required for the Production Team group for the production folder.

Select the development folder.

Go to the "Permissions" tab.

Click on "Add" and enter the email address of the Development Team Google Group.

Assign the "Viewer" role.

Repeat for the production folder, assigning appropriate roles to the Production Team Google Group. By following these steps, you create a clear separation between development and production environments and manage permissions efficiently using Google Groups and folders.

Reference::

Google Cloud IAM Documentation

Google Cloud Resource Manager Documentation

Question: 40

You want to evaluate GCP for PCI compliance. You need to identify Google's inherent controls.

Which document should you review to find the information?

- A. Google Cloud Platform: Customer Responsibility Matrix
- B. PCI DSS Requirements and Security Assessment Procedures
- C. PCI SSC Cloud Computing Guidelines
- D. Product documentation for Compute Engine

Answer: A

Explanation:

To evaluate Google Cloud Platform (GCP) for PCI compliance and identify Google's inherent controls, you should review the "Google Cloud Platform: Customer Responsibility Matrix". This document provides detailed information about the shared responsibility model, outlining the security controls managed by Google and those that are the responsibility of the customer.

Steps to access and use the document:

Access the Document:

Go to the Google Cloud compliance resource center.

Locate the "Customer Responsibility Matrix" for PCI DSS compliance.

Review Inherent Controls:

The document lists various controls and specifies whether they are managed by Google, the customer, or both.

It covers different aspects such as infrastructure security, data protection, and compliance requirements.

Analyze PCI Compliance:

Use the matrix to understand which PCI DSS requirements are inherently addressed by Google Cloud. Identify the controls you need to implement and manage as a customer to ensure full compliance. By reviewing this document, you can gain a comprehensive understanding of the inherent controls provided by Google Cloud and the responsibilities you must fulfill to achieve PCI compliance.

Reference::

Google Cloud Compliance Documentation

PCI DSS Compliance on Google Cloud

Question: 41

Your company runs a website that will store PII on Google Cloud Platform. To comply with data privacy regulations, this data can only be stored for a specific amount of time and must be fully deleted after this specific period. Data that has not yet

reached the time period should not be deleted. You want to automate the process of complying with this regulation.

What should you do?

- A. Store the data in a single Persistent Disk, and delete the disk at expiration time.
- B. Store the data in a single BigQuery table and set the appropriate table expiration time.
- C. Store the data in a Cloud Storage bucket, and configure the bucket's Object Lifecycle Management feature.
- D. Store the data in a single BigTable table and set an expiration time on the column families.

Answer: C

Explanation:

"To support common use cases like setting a Time to Live (TTL) for objects, retaining noncurrent versions of objects, or "downgrading" storage classes of objects to help manage costs, Cloud Storage offers the Object Lifecycle Management feature. This page describes the feature as well as the options available when using it. To learn how to enable Object Lifecycle Management, and for examples of lifecycle policies, see Managing Lifecycles."

<https://cloud.google.com/storage/docs/lifecycle>

Question: 42

A DevOps team will create a new container to run on Google Kubernetes Engine. As the application will be internet-facing, they want to minimize the attack surface of the container.

What should they do?

- A. Use Cloud Build to build the container images.
- B. Build small containers using small base images.
- C. Delete non-used versions from Container Registry.
- D. Use a Continuous Delivery tool to deploy the application.

Answer: B

Explanation:

To minimize the attack surface of the container for an internet-facing application running on Google Kubernetes Engine (GKE), the best practice is to build small containers using small base images. This approach helps in the following ways:

Reduce Vulnerabilities: Smaller base images contain fewer packages and dependencies, which minimizes the potential vulnerabilities that an attacker could exploit.

Improved Security: Using minimal base images such as distroless or Alpine Linux ensures that only the necessary components are included, reducing the attack surface significantly.

Easier Maintenance: Small containers are easier to maintain and update, ensuring that security patches can be applied quickly without dealing with unnecessary components.

Steps to Implement:

Choose a Minimal Base Image:

Use base images like gcr.io/distroless/base or alpine.

FROM gcr.io/distroless/base COPY myapp /myapp CMD ["/myapp"]

Optimize Container Image:

Remove unnecessary tools and libraries.

Use multi-stage builds to keep the final image small.

Regularly Update Base Images:

Keep the base images up-to-date with the latest security patches.

Reference::

[Distroless Images](#)

Best Practices for Building Containers

Question: 43

While migrating your organization's infrastructure to GCP, a large number of users will need to access GCP Console. The Identity Management team already has a well-established way to manage your users and want to keep using your existing Active Directory or LDAP server along with the existing SSO password.

What should you do?

- A. Manually synchronize the data in Google domain with your existing Active Directory or LDAP server.
- B. Use Google Cloud Directory Sync to synchronize the data in Google domain with your existing Active Directory or LDAP server.
- C. Users sign in directly to the GCP Console using the credentials from your on-premises Kerberos compliant identity provider.
- D. Users sign in using OpenID (OIDC) compatible IdP, receive an authentication token, then use that token to log in to the GCP Console.

Answer: B

Explanation:

To allow a large number of users to access the GCP Console while keeping the existing Active Directory or LDAP server for identity management, use Google Cloud Directory Sync (GCDS). Install GCDS:

Download and install Google Cloud Directory Sync from here.

Configure GCDS:

Set up the synchronization by specifying the LDAP server details and the Google domain.

Map the LDAP attributes to Google attributes to ensure user data is synchronized correctly.

Run Synchronization:

Perform an initial synchronization to populate the Google domain with existing users from the LDAP server.

Schedule regular synchronizations to keep the data up-to-date.

Benefits:

Automated Sync: Ensures that user data is consistently updated without manual intervention.

Secure Access: Users can log in to the GCP Console using their existing credentials, enhancing security and user experience.

Reference::

Google Cloud Directory Sync Documentation

GCDS Administration Guide

Question: 44

Your company is using GSuite and has developed an application meant for internal usage on Google App Engine. You need to make sure that an external user cannot gain access to the application even when an employee's password has been compromised.

What should you do?

- A. Enforce 2-factor authentication in GSuite for all users.
- B. Configure Cloud Identity-Aware Proxy for the App Engine Application.
- C. Provision user passwords using GSuite Password Sync.
- D. Configure Cloud VPN between your private network and GCP.

Answer: B

Explanation:

To ensure that an external user cannot gain access to an internal application on Google App Engine even if an employee's password is compromised, configure Cloud Identity-Aware Proxy (IAP).

Enable IAP:

Go to the Cloud Console, navigate to the App Engine application, and select "Identity-Aware Proxy".

Enable IAP for the application.

Configure Access Policies:

Set up access policies to restrict who can access the application.

Use IAM roles to grant access only to specific users or groups.

Enforce Authentication:

IAP enforces Google authentication, ensuring that users must log in with their GSuite credentials.

Enable Multi-Factor Authentication (MFA):

Enforce 2FA for all GSuite users to add an extra layer of security.

Advantages:

Protection against Compromised Credentials: Even if passwords are compromised, attackers cannot access the application without passing IAP authentication.

Centralized Access Management: Easily manage and monitor access through IAM and IAP policies.

Reference::

Identity-Aware Proxy Overview

Setting up IAP

Question: 45

A large financial institution is moving its Big Data analytics to Google Cloud Platform. They want to have maximum control over the encryption process of data stored at rest in BigQuery.

What technique should the institution use?

- A. Use Cloud Storage as a federated Data Source.
- B. Use a Cloud Hardware Security Module (Cloud HSM).
- C. Customer-managed encryption keys (CMEK).
- D. Customer-supplied encryption keys (CSEK).

Answer: C

Explanation:

If you want to manage the key encryption keys used for your data at rest, instead of having Google manage the keys, use Cloud Key Management Service to manage your keys. This scenario is known as customer-managed encryption keys (CMEK).

<https://cloud.google.com/bigquery/docs/encryption-at-rest>

Reference:: <https://cloud.google.com/bigquery/docs/encryption-at-rest>

Question: 46

A company is deploying their application on Google Cloud Platform. Company policy requires longterm data to be stored using a solution that can automatically replicate data over at least two geographic places.

Which Storage solution are they allowed to use?

- A. Cloud Bigtable
- B. Cloud BigQuery
- C. Compute Engine SSD Disk
- D. Compute Engine Persistent Disk

Answer: A

Explanation:

Cloud Bigtable is a fully managed NoSQL database service designed to handle large analytical and operational workloads. One of its key features is the ability to replicate data across multiple geographic locations automatically, ensuring high availability and resilience. Here's a detailed explanation:

Replication: Cloud Bigtable supports multi-cluster routing and replication across different geographic regions. This means that data can be replicated across multiple zones within a region or even across regions, providing geo-redundancy.

Automatic Handling: Once configured, Bigtable automatically manages replication without requiring manual intervention. This is in line with the company's policy for long-term data storage that necessitates automatic replication over at least two geographic places.

Use Case Suitability: Bigtable is ideal for applications that require low-latency access to large amounts of data, which makes it suitable for various use cases including analytical applications, IoT, and financial data processing.

Configuration: Setting up replication involves creating instances in multiple zones and configuring them to replicate data. Google Cloud's management interface and APIs make this straightforward to configure and monitor.

Reference::

Google Cloud Bigtable Documentation

Google Cloud Storage Options

Question: 47

A large e-retailer is moving to Google Cloud Platform with its ecommerce website. The company wants to ensure payment information is encrypted between the customer's browser and GCP when the customers checkout online.

What should they do?

- A. Configure an SSL Certificate on an L7 Load Balancer and require encryption.
- B. Configure an SSL Certificate on a Network TCP Load Balancer and require encryption.
- C. Configure the firewall to allow inbound traffic on port 443, and block all other inbound traffic.
- D. Configure the firewall to allow outbound traffic on port 443, and block all other outbound traffic.

Answer: A

Explanation:

To ensure that payment information is encrypted between the customer's browser and Google Cloud Platform during checkout, the company should configure an SSL certificate on an L7 (Layer 7) Load Balancer. Here's why this is the best solution:

SSL/TLS Termination: An L7 Load Balancer can handle SSL/TLS termination, which means it can decrypt HTTPS traffic, offloading the work from the backend servers. This is essential for handling encrypted connections securely.

HTTPS Configuration: By configuring an SSL certificate, the load balancer ensures that all traffic between the customer's browser and the application is encrypted using HTTPS.

Security Best Practices: Using an L7 Load Balancer with an SSL certificate aligns with best practices for securing web applications, particularly for e-commerce sites handling sensitive payment information.

Managed Certificates: Google Cloud offers managed SSL certificates, which simplifies the process of obtaining, deploying, and renewing SSL certificates.

Implementation Steps:

Obtain an SSL certificate.

Configure the L7 Load Balancer in the GCP Console.

Associate the SSL certificate with the load balancer.

Ensure that the backend services are configured to handle HTTPS traffic.

Reference::

Google Cloud Load Balancing Documentation

Setting up HTTPS Load Balancing

Question: 48

Applications often require access to "secrets" - small pieces of sensitive data at build or run time. The administrator managing these secrets on GCP wants to keep a track of "who did what, where, and when?" within their GCP projects.

Which two log streams would provide the information that the administrator is looking for? (Choose two.)

- A. Admin Activity logs
- B. System Event logs
- C. Data Access logs
- D. VPC Flow logs
- E. Agent logs

Answer: AC

Explanation:

To keep track of "who did what, where, and when?" within GCP projects, the administrator should focus on Admin Activity logs and Data Access logs. Here's a detailed explanation of why these two log streams are essential:

Admin Activity Logs:

These logs capture administrative actions performed in your Google Cloud resources. This includes actions like creating, modifying, or deleting resources.

Admin Activity logs provide detailed information about the user who performed the action, the resource that was affected, the action performed, and the timestamp.

Data Access Logs:

These logs capture read and write operations on data within your Google Cloud services. This includes actions like accessing or modifying data stored in databases, storage buckets, etc. Data Access logs help track the access patterns of users and services to sensitive data, providing insights into who accessed which data and when.

Steps to Enable and Access Logs:

Navigate to the Google Cloud Console.

Go to Logging in the left-hand menu.

Enable Admin Activity and Data Access logs if not already enabled.

Use Logs Explorer to filter and view specific logs based on your requirements.

By monitoring both Admin Activity and Data Access logs, administrators can gain comprehensive visibility into the actions performed on their GCP resources and data, ensuring robust security and compliance tracking.

Reference::

Google Cloud Logging Documentation

Audit Logs Overview

Question: 49

You are in charge of migrating a legacy application from your company datacenters to GCP before the current maintenance contract expires. You do not know what ports the application is using and no documentation is available for you to check. You want to complete the migration without putting your environment at risk.

What should you do?

- A. Migrate the application into an isolated project using a "Lift & Shift" approach. Enable all internal TCP traffic using VPC Firewall rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.
- B. Migrate the application into an isolated project using a "Lift & Shift" approach in a custom network. Disable all traffic within the VPC and look at the Firewall logs to determine what traffic should be allowed for the application to work properly.

- C. Refactor the application into a micro-services architecture in a GKE cluster. Disable all traffic from outside the cluster using Firewall Rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.
- D. Refactor the application into a micro-services architecture hosted in Cloud Functions in an isolated project. Disable all traffic from outside your project using Firewall Rules. Use VPC Flow logs to determine what traffic should be allowed for the application to work properly.

Answer: A

Explanation:

To migrate a legacy application to GCP without knowing what ports it uses and ensuring the environment is secure, the best approach is to use a "Lift & Shift" method in an isolated project and analyze the traffic using VPC Flow logs. Here's a step-by-step explanation:

Isolated Project:

Create a new, isolated project within your GCP environment to host the legacy application. This isolation ensures that any potential misconfigurations do not affect other projects.

Lift & Shift:

Migrate the application as-is (lift and shift) to the new isolated project. This involves moving the application without altering its architecture.

Enable Internal TCP Traffic:

Configure VPC Firewall rules to allow all internal TCP traffic within the VPC network. This step ensures that the application components can communicate internally without interruption. Use VPC Flow Logs:

Enable VPC Flow logs to capture information about the traffic to and from your application. VPC Flow logs provide details about the source, destination, port, and protocol of the traffic.

Analyze Traffic:

Analyze the VPC Flow logs to identify the necessary ports and protocols used by the application.

Based on this analysis, create specific firewall rules to allow only the required traffic, thereby tightening security.

Implementation Steps:

Navigate to the VPC network section in the GCP Console.

Create a new VPC or use an existing one, and configure firewall rules to allow internal TCP traffic.

Enable VPC Flow logs from the VPC network settings.

Migrate your application to the new project.

Monitor and analyze the VPC Flow logs to refine your firewall rules.

By following these steps, you can safely migrate the application, understand its network requirements, and secure it appropriately in the new GCP environment.

Reference::

Google Cloud VPC Documentation

VPC Flow Logs Documentation

Question: 50

Your company has deployed an application on Compute Engine. The application is accessible by clients on port 587. You need to balance the load between the different instances running the application. The connection should be secured using TLS, and terminated by the Load Balancer. What type of Load Balancing should you use?

- A. Network Load Balancing
- B. HTTP(S) Load Balancing
- C. TCP Proxy Load Balancing
- D. SSL Proxy Load Balancing

Answer: D

Explanation:

<https://cloud.google.com/load-balancing/docs/ssl> - SSL Proxy Load Balancing is a reverse proxy load balancer that distributes SSL traffic coming from the internet to virtual machine (VM) instances in your Google Cloud VPC network.

Reference:: <https://cloud.google.com/load-balancing/docs/ssl/>

Question: 51

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:

Objective: You want to limit the images that can be used as the source for boot disks to a set of images stored in a dedicated project.

Solution: Use the Organization Policy Service.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the Organization Policies page.

Step 3: Create a new policy by clicking on "Create Policy".

Step 4: Select the constraint compute.trustedimageProjects.

Step 5: Set the policy to ALLOW and specify the project ID where the trusted images are stored in the whitelist.

Step 6: Save and apply the policy.

By creating a compute.trustedimageProjects constraint at the organization level and specifying the trusted project in the

allow list, you ensure that only images from this project can be used for boot disks across the organization.

Reference::

GCP Organization Policy Service Documentation

Compute Trusted Image Projects Constraint

Question: 52

Your team needs to prevent users from creating projects in the organization. Only the DevOps team should be allowed to create projects on behalf of the requester.

Which two tasks should your team perform to handle this request? (Choose two.)

- A. Remove all users from the Project Creator role at the organizational level.
- B. Create an Organization Policy constraint, and apply it at the organizational level.
- C. Grant the Project Editor role at the organizational level to a designated group of users.
- D. Add a designated group of users to the Project Creator role at the organizational level.
- E. Grant the billing account creator role to the designated DevOps team.

Answer: AD

Explanation:

Objective: Prevent users from creating projects while allowing only the DevOps team to create projects.

Solution: Modify IAM roles and permissions.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the IAM & Admin page.

Step 3: At the organizational level, find and remove all users from the Project Creator role.

Step 4: Create or identify a group for the DevOps team.

Step 5: Assign the Project Creator role to the DevOps team group at the organizational level.

By removing all users from the Project Creator role and granting it only to the DevOps team, you ensure that only the designated team can create projects.

Reference::

GCP IAM Documentation

Project Creator Role

Question: 53

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:

Compute Engine doesn't automatically update the OS or the software on your deployed instances.

You will need to patch or update your deployed Compute Engine instances when necessary. However, in the cloud it is not recommended that you patch or update individual running instances. Instead it is best to patch the image that was used to launch the instance and then replace each affected instance with a new copy.

Question: 54

Your team needs to make sure that their backend database can only be accessed by the frontend application and no other instances on the network.

How should your team design this network?

- A. Create an ingress firewall rule to allow access only from the application to the database using firewall tags.
- B. Create a different subnet for the frontend application and database to ensure network isolation.
- C. Create two VPC networks, and connect the two networks using Cloud VPN gateways to ensure network isolation.
- D. Create two VPC networks, and connect the two networks using VPC peering to ensure network isolation.

Answer: A

Explanation:

"However, even though it is possible to use tags for target filtering in this manner, we recommend that you use service accounts where possible. Target tags are not access-controlled and can be

changed by someone with the instanceAdmin role while VMs are in service. Service accounts are access-controlled, meaning that a specific user must be explicitly authorized to use a service account. There can only be one service account per instance, whereas there can be multiple tags. Also, service accounts assigned to a VM can only be changed when the VM is stopped"

Question: 55

An organization receives an increasing number of phishing emails.

Which method should be used to protect employee credentials in this situation?

- A. Multifactor Authentication
- B. A strict password policy
- C. Captcha on login pages
- D. Encrypted emails

Answer: A

Explanation:

<https://cloud.google.com/blog/products/g-suite/7-ways-admins-can-help-secure-accounts-against-phishing-g-suite>

<https://www.duocircle.com/content/email-security-services/email-security-in-cryptography#:~:text=Customer%20Login->

[,Email%20Security%20In%20Cryptography%20Is%20One%20Of%20The%20Most,Measures%20To%20Prevent%20Phishing%20Attempts&text=Cybercriminals%20love%20emails%20the%20most,networks%20all%20over%20the%20world.](https://www.duocircle.com/content/email-security-services/email-security-in-cryptography#:~:text=Customer%20Login-Email%20Security%20In%20Cryptography%20Is%20One%20Of%20The%20Most,Measures%20To%20Prevent%20Phishing%20Attempts&text=Cybercriminals%20love%20emails%20the%20most,networks%20all%20over%20the%20world.)

Question: 56

A customer is collaborating with another company to build an application on Compute Engine. The customer is building the application tier in their GCP Organization, and the other company is building the storage tier in a different GCP Organization. This is a 3-tier web application. Communication between portions of the application must not traverse the public internet by any means.

Which connectivity option should be implemented?

- A. VPC peering
- B. Cloud VPN
- C. Cloud Interconnect
- D. Shared VPC

Answer: A

Explanation:

Objective: Ensure private communication between application tiers in different GCP Organizations. Solution: Use VPC peering to enable private communication without traversing the public internet. Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the VPC Network Peering page.

Step 3: Create a new VPC peering connection in the project hosting the application tier.

Step 4: Specify the VPC network in the other organization (hosting the storage tier) to peer with.

Step 5: Accept the peering request in the other project.

Step 6: Configure the necessary routes and firewall rules to allow traffic between the peered VPC networks.

VPC peering allows you to connect two VPC networks privately and directly, ensuring that traffic between them does not traverse the public internet.

Reference::

GCP VPC Peering Documentation

VPC Network Peering Guide

Question: 57

Your team wants to make sure Compute Engine instances running in your production project do not have public IP addresses. The frontend application Compute Engine instances will require public IPs. The product engineers have the Editor role to modify resources. Your team wants to enforce this requirement.

How should your team meet these requirements?

- A. Enable Private Access on the VPC network in the production project.
- B. Remove the Editor role and grant the Compute Admin IAM role to the engineers.
- C. Set up an organization policy to only permit public IPs for the front-end Compute Engine instances.
- D. Set up a VPC network with two subnets: one with public IPs and one without public IPs.

Answer: C

Explanation:

Objective: Ensure only front-end Compute Engine instances have public IPs, while others do not.

Solution: Use an organization policy to enforce this requirement.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the Organization Policies page.

Step 3: Create a new policy with the constraint `constraints/compute.requireOsLogin` (or similar constraint to manage public IPs).

Step 4: Define the conditions to allow public IPs only for the front-end instances.

Step 5: Apply the policy to the organization or specific projects as necessary.

By setting up an organization policy with specific conditions, you can control which instances are allowed to have public IPs based on their role or other attributes.

Reference::

GCP Organization Policies Documentation

Compute Engine Network Configurations

Question: 58

Which two security characteristics are related to the use of VPC peering to connect two VPC networks? (Choose two.)

- A. Central management of routes, firewalls, and VPNs for peered networks
- B. Non-transitive peered networks; where only directly peered networks can communicate
- C. Ability to peer networks that belong to different Google Cloud Platform organizations
- D. Firewall rules that can be created with a tag from one peered network to another peered network
- E. Ability to share specific subnets across peered networks

Answer: BC

Explanation:

Objective: Understand the security characteristics of VPC peering.

Security Characteristics:

Non-transitive Peering: VPC peering connections are non-transitive. This means that peering is strictly between two VPC networks. If VPC A is peered with VPC B, and VPC B is peered with VPC C, VPC A cannot communicate with VPC C unless a direct peering connection is established.

Inter-Organization Peering: VPC peering allows you to connect VPC networks across different Google Cloud Platform organizations, facilitating private communication between distinct organizational units.

These characteristics ensure controlled and secure connectivity between VPC networks while preventing unintended data exposure.

Reference::

GCP VPC Peering Documentation

VPC Network Peering Overview

Question: 59

A patch for a vulnerability has been released, and a DevOps team needs to update their running containers in Google Kubernetes Engine (GKE).

How should the DevOps team accomplish this?

- A. Use Puppet or Chef to push out the patch to the running container.
- B. Verify that auto upgrade is enabled; if so, Google will upgrade the nodes in a GKE cluster.
- C. Update the application code or apply a patch, build a new image, and redeploy it.
- D. Configure containers to automatically upgrade when the base image is available in Container Registry.

Answer: C

Explanation:

When a vulnerability patch is released for a running container in Google Kubernetes Engine (GKE), the recommended approach is to update the application code or apply the patch directly to the codebase. Then, a new container image should be built incorporating these changes. After building the new image, it should be deployed to replace the running containers.

This method ensures that the containers run the updated, secure code.

Steps:

Update Application Code: Modify the application code or dependencies to incorporate the vulnerability patch.

Build New Image: Use a tool like Docker to build a new container image with the updated code. **Push New Image:** Push the new container image to the Container Registry.

Update Deployments: Update the Kubernetes deployment to use the new image. This can be done by modifying the image tag in the deployment YAML file.

Redeploy Containers: Apply the updated deployment configuration using `kubectl apply -f <deployment-file>.yaml`, which will

redeploy the containers with the new image.

Reference::

Google Cloud: Container security

Kubernetes: Updating an application

Question: 60

A company is running their webshop on Google Kubernetes Engine and wants to analyze customer transactions in BigQuery.

You need to ensure that no credit card numbers are stored in BigQuery

What should you do?

- A. Create a BigQuery view with regular expressions matching credit card numbers to query and delete affected rows.
- B. Use the Cloud Data Loss Prevention API to redact related infoTypes before data is ingested into BigQuery.
- C. Leverage Security Command Center to scan for the assets of type Credit Card Number in BigQuery.
- D. Enable Cloud Identity-Aware Proxy to filter out credit card numbers before storing the logs in BigQuery.

Answer: B

Explanation:

<https://cloud.google.com/bigquery/docs/scan-with-dlp>

Cloud Data Loss Prevention API allows to detect and redact or remove sensitive data before the comments or reviews are published. Cloud DLP will read information from BigQuery, Cloud Storage or Datastore and scan it for sensitive data.

Question: 61

A customer wants to deploy a large number of 3-tier web applications on Compute Engine.

How should the customer ensure authenticated network separation between the different tiers of the application?

- A. Run each tier in its own Project, and segregate using Project labels.
- B. Run each tier with a different Service Account (SA), and use SA-based firewall rules.
- C. Run each tier in its own subnet, and use subnet-based firewall rules.
- D. Run each tier with its own VM tags, and use tag-based firewall rules.

Answer: B

Explanation:

"Isolate VMs using service accounts when possible" "even though it is possible to use tags for target filtering in this manner, we recommend that you use service accounts where possible. Target tags are not access-controlled and can be changed by someone with the instanceAdmin role while VMs are in service. Service accounts are access-controlled, meaning that a specific user must be explicitly authorized to use a service account. There can only be one service account per instance, whereas there can be multiple tags. Also, service accounts assigned to a VM can only be changed when the VM is stopped."

<https://cloud.google.com/solutions/best-practices-vpc-design#isolate-vms-service-accounts>

Question: 62

A manager wants to start retaining security event logs for 2 years while minimizing costs. You write a filter to select the appropriate log entries.

Where should you export the logs?

- A. BigQuery datasets
- B. Cloud Storage buckets
- C. StackDriver logging
- D. Cloud Pub/Sub topics

Answer: B

Explanation:

To retain security event logs for 2 years while minimizing costs, exporting the logs to Cloud Storage buckets is the most cost-effective solution. Cloud Storage provides scalable and durable storage at a lower cost compared to BigQuery, which is more suited for analytics and querying, or Cloud Pub/Sub, which is designed for messaging and stream processing.

Steps:

Create a Cloud Storage Bucket: Set up a new Cloud Storage bucket configured with appropriate retention policies.

Set Up Log Export: Use the Google Cloud Console or gcloud command-line tool to create a sink that exports the selected log entries to the Cloud Storage bucket.

Configure Retention Policy: Set the retention policy on the Cloud Storage bucket to ensure that logs are retained for the required period of 2 years.

Reference::

Google Cloud: Exporting logs

Google Cloud Storage pricing

Question: 63

For compliance reasons, an organization needs to ensure that in-scope PCI Kubernetes Pods reside on “in- scope” Nodes only. These Nodes can only contain the “in-scope” Pods.

How should the organization achieve this objective?

- A. Add a nodeSelector field to the pod configuration to only use the Nodes labeled inscope: true.
- B. Create a node pool with the label inscope: true and a Pod Security Policy that only allows the Pods to run on Nodes with that label.
- C. Place a taint on the Nodes with the label inscope: true and effect NoSchedule and a toleration to match in the Pod configuration.
- D. Run all in-scope Pods in the namespace “in-scope-pci”.

Answer: A

Explanation:

nodeSelector is the simplest recommended form of node selection constraint. You can add the nodeSelector field to your Pod specification and specify the node labels you want the target node to have. Kubernetes only schedules the Pod onto nodes that have each of the labels you specify. => <https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/#nodeselector> Tolerations are applied to pods. Tolerations allow the scheduler to schedule pods with matching taints. Tolerations allow scheduling but don't guarantee scheduling: the scheduler also evaluates other parameters as part of its function. => <https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/>

Question: 64

In an effort for your company messaging app to comply with FIPS 140-2, a decision was made to use GCP compute and network services. The messaging app architecture includes a Managed Instance Group (MIG) that controls a cluster of Compute Engine instances. The instances use Local SSDs for data caching and UDP for instance-to-instance communications. The app development team is willing to make any changes necessary to comply with the standard

Which options should you recommend to meet the requirements?

- A. Encrypt all cache storage and VM-to-VM communication using the BoringCrypto module.
- B. Set Disk Encryption on the Instance Template used by the MIG to customer-managed key and use BoringSSL for all data transit between instances.
- C. Change the app instance-to-instance communications from UDP to TCP and enable BoringSSL on clients' TLS connections.
- D. Set Disk Encryption on the Instance Template used by the MIG to Google-managed Key and use BoringSSL library on all instance-to-instance communications.

Answer: B

Explanation:

To comply with FIPS 140-2 for the messaging app, you need to ensure that both data at rest and data in transit are encrypted according to the standard. Using customer-managed encryption keys (CMEK) ensures that you have control over the encryption keys, and BoringSSL is a library that meets FIPS 140-2 standards for encrypting data in transit.

Steps:

Encrypt Local SSDs: Modify the instance template for the Managed Instance Group (MIG) to use customer-managed encryption keys (CMEK) for encrypting Local SSDs.

Enable BoringSSL: Update the application to use the BoringSSL library for all instance-to-instance communication to ensure that all data in transit is encrypted according to FIPS 140-2 standards. Reference::

Google Cloud: Customer-managed encryption keys (CMEK)

BoringSSL documentation

Question: 65

A customer has an analytics workload running on Compute Engine that should have limited internet access.

Your team created an egress firewall rule to deny (priority 1000) all traffic to the internet.

The Compute Engine instances now need to reach out to the public repository to get security updates. What should your team do?

- A. Create an egress firewall rule to allow traffic to the CIDR range of the repository with a priority greater than 1000.
- B. Create an egress firewall rule to allow traffic to the CIDR range of the repository with a priority less than 1000.
- C. Create an egress firewall rule to allow traffic to the hostname of the repository with a priority greater than 1000.
- D. Create an egress firewall rule to allow traffic to the hostname of the repository with a priority less than 1000.

Answer: B

Explanation:

To allow Compute Engine instances to access public repositories for security updates while an egress firewall rule is in place to deny all internet traffic, you need to create a more specific egress rule that permits traffic to the CIDR range of the repository. The priority of this rule should be lower (i.e., a higher priority number) than the deny rule.

Steps:

Identify the CIDR Range: Determine the CIDR range of the public repository from which the security updates will be fetched.

Create Egress Firewall Rule: Create a new egress firewall rule allowing traffic to the identified CIDR range with a priority less than 1000.

Apply Firewall Rule: Use the Google Cloud Console or gcloud command-line tool to apply the new firewall rule.

Reference::

Google Cloud: Firewall rules

Creating firewall rules

Question: 66

You want data on Compute Engine disks to be encrypted at rest with keys managed by Cloud Key Management Service (KMS). Cloud Identity and Access Management (IAM) permissions to these keys must be managed in a grouped way because the permissions should be the same for all keys.

What should you do?

- A. Create a single KeyRing for all persistent disks and all Keys in this KeyRing. Manage the IAM permissions at the Key level.
- B. Create a single KeyRing for all persistent disks and all Keys in this KeyRing. Manage the IAM permissions at the KeyRing level.
- C. Create a KeyRing per persistent disk, with each KeyRing containing a single Key. Manage the IAM permissions at the

Key level.

D. Create a KeyRing per persistent disk, with each KeyRing containing a single Key. Manage the IAM permissions at the KeyRing level.

Answer: B

Explanation:

Managing IAM permissions at the KeyRing level is more efficient and scalable compared to managing them at the individual Key level. By creating a single KeyRing and placing all encryption keys within it, you can apply uniform IAM permissions to the entire KeyRing, simplifying the management of permissions.

Steps:

Create a KeyRing: Set up a single KeyRing in Cloud KMS for all the encryption keys required for the persistent disks.

Create Encryption Keys: Generate the necessary encryption keys within this KeyRing.

Set IAM Permissions: Assign IAM roles and permissions to the KeyRing to manage access control at this level, ensuring that all keys within the KeyRing inherit these permissions.

Reference::

Google Cloud: Cloud Key Management Service (KMS)

Managing access to resources

Question: 67

A company is backing up application logs to a Cloud Storage bucket shared with both analysts and the administrator.

Analysts should only have access to logs that do not contain any personally identifiable information (PII). Log files containing

PII should be stored in another bucket that is only accessible by the administrator.

What should you do?

A. Use Cloud Pub/Sub and Cloud Functions to trigger a Data Loss Prevention scan every time a file is uploaded to the shared bucket. If the scan detects PII, have the function move into a Cloud Storage bucket only accessible by the administrator.

B. Upload the logs to both the shared bucket and the bucket only accessible by the administrator. Create a job trigger using the Cloud Data Loss Prevention API. Configure the trigger to delete any files from the shared bucket that contain PII.

C. On the bucket shared with both the analysts and the administrator, configure Object Lifecycle Management to delete objects that contain any PII.

D. On the bucket shared with both the analysts and the administrator, configure a Cloud Storage Trigger that is only triggered when PII data is uploaded. Use Cloud Functions to capture the trigger and delete such files.

Answer: A

Explanation:

To ensure that PII data is separated from non-PII data, using Cloud Pub/Sub and Cloud Functions to trigger a scan by the Data Loss Prevention (DLP) API is an effective approach. This method allows for automated detection and handling of

PII.

Steps:

Set Up Cloud Pub/Sub: Configure a Cloud Pub/Sub topic to receive notifications whenever a file is uploaded to the shared Cloud Storage bucket.

Deploy Cloud Functions: Create a Cloud Function that is triggered by the Pub/Sub topic. This function will invoke the DLP API to scan the uploaded file for PII.

Move Detected PII Files: If the scan detects PII, the Cloud Function will move the file to a secure Cloud Storage bucket accessible only by the administrator.

Set Permissions: Ensure that appropriate permissions are set on the Cloud Storage buckets to restrict access to files containing PII.

Reference::

Google Cloud: Data Loss Prevention

Cloud Functions documentation

Question: 68

A customer terminates an engineer and needs to make sure the engineer's Google account is automatically deprovisioned.

What should the customer do?

- A. Use the Cloud SDK with their directory service to remove their IAM permissions in Cloud Identity.
- B. Use the Cloud SDK with their directory service to provision and deprovision users from Cloud Identity.
- C. Configure Cloud Directory Sync with their directory service to provision and deprovision users from Cloud Identity.
- D. Configure Cloud Directory Sync with their directory service to remove their IAM permissions in Cloud Identity.

Answer: C

Explanation:

https://cloud.google.com/identity/solutions/automate-user-provisioning#cloud_identity_automated_provisioning

"Cloud Identity has a catalog of automated provisioning connectors, which act as a bridge between Cloud Identity and third-party cloud apps."

Question: 69

An organization is evaluating the use of Google Cloud Platform (GCP) for certain IT workloads. A well-established directory service is used to manage user identities and lifecycle management. This

directory service must continue for the organization to use as the "source of truth" directory for identities.

Which solution meets the organization's requirements?

- A. Google Cloud Directory Sync (GCDS)

- B. Cloud Identity
- C. Security Assertion Markup Language (SAML)
- D. Pub/Sub

Answer: A

Explanation:

With Google Cloud Directory Sync (GCDS), you can synchronize the data in your Google Account with your Microsoft Active Directory or LDAP server. GCDS doesn't migrate any content (such as email messages, calendar events, or files) to your Google Account. You use GCDS to synchronize your Google users, groups, and shared contacts to match the information in your LDAP server. <https://support.google.com/a/answer/106368?hl=en>

Question: 70

Which international compliance standard provides guidelines for information security controls applicable to the provision and use of cloud services?

- A. ISO 27001
- B. ISO 27002
- C. ISO 27017
- D. ISO 27018

Answer: C

Explanation:

Create a new Service Account that should be able to list the Compute Engine instances in the project. You want to follow Google-recommended practices.

<https://cloud.google.com/security/compliance/iso-27017>

Question: 71

You will create a new Service Account that should be able to list the Compute Engine instances in the project. You want to follow Google-recommended practices.

What should you do?

- A. Create an Instance Template, and allow the Service Account Read Only access for the Compute Engine Access

Scope.

- B. Create a custom role with the permission compute.instances.list and grant the Service Account this role.
- C. Give the Service Account the role of Compute Viewer, and use the new Service Account for all instances.
- D. Give the Service Account the role of Project Viewer, and use the new Service Account for all instances.

Answer: B

Explanation:

Objective: Create a Service Account that can list Compute Engine instances in the project following Google-recommended practices.

Solution: Create a custom role and assign it to the Service Account.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the IAM & Admin page and select "Roles".

Step 3: Click on "Create Role" and define a new role with a suitable name and description.

Step 4: Add the permission compute.instances.list to the custom role.

Step 5: Save the custom role.

Step 6: Go to the "Service Accounts" section.

Step 7: Create a new Service Account or select an existing one.

Step 8: Assign the newly created custom role to the Service Account.

By creating a custom role with the specific permission to list Compute Engine instances, you follow the principle of least privilege, which is a recommended security practice.

Reference::

Creating and Managing Custom Roles

Best Practices for IAM

Question: 72

In a shared security responsibility model for IaaS, which two layers of the stack does the customer share responsibility for? (Choose two.)

- A. Hardware
- B. Network Security
- C. Storage Encryption
- D. Access Policies
- E. Boot

Answer: BD

Explanation:

Objective: Identify the layers of the stack the customer shares responsibility for in a shared security responsibility model for IaaS.

Solution: Understand the shared responsibility model.

Network Security: Customers are responsible for configuring network security, such as setting up firewalls, managing VPCs, and ensuring secure network traffic.

Access Policies: Customers are responsible for managing access policies, including IAM roles and permissions, to ensure only authorized users can access resources.

In IaaS, the cloud provider is typically responsible for the underlying infrastructure, while the customer is responsible for securing their applications and data on the cloud.

Reference::

Shared Responsibility Model

Google Cloud Security Overview

Question: 73

An organization is starting to move its infrastructure from its on-premises environment to Google Cloud Platform (GCP). The first step the organization wants to take is to migrate its ongoing data backup and disaster recovery solutions to GCP. The organization's on-premises production environment is going to be the next phase for migration to GCP. Stable networking connectivity between the on-premises environment and GCP is also being implemented.

Which GCP solution should the organization use?

- A. BigQuery using a data pipeline job with continuous updates via Cloud VPN
- B. Cloud Storage using a scheduled task and gsutil via Cloud Interconnect
- C. Compute Engines Virtual Machines using Persistent Disk via Cloud Interconnect
- D. Cloud Datastore using regularly scheduled batch upload jobs via Cloud VPN

Answer: B

Explanation:

Objective: Migrate ongoing data backup and disaster recovery solutions to GCP.

Solution: Use Cloud Storage with scheduled tasks and gsutil.

Steps:

Step 1: Set up a Cloud Interconnect to ensure stable networking connectivity between the on-premises environment and GCP.

Step 2: Create a Cloud Storage bucket to store backups.

Step 3: Use gsutil, a command-line tool for Cloud Storage, to create scripts for data transfer.

Step 4: Schedule these scripts using cron jobs or another scheduling tool to automate the backup process.

Using Cloud Storage with scheduled tasks and gsutil ensures efficient and reliable backup and disaster recovery while leveraging stable connectivity provided by Cloud Interconnect.

Reference::

Cloud Storage Documentation

gsutil Tool Documentation

Cloud Interconnect Documentation

Question: 74

What are the steps to encrypt data using envelope encryption?

- A. Generate a data encryption key (DEK) locally.
Use a key encryption key (KEK) to wrap the DEK. Encrypt data with the KEK.
Store the encrypted data and the wrapped KEK.
- B. Generate a key encryption key (KEK) locally.
Use the KEK to generate a data encryption key (DEK). Encrypt data with the DEK.
Store the encrypted data and the wrapped DEK.
- C. Generate a data encryption key (DEK) locally.
Encrypt data with the DEK.
Use a key encryption key (KEK) to wrap the DEK. Store the encrypted data and the wrapped DEK.
- D. Generate a key encryption key (KEK) locally.
Generate a data encryption key (DEK) locally. Encrypt data with the KEK.
Store the encrypted data and the wrapped DEK.

Answer: C

Explanation:

Objective: Encrypt data using envelope encryption.

Solution: Follow the envelope encryption process.

Steps:

Step 1: Generate a Data Encryption Key (DEK) locally. The DEK is used to encrypt the actual data.

Step 2: Encrypt the data using the DEK.

Step 3: Use a Key Encryption Key (KEK) to wrap the DEK. The KEK is used to encrypt the DEK.

Step 4: Store the encrypted data and the wrapped DEK. This ensures that the data can be securely decrypted in the future using the KEK to unwrap the DEK.

Envelope encryption enhances security by adding an additional layer of encryption to the data encryption key, which is particularly useful for managing large volumes of encrypted data. Reference::

Envelope Encryption Overview

Google Cloud Key Management Service Documentation

Question: 75

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 (Cloud IAP) provides a way to control access to your web applications and resources running on Google Cloud. It works by verifying the identity of a user trying to access the application and supports multi-factor authentication (MFA). Cloud IAP can restrict access to users on the corporate network and also supports access over the internet securely.

Steps:

Enable Cloud IAP: In the Google Cloud Console, navigate to the IAP section and enable IAP for your web application.

Configure OAuth Consent Screen: Set up the OAuth consent screen to manage how users grant access.

Set Up Authentication: Use Google Identity Platform to manage users and enable two-factor authentication.

Add Users: Grant users access to the application by adding their identities in the IAP settings. Reference::

Google Cloud: Identity-Aware Proxy

Setting up IAP

Question: 76

Your company is storing sensitive data in Cloud Storage. You want a key generated on-premises to be used in the encryption process.

What should you do?

- A. Use the Cloud Key Management Service to manage a data encryption key (DEK).
- B. Use the Cloud Key Management Service to manage a key encryption key (KEK).
- C. Use customer-supplied encryption keys to manage the data encryption key (DEK).
- D. Use customer-supplied encryption keys to manage the key encryption key (KEK).

Answer: C

Explanation:

This is a Customer-supplied encryption keys (CSEK). We generate our own encryption key and manage it on-premises. A KEK never leaves Cloud KMS. There is no KEK or KMS on-premises. Encryption at rest by default, with various key management options <https://cloud.google.com/security/encryption-at-rest>

Reference:: <https://cloud.google.com/security/encryption-at-rest/default-encryption/>

Question: 77

Last week, a company deployed a new App Engine application that writes logs to BigQuery. No other workloads are running in the project. You need to validate that all data written to BigQuery was done using the App Engine Default Service Account.

What should you do?

- A. 1. Use StackDriver Logging and filter on BigQuery Insert Jobs.
2. Click on the email address in line with the App Engine Default Service Account in the authentication field.
3. Click Hide Matching Entries.
4. Make sure the resulting list is empty.
- B. 1. Use StackDriver Logging and filter on BigQuery Insert Jobs.
2. Click on the email address in line with the App Engine Default Service Account in the authentication field.
3. Click Show Matching Entries.
4. Make sure the resulting list is empty.
- C. 1. In BigQuery, select the related dataset.
2. Make sure the App Engine Default Service Account is the only account that can write to the dataset.
- D. 1. Go to the IAM section on the project.
2. Validate that the App Engine Default Service Account is the only account that has a role that can write to BigQuery.

Answer: A

Explanation:

To validate that all data written to BigQuery was done using the App Engine Default Service Account, you can use StackDriver Logging (now known as Cloud Logging) to filter and inspect the logs for

BigQuery insert jobs. By hiding the entries matching the App Engine Default Service Account, you can ensure that no other service account has written to BigQuery if the resulting list is empty.

Steps:

Open Cloud Logging: Navigate to Cloud Logging in the Google Cloud Console.

Filter Logs: Apply a filter to display logs for BigQuery insert jobs.

Inspect Entries: Click on the email address that corresponds to the App Engine Default Service Account in the authentication field.

Hide Matching Entries: Select the option to hide matching entries.

Validate: Check if the resulting list is empty, confirming that no other service account has performed write operations to BigQuery.

Reference::

Google Cloud Logging

Monitoring BigQuery logs

Question: 78

Your team wants to limit users with administrative privileges at the organization level.

Which two roles should your team restrict? (Choose two.)

- A. Organization Administrator
- B. Super Admin
- C. GKE Cluster Admin
- D. Compute Admin
- E. Organization Role Viewer

Answer: AB

Explanation:

The Organization Administrator and Super Admin roles have extensive administrative privileges at the organization level. Restricting these roles is crucial to limit the number of users who have the ability to manage critical resources and configurations within the organization, thereby enhancing security and minimizing potential risks.

Organization Administrator: Has comprehensive permissions to manage all aspects of the Google Cloud organization, including projects, folders, and IAM policies.

Super Admin: In Google Workspace (formerly G Suite), the Super Admin has access to all administrative features and can manage user accounts, services, and settings across the organization. Reference::

Google Cloud IAM roles

Managing super admin roles in Google Workspace

Question: 79

An organization's security and risk management teams are concerned about where their responsibility lies for certain production workloads they are running in Google Cloud Platform (GCP), and where Google's responsibility lies. They are mostly running workloads using Google Cloud's Platform-as-a-Service (PaaS) offerings, including App Engine primarily.

Which one of these areas in the technology stack would they need to focus on as their primary responsibility when using App Engine?

- A. Configuring and monitoring VPC Flow Logs
- B. Defending against XSS and SQLi attacks
- C. Manage the latest updates and security patches for the Guest OS
- D. Encrypting all stored data

Answer: B

Explanation:

When using Google Cloud's Platform-as-a-Service (PaaS) offerings like App Engine, Google manages the infrastructure, including the underlying OS, runtime, and scaling. However, securing the application code itself, such as defending against cross-site scripting (XSS) and SQL injection (SQLi) attacks, remains the responsibility of the user. This involves implementing secure coding practices, validating inputs, and employing appropriate security measures within the application.

Reference::

Google Cloud: Shared responsibility model

App Engine security

Question: 80

An engineering team is launching a web application that will be public on the internet. The web application is hosted in multiple GCP regions and will be directed to the respective backend based on the URL request.

Your team wants to avoid exposing the application directly on the internet and wants to deny traffic from a specific list of malicious IP addresses

Which solution should your team implement to meet these requirements?

- A. Cloud Armor
- B. Network Load Balancing
- C. SSL Proxy Load Balancing
- D. NAT Gateway

Answer: A

Explanation:

Google Cloud Armor provides protection against DDoS attacks and allows you to define security policies to control access to your application. It enables you to block traffic from specific IP addresses or ranges, making it suitable for denying traffic from a list of malicious IP addresses while protecting your application from being directly exposed to the internet.

Steps:

Set Up Cloud Armor: Enable Cloud Armor in your Google Cloud Console.

Create Security Policies: Define security policies that specify the rules for allowing or denying traffic based on IP addresses.

Attach Policies to Backend Services: Apply these security policies to the backend services of your web application.

Reference::

Google Cloud Armor documentation
Creating and managing security policies

Question: 81

A customer is running an analytics workload on Google Cloud Platform (GCP) where Compute Engine instances are accessing data stored on Cloud Storage. Your team wants to make sure that this workload will not be able to access, or be accessed from, the internet.

Which two strategies should your team use to meet these requirements? (Choose two.)

- A. Configure Private Google Access on the Compute Engine subnet
- B. Avoid assigning public IP addresses to the Compute Engine cluster.
- C. Make sure that the Compute Engine cluster is running on a separate subnet.
- D. Turn off IP forwarding on the Compute Engine instances in the cluster.
- E. Configure a Cloud NAT gateway.

Answer: AB

Explanation:

Objective: Ensure that the analytics workload on Compute Engine instances accessing Cloud Storage does not interact with the public internet.

Solution:

Private Google Access: This allows Compute Engine instances that only have internal IP addresses to reach Google APIs and services through a private connection without the need for a public IP address.

No Public IP Addresses: By avoiding public IP addresses for the instances, you ensure that they are not accessible from the internet and do not initiate internet connections.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the VPC Network page and select the subnet where the Compute Engine instances are located.

Step 3: Enable Private Google Access for the subnet.

Step 4: Ensure that when launching the Compute Engine instances, no public IP addresses are assigned to them.

Reference::

Configuring Private Google Access

Preventing External IP Address Assignment

Question: 82

A customer wants to run a batch processing system on VMs and store the output files in a Cloud Storage bucket. The networking and security teams have decided that no VMs may reach the public internet.

How should this be accomplished?

- A. Create a firewall rule to block internet traffic from the VM.
- B. Provision a NAT Gateway to access the Cloud Storage API endpoint.
- C. Enable Private Google Access on the VPC.
- D. Mount a Cloud Storage bucket as a local filesystem on every VM.

Answer: C

Explanation:

Objective: Ensure VMs can access Cloud Storage without reaching the public internet.

Solution: Enable Private Google Access on the VPC network, allowing VMs with only internal IP addresses to access Google APIs and services privately.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Go to the VPC Network section.

Step 3: Select the relevant VPC network and subnet.

Step 4: Enable Private Google Access for the subnet.

Private Google Access ensures that instances can access Google APIs and services (such as Cloud Storage) over a private network connection, without requiring a public IP address.

Reference::

Configuring Private Google Access

Best Practices for Secure Access

Question: 83

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 `rowLimit` value on BigQuery data hosted outside the US and set appropriate `bytesLimitPerFile` value on multiregional Cloud Storage buckets.
- B. Set appropriate `rowLimit` value on BigQuery data hosted outside the US, and minimize transformation units on multiregional Cloud Storage buckets.
- C. Use `rowLimit` and `bytesLimitPerFile` to sample data and use `CloudStorageRegexFileSet` to limit scans.
- D. Use `FindingLimits` and `TimespanConfig` to sample data and minimize transformation units.

Answer: C

Explanation:

Objective: Optimize the usage of Cloud Data Loss Prevention (DLP) API to reduce costs.

Solution:

`rowLimit` and `bytesLimitPerFile`: These parameters help in sampling data instead of scanning the entire dataset, thereby reducing the amount of data processed.

`CloudStorageRegexFileSet`: This feature allows you to specify a subset of files to be scanned using regular expressions, limiting the scope and volume of data scanned.

Steps:

Step 1: Set appropriate `rowLimit` values for BigQuery data scans to sample rows instead of scanning entire tables.

Step 2: Set `bytesLimitPerFile` values for Cloud Storage buckets to limit the number of bytes scanned per file.

Step 3: Use `CloudStorageRegexFileSet` to specify the subset of files to be scanned based on patterns that match the filenames.

By combining these strategies, you effectively reduce the scope and volume of data processed by the DLP API, leading to cost savings.

Reference::

DLP API Best Practices

Configuring Finding Limits

Question: 84

Your team uses a service account to authenticate data transfers from a given Compute Engine virtual machine instance of to a specified Cloud Storage bucket. An engineer accidentally deletes the service account, which breaks application functionality. You want to recover the application as quickly as possible without compromising security.

What should you do?

- A. Temporarily disable authentication on the Cloud Storage bucket.

- B. Use the undelete command to recover the deleted service account.
- C. Create a new service account with the same name as the deleted service account.
- D. Update the permissions of another existing service account and supply those credentials to the applications.

Answer: B

Explanation:

Objective: Quickly recover a deleted service account used for data transfers.

Solution: Use the undelete command available in the gcloud command-line tool to recover the service account.

Steps:

Step 1: Open the Cloud Shell in the Google Cloud Console.

Step 2: Run the following command to list the deleted service accounts:

```
gcloud iam service-accounts list --filter="deleted: true"
```

Step 3: Identify the name and ID of the deleted service account.

Step 4: Use the undelete command to recover the service account:

```
gcloud iam service-accounts undelete [SERVICE_ACCOUNT_ID]
```

Step 5: Verify that the service account has been restored and reassign any necessary permissions.

By using the undelete command, you can quickly restore the service account and resume application functionality without compromising security.

Reference::

Restoring a Deleted Service Account `gcloud iam service-accounts undelete`

Question: 85

You are the Security Admin in your company. You want to synchronize all security groups that have an email address from your LDAP directory in Cloud IAM.

What should you do?

- A. Configure Google Cloud Directory Sync to sync security groups using LDAP search rules that have “user email address” as the attribute to facilitate one-way sync.
- B. Configure Google Cloud Directory Sync to sync security groups using LDAP search rules that have “user email address” as the attribute to facilitate bidirectional sync.
- C. Use a management tool to sync the subset based on the email address attribute. Create a group in the Google domain. A group created in a Google domain will automatically have an explicit Google Cloud Identity and Access Management (IAM) role.
- D. Use a management tool to sync the subset based on group object class attribute. Create a group in the Google domain. A group created in a Google domain will automatically have an explicit Google Cloud Identity and Access Management (IAM) role.

Answer: A

Explanation:

Objective: Synchronize security groups with email addresses from an LDAP directory to Cloud IAM. Solution: Use Google Cloud Directory Sync (GCDS) to perform one-way synchronization based on LDAP search rules.

Steps:

Step 1: Download and install Google Cloud Directory Sync (GCDS) on a secure server.

Step 2: Configure GCDS with the LDAP server details and authentication.

Step 3: Define LDAP search rules to filter security groups based on the "user email address" attribute.

Step 4: Map LDAP security groups to Google Cloud IAM roles.

Step 5: Set up a synchronization schedule to keep the groups in sync.

Step 6: Perform a test sync to ensure that the configuration is correct.

Step 7: Activate the synchronization to keep the LDAP directory and Cloud IAM in sync.

Using GCDS for one-way synchronization ensures that the security groups in Cloud IAM are consistently updated based on the LDAP directory, maintaining alignment with the organization's security policies.

Reference::

Google Cloud Directory Sync Documentation

Setting Up Google Cloud Directory Sync

Question: 86

You are part of a security team investigating a compromised service account key. You need to audit which new resources were created by the service account.

What should you do?

- A. Query Data Access logs.
- B. Query Admin Activity logs.
- C. Query Access Transparency logs.
- D. Query Stackdriver Monitoring Workspace.

Answer: B

Explanation:

Admin activity logs are always created to log entries for API calls or other actions that modify the configuration or metadata of resources. For example, these logs record when users create VM instances or change Identity and Access Management permissions.

Reference:: <https://cloud.google.com/iam/docs/audit-logging/examples-service-accounts>

Question: 87

You have an application where the frontend is deployed on a managed instance group in subnet A and the data layer is stored on a mysql Compute Engine virtual machine (VM) in subnet B on the same VPC. Subnet A and Subnet B hold several other Compute Engine VMs. You only want to allow the application frontend to access the data in the application's mysql instance on port 3306.

What should you do?

- A. Configure an ingress firewall rule that allows communication from the src IP range of subnet A to the tag "data-tag" that is applied to the mysql Compute Engine VM on port 3306.
- B. Configure an ingress firewall rule that allows communication from the frontend's unique service account to the unique service account of the mysql Compute Engine VM on port 3306.
- C. Configure a network tag "fe-tag" to be applied to all instances in subnet A and a network tag "data-tag" to be applied to all instances in subnet B. Then configure an egress firewall rule that allows communication from Compute Engine VMs tagged with data-tag to destination Compute Engine VMs tagged fe-tag.
- D. Configure a network tag "fe-tag" to be applied to all instances in subnet A and a network tag "data-tag" to be applied to all instances in subnet B. Then configure an ingress firewall rule that allows communication from Compute Engine VMs tagged with fe-tag to destination Compute Engine VMs tagged with data-tag.

Answer: A

Explanation:

To restrict access to the MySQL instance to only the frontend application while other VMs are present in the subnets, creating an ingress firewall rule is the most appropriate approach. This rule will specifically allow traffic from subnet A (where the frontend application resides) to the MySQL instance in subnet B on port 3306, using network tags to target the specific MySQL VM.

Steps:

Create Network Tags: Apply a network tag (e.g., "data-tag") to the MySQL VM in subnet B.

Create Ingress Firewall Rule: Configure an ingress firewall rule with the following settings:

Source IP Range: Subnet A's IP range.

Target Tag: "data-tag".

Allowed Protocol/Ports: TCP:3306 (for MySQL).

This setup ensures that only instances in subnet A can communicate with the MySQL instance on port 3306.

Reference::

Google Cloud: Configuring firewall rules

Question: 88

Your company operates an application instance group that is currently deployed behind a Google Cloud load balancer in us-central-1 and is configured to use the Standard Tier network. The infrastructure team wants to expand to a second Google Cloud region, us-east-2. You need to set up a single external IP address to distribute new requests to the instance groups in both regions.

What should you do?

- A. Change the load balancer backend configuration to use network endpoint groups instead of instance groups.
- B. Change the load balancer frontend configuration to use the Premium Tier network, and add the new instance group.
- C. Create a new load balancer in us-east-2 using the Standard Tier network, and assign a static external IP address.
- D. Create a Cloud VPN connection between the two regions, and enable Google Private Access.

Answer: B

Explanation:

The Standard Tier network only provides regional load balancing, while the Premium Tier supports global load balancing with a single anycast IP address. To distribute requests across multiple regions, you need to use the Premium Tier and update the load balancer configuration accordingly.

Steps:

Upgrade to Premium Tier: Update the load balancer to use the Premium Tier network in the Google Cloud Console.

Add New Instance Group: Add the instance group in the new region (us-east-2) to the backend configuration of the existing load balancer.

Verify Configuration: Ensure that the frontend configuration of the load balancer uses a single external IP address for global distribution.

Reference::

Google Cloud: Global load balancing

Question: 89

You are the security admin of your company. You have 3,000 objects in your Cloud Storage bucket.

You do not want to manage access to each object individually. You also do not want the uploader of an object to always have full control of the object. However, you want to use Cloud Audit Logs to manage access to your bucket.

What should you do?

- A. Set up an ACL with OWNER permission to a scope of allUsers.
- B. Set up an ACL with READER permission to a scope of allUsers.
- C. Set up a default bucket ACL and manage access for users using IAM.
- D. Set up Uniform bucket-level access on the Cloud Storage bucket and manage access for users using IAM.

Answer: D

Explanation:

Uniform bucket-level access allows you to manage permissions at the bucket level, rather than at the object level. This simplifies permission management and ensures that access to objects is controlled consistently via IAM roles, without allowing uploaders full control over the objects.

Steps:

Enable Uniform Bucket-Level Access: In the Google Cloud Console, enable uniform bucket-level access for the Cloud Storage bucket.

Configure IAM Policies: Assign appropriate IAM roles to users and groups to control access to the bucket.

Audit Logging: Enable Cloud Audit Logs to track access and modifications to the bucket.

Reference::

Google Cloud: Uniform bucket-level access

Managing access with IAM

Question: 90

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: D

Explanation:

Setting up separate service perimeters for dev, staging, and prod environments allows for more granular control and monitoring. Automating the addition of new projects to the respective perimeters ensures that all projects are consistently secured without manual intervention.

Steps:

Set Up Service Perimeters: Use Access Context Manager to define and configure three separate service perimeters for dev, staging, and prod.

Deploy Monitoring Function: Create a Cloud Function that monitors the "implementation" folder for new projects using Stackdriver (Cloud Monitoring) and Cloud Pub/Sub.

Automate Perimeter Updates: Configure the Cloud Function to execute Terraform scripts that automatically add new projects to the appropriate service perimeter.

Reference::

Google Cloud: Access Context Manager

Service perimeter automation

Question: 91

You need to provide a corporate user account in Google Cloud for each of your developers and operational staff who need direct access to GCP resources. Corporate policy requires you to maintain the user identity in a third-party identity management provider and leverage single sign-on. You learn that a significant number of users are using their corporate domain email addresses for personal Google accounts, and you need to follow Google recommended practices to convert existing unmanaged users to managed accounts.

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

- A. Use Google Cloud Directory Sync to synchronize your local identity management system to Cloud Identity.
- B. Use the Google Admin console to view which managed users are using a personal account for their recovery email.
- C. Add users to your managed Google account and force users to change the email addresses associated with their personal accounts.
- D. Use the Transfer Tool for Unmanaged Users (TTUU) to find users with conflicting accounts and ask them to transfer their personal Google accounts.
- E. Send an email to all of your employees and ask those users with corporate email addresses for personal Google accounts to delete the personal accounts immediately.

Answer: AD

Explanation:

To manage user accounts and ensure they comply with corporate policies, using Google Cloud Directory Sync (GCDS) allows synchronization between your local identity system and Cloud Identity. The Transfer Tool for Unmanaged Users (TTUU) helps identify and manage conflicting accounts by allowing users to transfer their personal accounts to managed accounts.

Steps:

Synchronize Identities: Use GCDS to sync users from your local identity management system to Cloud Identity, ensuring that all corporate user accounts are managed.

Identify Conflicting Accounts: Use TTUU to find users who have personal Google accounts using corporate email addresses.

Manage Conflicting Accounts: Request users to transfer their personal accounts to managed accounts using TTUU, ensuring all accounts are under corporate control.

Reference::

Google Cloud Directory Sync

Transfer Tool for Unmanaged Users

Question: 92

You are on your company's development team. You noticed that your web application hosted in staging on GKE dynamically includes user data in web pages without first properly validating the inputted data.

- a. This could allow an attacker to execute gibberish commands and display arbitrary content in a victim user's browser in a production environment.

How should you prevent and fix this vulnerability?

- A. Use Cloud IAP based on IP address or end-user device attributes to prevent and fix the vulnerability.
- B. Set up an HTTPS load balancer, and then use Cloud Armor for the production environment to prevent the potential XSS attack.
- C. Use Web Security Scanner to validate the usage of an outdated library in the code, and then use a secured version of the included library.
- D. Use Web Security Scanner in staging to simulate an XSS injection attack, and then use a templating system that supports contextual auto-escaping.

Answer: D

Explanation:

There is mention about simulating in Web Security Scanner. "Web Security Scanner cross-site scripting (XSS) injection testing *simulates* an injection attack by inserting a benign test string into user-editable fields and then performing various user actions." <https://cloud.google.com/security-command-center/docs/how-to-remediate-web-security-scanner-findings#xss>

Reference:: <https://cloud.google.com/security-scanner/docs/remediate-findings>

Question: 93

You are part of a security team that wants to ensure that a Cloud Storage bucket in Project A can only be readable from Project B. You also want to ensure that data in the Cloud Storage bucket cannot be accessed from or copied to Cloud Storage buckets outside the network, even if the user has the correct credentials.

What should you do?

- A. Enable VPC Service Controls, create a perimeter with Project A and B, and include Cloud Storage service.
- B. Enable Domain Restricted Sharing Organization Policy and Bucket Policy Only on the Cloud Storage bucket.
- C. Enable Private Access in Project A and B networks with strict firewall rules to allow communication between the networks.
- D. Enable VPC Peering between Project A and B networks with strict firewall rules to allow communication between the networks.

Answer: A

Explanation:

Objective: Ensure that a Cloud Storage bucket in Project A can only be readable from Project B and prevent data access or copying to Cloud Storage buckets outside the network, even with correct credentials.

Solution: Use VPC Service Controls to create a security perimeter.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the VPC Service Controls page.

Step 3: Create a new service perimeter.

Step 4: Add Project A and Project B to the service perimeter.

Step 5: Include Cloud Storage service in the perimeter configuration.

Step 6: Define access levels to ensure that only resources within the perimeter can access the Cloud Storage bucket.

By setting up a VPC Service Controls perimeter, you can enforce security boundaries that restrict data access and movement to within defined projects, providing an extra layer of protection beyond IAM permissions.

Reference::

VPC Service Controls Overview

Configuring VPC Service Controls

Question: 94

You are responsible for protecting highly sensitive data in BigQuery. Your operations teams need access to this data, but given privacy regulations, you want to ensure that they cannot read the sensitive fields such as email addresses and first names. These specific sensitive fields should only be available on a need-to-know basis to the HR team. What should you do?

- A. Perform data masking with the DLP API and store that data in BigQuery for later use.
- B. Perform data redaction with the DLP API and store that data in BigQuery for later use.
- C. Perform data inspection with the DLP API and store that data in BigQuery for later use.
- D. Perform tokenization for Pseudonymization with the DLP API and store that data in BigQuery for later use.

Answer: D

Explanation:

Pseudonymization is a de-identification technique that replaces sensitive data values with cryptographically generated tokens. Pseudonymization is widely used in industries like finance and healthcare to help reduce the risk of data in use, narrow compliance scope, and minimize the exposure of sensitive data to systems while preserving data utility and accuracy.

<https://cloud.google.com/dlp/docs/pseudonymization>

Question: 95

You are a Security Administrator at your organization. You need to restrict service account creation capability within production environments. You want to accomplish this centrally across the organization. What should you do?

- A. Use Identity and Access Management (IAM) to restrict access of all users and service accounts that have access to the production environment.
- B. Use organization policy constraints/iam.disableServiceAccountKeyCreation boolean to disable the creation of new service accounts.
- C. Use organization policy constraints/iam.disableServiceAccountKeyUpload boolean to disable the creation of new service accounts.
- D. Use organization policy constraints/iam.disableServiceAccountCreation boolean to disable the creation of new service accounts.

Answer: D

Explanation:

Reference:: <https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-accounts>

You can use the iam.disableServiceAccountCreation boolean constraint to disable the creation of new service accounts. This

allows you to centralize management of service accounts while not restricting the other permissions your developers have on projects.

https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-accounts#disable_service_account_creation

Question: 96

You are the project owner for a regulated workload that runs in a project you own and manage as an Identity and Access Management (IAM) admin. For an upcoming audit, you need to provide access reviews evidence. Which tool should you use?

- A. Policy Troubleshooter
- B. Policy Analyzer
- C. IAM Recommender
- D. Policy Simulator

Answer: B

Explanation:

Objective: Provide evidence of access reviews for an upcoming audit.

Solution: Use Policy Analyzer to review and report on IAM policies.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the Policy Analyzer tool.

Step 3: Select the project for which you need to review access policies.

Step 4: Use the tool to generate reports on IAM roles and permissions.

Step 5: Export the reports as evidence for the audit.

Policy Analyzer provides detailed insights into IAM policies, helping you to review access configurations and generate necessary reports for compliance and auditing purposes. Reference::

Policy Analyzer Documentation

Question: 97

Your organization has implemented synchronization and SAML federation between Cloud Identity and Microsoft Active Directory. You want to reduce the risk of Google Cloud user accounts being compromised. What should you do?

- A. Create a Cloud Identity password policy with strong password settings, and configure 2-Step Verification with security keys in the Google Admin console.
- B. Create a Cloud Identity password policy with strong password settings, and configure 2-Step Verification with verification codes via text or phone call in the Google Admin console.
- C. Create an Active Directory domain password policy with strong password settings, and configure post-SSO (single sign-on) 2-Step Verification with security keys in the Google Admin console.
- D. Create an Active Directory domain password policy with strong password settings, and configure post-SSO (single sign-on) 2-Step Verification with verification codes via text or phone call in the Google Admin console.

Answer: C

Explanation:

Objective: Reduce the risk of Google Cloud user accounts being compromised.

Solution: Implement strong password policies and post-SSO 2-Step Verification using security keys.

Steps:

Step 1: In Active Directory, configure a domain password policy with strong settings (e.g., complexity, length, expiration).

Step 2: In the Google Admin console, navigate to the Security settings.

Step 3: Enable 2-Step Verification and configure it to use security keys for post-SSO verification.

Step 4: Ensure all users enroll in the 2-Step Verification with security keys.

Using strong password policies in Active Directory along with security keys for 2-Step Verification post-SSO provides enhanced security against account compromises.

Reference::

[Active Directory Password Policies](#)

Google Admin Console 2-Step Verification

Question: 98

You have been tasked with implementing external web application protection against common web application attacks for a public application on Google Cloud. You want to validate these policy changes before they are enforced. What service should you use?

- A. Google Cloud Armor's preconfigured rules in preview mode
- B. Prepopulated VPC firewall rules in monitor mode
- C. The inherent protections of Google Front End (GFE)
- D. Cloud Load Balancing firewall rules
- E. VPC Service Controls in dry run mode

Answer: A

Explanation:

Objective: Implement external web application protection and validate policy changes before enforcement.

Solution: Use Google Cloud Armor's preconfigured rules in preview mode.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the Google Cloud Armor section.

Step 3: Create or select a security policy.

Step 4: Apply preconfigured rules to the policy.

Step 5: Enable preview mode to simulate the effects of the rules without enforcing them.

Step 6: Monitor the logs to validate the policy changes.

Google Cloud Armor's preview mode allows you to test and validate the impact of security policies on your application traffic before applying them, ensuring that they work as intended without disrupting the service.

Reference::

Question: 99

You are asked to recommend a solution to store and retrieve sensitive configuration data from an application that runs on Compute Engine. Which option should you recommend?

- A. Cloud Key Management Service
- B. Compute Engine guest attributes
- C. Compute Engine custom metadata
- D. Secret Manager

Answer: D

Explanation:

Objective: Store and retrieve sensitive configuration data for an application running on Compute Engine.

Solution: Use Secret Manager to securely store and manage access to sensitive configuration data. **Steps:**

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the Secret Manager section.

Step 3: Create a new secret and add the sensitive configuration data.

Step 4: Set appropriate IAM policies to control access to the secret.

Step 5: Update the application to retrieve the secret from Secret Manager using the appropriate client libraries or APIs. Secret Manager provides a secure and centralized way to manage sensitive information, with finegrained access control and audit logging capabilities.

Reference::

Secret Manager Documentation

Storing and Accessing Secrets

Question: 100

You need to implement an encryption at-rest strategy that reduces key management complexity for non-sensitive data and protects sensitive data while providing the flexibility of controlling the key residency and rotation schedule. FIPS 140-2 L1 compliance is required for all data types. What should you do?

- A. Encrypt non-sensitive data and sensitive data with Cloud External Key Manager.
- B. Encrypt non-sensitive data and sensitive data with Cloud Key Management Service
- C. Encrypt non-sensitive data with Google default encryption, and encrypt sensitive data with Cloud External Key Manager.
- D. Encrypt non-sensitive data with Google default encryption, and encrypt sensitive data with Cloud Key Management Service.

Answer: D

Explanation:

Objective: Implement an encryption at-rest strategy that balances key management complexity and control for sensitive and non-sensitive data, ensuring FIPS 140-2 L1 compliance.

Solution: Use Google default encryption for non-sensitive data and Cloud Key Management Service (KMS) for sensitive data.

Steps:

Step 1: Store non-sensitive data using Google Cloud's default encryption, which automatically encrypts data at rest without additional configuration.

Step 2: For sensitive data, use Cloud KMS to create and manage encryption keys.

Step 3: Configure key rotation policies for the keys managed by Cloud KMS to meet compliance requirements.

Step 4: Ensure that all data encryption keys used by Cloud KMS comply with FIPS 140-2 Level 1 standards.

By using Google default encryption for non-sensitive data and Cloud KMS for sensitive data, you can manage encryption efficiently while maintaining control over key residency and rotation for sensitive data.

Reference::

Google Cloud Default Encryption

Cloud Key Management Service

FIPS 140-2 Compliance

Question: 101

Your company wants to determine what products they can build to help customers improve their credit scores depending on their age range. To achieve this, you need to join user information in the company's banking app with customers' credit score data received from a third party. While using this raw data will allow you to complete this task, it exposes sensitive data, which could be propagated into new systems.

This risk needs to be addressed using de-identification and tokenization with Cloud Data Loss Prevention while maintaining the referential integrity across the database. Which cryptographic token format should you use to meet these requirements?

- A. Deterministic encryption
- B. Secure, key-based hashes
- C. Format-preserving encryption
- D. Cryptographic hashing

Answer: A

Explanation:

"This encryption method is reversible, which helps to maintain referential integrity across your database and has no character-set limitations." <https://cloud.google.com/blog/products/identity-security/take-charge-of-your-data-how-tokenization-makes-data-usable-without-sacrificing-privacy>

<https://cloud.google.com/dlp/docs/pseudonymization>

FPE provides fewer security guarantees compared to other deterministic encryption methods such as AES-SIV. For these reasons, Google strongly recommends using deterministic encryption with AES-SIV instead of FPE for all security sensitive

use cases. Other methods like deterministic encryption using AES-SIV provide these stronger security guarantees and are recommended for tokenization use cases unless length and character set preservation are strict requirements—for example, for backward compatibility with a legacy data system.

Question: 102

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: CD

Explanation:

<https://cloud.google.com/billing/docs/how-to/billing-access#overview-of-cloud-billing-roles-in-cloud-iam>

Billing Account Costs Manager (roles/billing.costsManager)

- Manage budgets and view and export cost information of billing accounts (but not pricing information)

Billing Account Viewer (roles/billing.viewer)

- View billing account cost information and transactions.

Question: 103

You are designing a new governance model for your organization's secrets that are stored in Secret Manager. Currently, secrets for Production and Non-Production applications are stored and accessed using service accounts. Your

proposed solution must: Provide granular access to secrets

Give you control over the rotation schedules for the encryption keys that wrap your secrets Maintain environment separation

Provide ease of management

Which approach should you take?

- A.
 1. Use separate Google Cloud projects to store Production and Non-Production secrets.
 2. Enforce access control to secrets using project-level identity and Access Management (IAM) bindings.
 3. Use customer-managed encryption keys to encrypt secrets.
- B.
 1. Use a single Google Cloud project to store both Production and Non-Production secrets.

- 2. Enforce access control to secrets using secret-level Identity and Access Management (IAM) bindings.
 - 3. Use Google-managed encryption keys to encrypt secrets.
- C.
- 1. Use separate Google Cloud projects to store Production and Non-Production secrets.
 - 2. Enforce access control to secrets using secret-level Identity and Access Management (IAM) bindings.
 - 3. Use Google-managed encryption keys to encrypt secrets.

- D.
- 1. Use a single Google Cloud project to store both Production and Non-Production secrets.
 - 2. Enforce access control to secrets using project-level Identity and Access Management (IAM) bindings.
 - 3. Use customer-managed encryption keys to encrypt secrets.

Answer: A

Explanation:

Provide granular access to secrets: 2. Enforce access control to secrets using project-level identity and Access Management (IAM) bindings. Give you control over the rotation schedules for the encryption keys that wrap your secrets: 3. Use customer-managed encryption keys to encrypt secrets. Maintain environment separation: 1. Use separate Google Cloud projects to store Production and Non-Production secrets.

Question: 104

You are a security engineer at a finance company. Your organization plans to store data on Google Cloud, but your leadership team is worried about the security of their highly sensitive data. Specifically, your

company is concerned about internal Google employees' ability to access your company's data on Google Cloud. What solution should you propose?

- A. Use customer-managed encryption keys.
- B. Use Google's Identity and Access Management (IAM) service to manage access controls on Google Cloud.
- C. Enable Admin activity logs to monitor access to resources.
- D. Enable Access Transparency logs with Access Approval requests for Google employees.

Answer: D

Explanation:

<https://cloud.google.com/access-transparency> Access approval Explicitly approve access to your data or configurations on Google Cloud. Access Approval requests, when combined with Access Transparency logs, can be used to audit an end-to-end chain from support ticket to access request to approval, to eventual access.

Question: 105

You want to use the gcloud command-line tool to authenticate using a third-party single sign-on (SSO) SAML identity provider. Which options are necessary to ensure that authentication is supported by the third-party identity provider (IdP)? (Choose two.)

- A. SSO SAML as a third-party IdP

- B. Identity Platform
- C. OpenID Connect
- D. Identity-Aware Proxy
- E. Cloud Identity

Answer: A, C

Explanation:

To provide users with SSO-based access to selected cloud apps, Cloud Identity as your IdP supports the OpenID Connect (OIDC) and Security Assertion Markup Language 2.0 (SAML) protocols.

<https://cloud.google.com/identity/solutions/enable-ssso>

Question: 106

You work for a large organization where each business unit has thousands of users. You need to delegate management of access control permissions to each business unit. You have the following requirements:

Each business unit manages access controls for their own projects.

Each business unit manages access control permissions at scale.

Business units cannot access other business units' projects.

Users lose their access if they move to a different business unit or leave the company.

Users and access control permissions are managed by the on-premises directory service.

What should you do? (Choose two.)

- A. Use VPC Service Controls to create perimeters around each business unit's project.
- B. Organize projects in folders, and assign permissions to Google groups at the folder level.
- C. Group business units based on Organization Units (OUs) and manage permissions based on OUs.
- D. Create a project naming convention, and use Google's IAM Conditions to manage access based on the prefix of project names.
- E. Use Google Cloud Directory Sync to synchronize users and group memberships in Cloud Identity.

Answer: B, E

Explanation:

To delegate management of access control permissions to each business unit effectively, organizing projects into folders and assigning permissions to Google groups at the folder level allows for scalable and manageable access control. Using Google Cloud Directory Sync (GCDS) to synchronize users and groups from the on-premises directory service ensures that access controls are maintained and updated automatically as users change roles or leave the company.

Steps:

Organize Projects in Folders: Create a folder structure in the Google Cloud Resource Manager to organize projects by business unit.

Assign Permissions to Google Groups: Use IAM to assign necessary permissions to Google Groups at the folder level,

ensuring each business unit can manage access controls for their own projects. **Synchronize Users and Groups:** Use GCDS to sync users and group memberships from your onpremises directory service to Google Cloud Identity, ensuring that changes in the on-premises directory are reflected in Google Cloud.

Reference::

Google Cloud Resource Manager

Google Cloud Directory Sync

Question: 107

Your organization recently deployed a new application on Google Kubernetes Engine. You need to deploy a solution to protect the application. The solution has the following requirements: Scans must run at least once per week

Must be able to detect cross-site scripting vulnerabilities

Must be able to authenticate using Google accounts

Which solution should you use?

- A. Google Cloud Armor
- B. Web Security Scanner
- C. Security Health Analytics
- D. Container Threat Detection

Answer: B

Explanation:

Web Security Scanner is designed to scan your web applications deployed on Google Cloud for

common vulnerabilities, including cross-site scripting (XSS). It can authenticate using Google accounts and can be scheduled to run scans regularly.

Steps:

Enable Web Security Scanner: In the Google Cloud Console, enable Web Security Scanner for your project.

Configure Scan: Set up the scan configuration, specifying the target URLs, authentication details (Google accounts), and scan frequency (at least once per week).

Run and Monitor Scans: Run the scans and monitor the results for vulnerabilities, addressing any issues found.

Reference::

Web Security Scanner documentation

Question: 108

An organization is moving applications to Google Cloud while maintaining a few mission-critical applications on-premises.

The organization must transfer the data at a bandwidth of at least 50 Gbps. What should they use to ensure secure continued connectivity between sites?

- A. Dedicated Interconnect
- B. Cloud Router
- C. Cloud VPN
- D. Partner Interconnect

Answer: A

Explanation:

Dedicated Interconnect provides a high-bandwidth (up to 80 Gbps per connection) and secure connection between your on-premises network and Google Cloud. It ensures reliable and high-speed data transfer, meeting the requirement of at least 50 Gbps bandwidth.

Steps:

Set Up Dedicated Interconnect: Order a Dedicated Interconnect connection through the Google Cloud Console.

Configure VLAN Attachments: Set up VLAN attachments to segment traffic between your on-premises network and Google Cloud.

Establish BGP Sessions: Configure BGP sessions for dynamic routing and failover.

Reference::

Dedicated Interconnect documentation

Question: 109

Your organization has had a few recent DDoS attacks. You need to authenticate responses to domain name lookups.

Which Google Cloud service should you use?

- A. Cloud DNS with DNSSEC
- B. Cloud NAT
- C. HTTP(S) Load Balancing
- D. Google Cloud Armor

Answer: A

Explanation:

Cloud DNS with DNSSEC (Domain Name System Security Extensions) provides authentication for DNS responses, ensuring that they are legitimate and have not been tampered with. DNSSEC helps protect against DNS spoofing and cache poisoning attacks, which are common techniques used in DDoS attacks.

Steps:

Enable DNSSEC: In the Google Cloud Console, navigate to Cloud DNS and enable DNSSEC for your managed zones.

Configure Key Signing: Set up key signing keys (KSK) and zone signing keys (ZSK) to sign your DNS records.

Monitor DNSSEC Status: Regularly monitor the DNSSEC status and logs to ensure it is functioning correctly.

Reference::

Cloud DNS documentation

Question: 110

Your Security team believes that a former employee of your company gained unauthorized access to Google Cloud resources some time in the past 2 months by using a service account key. You need to confirm the unauthorized access and determine the user activity. What should you do?

- A. Use Security Health Analytics to determine user activity.

- B. Use the Cloud Monitoring console to filter audit logs by user.
- C. Use the Cloud Data Loss Prevention API to query logs in Cloud Storage.
- D. Use the Logs Explorer to search for user activity.

Answer: D

Explanation:

Objective: Ensure that a Cloud Storage bucket in Project A can only be readable from Project B and prevent data access or copying to Cloud Storage buckets outside the network, even with correct credentials.

Solution: Use VPC Service Controls to create a security perimeter.

Steps:

Step 1: Open the Google Cloud Console.

Step 2: Navigate to the VPC Service Controls page.

Step 3: Create a new service perimeter.

Step 4: Add Project A and Project B to the service perimeter.

Step 5: Include Cloud Storage service in the perimeter configuration.

Step 6: Define access levels to ensure that only resources within the perimeter can access the Cloud Storage bucket.

By setting up a VPC Service Controls perimeter, you can enforce security boundaries that restrict data access and movement to within defined projects, providing an extra layer of protection beyond IAM permissions.

Reference::

VPC Service Controls Overview

Configuring VPC Service Controls

Question: 111

Your company requires the security and network engineering teams to identify all network anomalies within and across VPCs, internal traffic from VMs to VMs, traffic between end locations on the internet and VMs, and traffic between VMs to Google Cloud services in production. Which method should you use?

- A. Define an organization policy constraint.
- B. Configure packet mirroring policies.
- C. Enable VPC Flow Logs on the subnet.
- D. Monitor and analyze Cloud Audit Logs.

Answer: B

Explanation:

Question: 112

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.
 - 1. Configure the option to suspend domain users not found in LDAP.
 - 2. Set up a recurring GCDS task.
- B.
 - 1. Configure the option to delete domain users not found in LDAP.
 - 2. Run GCDS after user and group lifecycle changes.
- C.
 - 1. Configure the LDAP search attributes to exclude manually created Cloud Identity users not found in LDAP.
 - 2. Set up a recurring GCDS task.
- D.
 - 1. Configure the LDAP search attributes to exclude manually created Cloud identity users not found in LDAP.
 - 2. Run GCDS after user and group lifecycle changes.

Answer: A

Explanation:

Question: 113

You are troubleshooting access denied errors between Compute Engine instances connected to a Shared VPC and BigQuery datasets. The datasets reside in a project protected by a VPC Service Controls perimeter. What should you do?

- A. Add the host project containing the Shared VPC to the service perimeter.
- B. Add the service project where the Compute Engine instances reside to the service perimeter.
- C. Create a service perimeter between the service project where the Compute Engine instances reside and the host project that contains the Shared VPC.
- D. Create a perimeter bridge between the service project where the Compute Engine instances reside and the perimeter that contains the protected BigQuery datasets.

Answer: A

Explanation:

Question: 114

You recently joined the networking team supporting your company's Google Cloud implementation.

You are tasked with familiarizing yourself with the firewall rules configuration and providing recommendations based on your networking and Google Cloud experience. What product should you recommend to detect firewall rules that are overlapped by attributes from other firewall rules with higher or equal priority?

- A. Security Command Center
- B. Firewall Rules Logging
- C. VPC Flow Logs
- D. Firewall Insights

Answer: D

Explanation:

<https://cloud.google.com/network-intelligence-center/docs/firewall-insights/concepts/overview#shadowed-firewall-rules>

Firewall Insights analyzes your firewall rules to detect firewall rules that are shadowed by other rules. A shadowed rule is a firewall rule that has all of its relevant attributes, such as its IP address and port ranges, overlapped by attributes from one or more rules with higher or equal priority, called shadowing rules.

Question: 115

The security operations team needs access to the security-related logs for all projects in their organization. They have the following requirements:

Follow the least privilege model by having only view access to logs.

Have access to Admin Activity logs.

Have access to Data Access logs.

Have access to Access Transparency logs.

Which Identity and Access Management (IAM) role should the security operations team be granted?

- A. roles/logging.privateLogViewer
- B. roles/logging.admin
- C. roles/viewer
- D. roles/logging.viewer

Answer: A

Explanation:

<https://cloud.google.com/logging/docs/access-control#considerations>

roles/logging.privateLogViewer (Private Logs Viewer) includes all the permissions contained by roles/logging.viewer, plus the ability to read Data Access audit logs in the _Default bucket.

Question: 116

You are exporting application logs to Cloud Storage. You encounter an error message that the log sinks don't support uniform bucket-level access policies. How should you resolve this error?

- A. Change the access control model for the bucket

- B. Update your sink with the correct bucket destination.
- C. Add the roles/logging.logWriter Identity and Access Management (IAM) role to the bucket for the log sink identity.
- D. Add the roles/logging.bucketWriter Identity and Access Management (IAM) role to the bucket for the log sink identity.

Answer: A

Explanation:

https://cloud.google.com/logging/docs/export/troubleshoot#errors_exporting_to_cloud_storage

<https://cloud.google.com/logging/docs/export/troubleshoot>

Unable to grant correct permissions to the destination: Even if the sink was successfully created with the correct service account permissions, this error message displays if the access control model for the Cloud Storage bucket was set to uniform access when the bucket was created. For existing Cloud Storage buckets, you can change the access control model for the first 90 days after bucket creation by using the Permissions tab. For new buckets, select the Fine-grained access control model during bucket creation. For details, see Creating Cloud Storage buckets.

Question: 117

You plan to deploy your cloud infrastructure using a CI/CD cluster hosted on Compute Engine. You want to minimize the risk of its credentials being stolen by a third party. What should you do?

- A. Create a dedicated Cloud Identity user account for the cluster. Use a strong self-hosted vault solution to store the user's temporary credentials.
- B. Create a dedicated Cloud Identity user account for the cluster. Enable the constraints/iam.disableServiceAccountCreation organization policy at the project level.
- C. Create a custom service account for the cluster. Enable the constraints/iam.disableServiceAccountKeyCreation organization policy at the project level.
- D. Create a custom service account for the cluster. Enable the constraints/iam.allowServiceAccountCredentialLifetimeExtension organization policy at the project level.

Answer: C

Explanation:

Disable service account key creation You can use the iam.disableServiceAccountKeyCreation boolean constraint to disable the creation of new external service account keys. This allows you to control the use of unmanaged long-term credentials for service accounts. When this constraint is set, usermanaged credentials cannot be created for service accounts in projects affected by the constraint. [https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-](https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-accounts#example_policy_boolean_constraint)

[accounts#example_policy_boolean_constraint](https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-accounts#example_policy_boolean_constraint)

Question: 118

You need to set up two network segments: one with an untrusted subnet and the other with a trusted subnet. You want to configure a virtual appliance such as a next-generation firewall (NGFW) to inspect all traffic between the two network segments. How should you design the network to inspect the traffic?

- A.
 - 1. Set up one VPC with two subnets: one trusted and the other untrusted.
 - 2. Configure a custom route for all traffic (0.0.0.0/0) pointed to the virtual appliance.
- B.
 - 1. Set up one VPC with two subnets: one trusted and the other untrusted.
 - 2. Configure a custom route for all RFC1918 subnets pointed to the virtual appliance.
- C.
 - 1. Set up two VPC networks: one trusted and the other untrusted, and peer them together.
 - 2. Configure a custom route on each network pointed to the virtual appliance.
- D.
 - 1. Set up two VPC networks: one trusted and the other untrusted.
 - 2. Configure a virtual appliance using multiple network interfaces, with each interface connected to one of the VPC networks.

Answer: D

Explanation:

Multiple network interfaces. The simplest way to connect multiple VPC networks through a virtual appliance is by using multiple network interfaces, with each interface connecting to one of the VPC networks. Internet and on-premises connectivity is provided over one or two separate network interfaces. With many NGFW products, internet connectivity is connected through an interface marked as untrusted in the NGFW software.

<https://cloud.google.com/architecture/best-practices-vpc-design#l7>

This architecture has multiple VPC networks that are bridged by an L7 next-generation firewall (NGFW) appliance, which functions as a multi-NIC bridge between VPC networks. An untrusted, outside VPC network is introduced to terminate hybrid interconnects and internet-based connections that terminate on the outside leg of the L7 NGFW for inspection. There are many variations on this design, but the key principle is to filter traffic through the firewall before the traffic reaches trusted VPC networks.

Question: 119

You are a member of your company's security team. You have been asked to reduce your Linux bastion host external attack surface by removing all public IP addresses. Site Reliability Engineers (SREs) require access to the bastion host from public locations so they can access the internal VPC while off-site. How should you enable this access?

- A. Implement Cloud VPN for the region where the bastion host lives.
- B. Implement OS Login with 2-step verification for the bastion host.
- C. Implement Identity-Aware Proxy TCP forwarding for the bastion host.
- D. Implement Google Cloud Armor in front of the bastion host.

Answer: C

Explanation:

Question: 120

You need to enable VPC Service Controls and allow changes to perimeters in existing environments without preventing access to resources. Which VPC Service Controls mode should you use?

- A. Cloud Run
- B. Native
- C. Enforced
- D. Dry run

Answer: D

Explanation:

Reference:: <https://cloud.google.com/vpc-service-controls/docs/service-perimeters>

In dry run mode, requests that violate the perimeter policy are not denied, only logged. Dry run mode is used to test perimeter configuration and to monitor usage of services without preventing access to resources.

<https://cloud.google.com/vpc-service-controls/docs/dry-run-mode>

Question: 121

You manage your organization's Security Operations Center (SOC). You currently monitor and detect network traffic anomalies in your Google Cloud VPCs based on packet header information. However, you want the capability to explore network flows and their payload to aid investigations. Which Google Cloud product should you use?

- A. Marketplace IDS
- B. VPC Flow Logs
- C. VPC Service Controls logs
- D. Packet Mirroring
- E. Google Cloud Armor Deep Packet Inspection

Answer: D

Explanation:

Reference:: <https://aws.amazon.com/blogs/aws/new-vpc-traffic-mirroring/>

Packet Mirroring clones the traffic of specified instances in your Virtual Private Cloud (VPC) network and forwards it for examination. Packet Mirroring captures all traffic and packet data, including payloads and headers.

<https://cloud.google.com/vpc/docs/packet-mirroring>

Question: 122

Your organization acquired a new workload. The Web and Application (App) servers will be running on Compute Engine in a newly created custom VPC. You are responsible for configuring a secure network communication solution that meets the following requirements:

Only allows communication between the Web and App tiers.

Enforces consistent network security when autoscaling the Web and App tiers.

Prevents Compute Engine Instance Admins from altering network traffic.

What should you do?

- A.
 - 1. Configure all running Web and App servers with respective network tags.
 - 2. Create an allow VPC firewall rule that specifies the target/source with respective network tags.
- B.
 - 1. Configure all running Web and App servers with respective service accounts.
 - 2. Create an allow VPC firewall rule that specifies the target/source with respective service accounts.
- C.
 - 1. Re-deploy the Web and App servers with instance templates configured with respective network tags.
 - 2. Create an allow VPC firewall rule that specifies the target/source with respective network tags.
- D.
 - 1. Re-deploy the Web and App servers with instance templates configured with respective service accounts.
 - 2. Create an allow VPC firewall rule that specifies the target/source with respective service accounts.

Answer: D

Explanation:

<https://cloud.google.com/vpc/docs/firewalls#service-accounts-vs-tags>

<https://cloud.google.com/vpc/docs/firewalls#service-accounts-vs-tags>

A service account represents an identity associated with an instance. Only one service account can be associated with an instance. You control access to the service account by controlling the grant of the Service Account User role for other IAM principals. For an IAM principal to start an instance by using a service account, that principal must have the Service Account User role to at least use that service account and appropriate permissions to create instances (for example, having the Compute Engine Instance Admin role to the project).

Question: 123

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: Dedicated Interconnect provides a direct physical connection between your on-premises network and Google's network, which is ideal for high-throughput, low-latency connections.

Request a Dedicated Interconnect from the Google Cloud Console, specifying the required bandwidth and location.

Once provisioned, set up the connection on your on-premises router and configure the BGP sessions to exchange routes with Google Cloud.

Configure private access using the restricted.googleapis.com domains in on-premises DNS configurations:

Configure your on-premises DNS server to resolve Google APIs to restricted.googleapis.com. This ensures that the traffic stays within the Google network and is not exposed to the public internet. Update your DNS settings to use restricted.googleapis.com for the necessary API endpoints.

This setup ensures that all Google Cloud API traffic is routed through the private link and subject to VPC Service Controls for additional security and compliance.

Reference::

Dedicated Interconnect Overview

Configuring DNS to use restricted.googleapis.com

Question: 124

You are working with protected health information (PHI) for an electronic health record system. The privacy officer is concerned that sensitive data is stored in the analytics system. You are tasked with anonymizing the sensitive data in a way that is not reversible. Also, the anonymized data should not preserve the character set and length. Which Google Cloud solution should you use?

- A. Cloud Data Loss Prevention with deterministic encryption using AES-SIV
- B. Cloud Data Loss Prevention with format-preserving encryption
- C. Cloud Data Loss Prevention with cryptographic hashing
- D. Cloud Data Loss Prevention with Cloud Key Management Service wrapped cryptographic keys

Answer: C

Explanation:

Use Cloud Data Loss Prevention (DLP) with cryptographic hashing:

Cloud DLP allows you to de-identify sensitive data using several techniques, including cryptographic hashing.

Choose a suitable hashing algorithm like SHA-256 for non-reversible anonymization.

This method converts the original data into a fixed-length hash that does not preserve the original data's format or

character set.

Set up a Cloud DLP job to scan your data sources, identify PHI, and apply the cryptographic hashing transformation.

Reference::

Cloud DLP Overview

De-identification with Cloud DLP

Question: 125

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:

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:

Use Cloud Build to automate your CI/CD pipeline.

Integrate Container Analysis to scan container images for vulnerabilities during the build process.

If vulnerabilities are found, configure the build to fail, preventing deployment of insecure containers.

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:

Use Binary Authorization to enforce deploy-time security policies.

Configure your CI/CD pipeline to generate attestations for container images that pass vulnerability scans.

Binary Authorization will then block deployments of any containers without valid attestations, ensuring only secure images are deployed.

Reference::

Cloud Build Overview

Container Analysis

Binary Authorization

Question: 126

Which type of load balancer should you use to maintain client IP by default while using the standard network tier?

- A. SSL Proxy
- B. TCP Proxy
- C. Internal TCP/UDP
- D. TCP/UDP Network

Answer: D

Explanation:

Use the TCP/UDP Network Load Balancer:

TCP/UDP Network Load Balancer maintains the client IP address by default when forwarding traffic to backends.

Configure a TCP/UDP Network Load Balancer with appropriate backend services and health checks. Ensure that the load balancer is using the standard network tier to comply with the requirements. Reference::

TCP/UDP Network Load Balancing

Network Service Tiers

Question: 127

You want to prevent users from accidentally deleting a Shared VPC host project. Which organizationlevel policy constraint should you enable?

- A. compute.restrictSharedVpcHostProjects
- B. compute.restrictXpnProjectLienRemoval
- C. compute.restrictSharedVpcSubnetworks
- D. compute.sharedReservationsOwnerProjects

Answer: B

Explanation:

Enable the compute.restrictXpnProjectLienRemoval organization-level policy constraint:

This constraint prevents users from removing liens from Shared VPC host projects.

By enabling this constraint, you ensure that the Shared VPC host project cannot be accidentally deleted, as liens prevent deletion without proper authorization.

Apply this constraint via the Google Cloud Console or using the gcloud command-line tool. Reference::

Organization Policy Constraints

Shared VPC

Question: 128

Users are reporting an outage on your public-facing application that is hosted on Compute Engine.

You suspect that a recent change to your firewall rules is responsible. You need to test whether your firewall rules are working properly. What should you do?

- A. Enable Firewall Rules Logging on the latest rules that were changed. Use Logs Explorer to analyze whether the rules are working correctly.
- B. Connect to a bastion host in your VPC. Use a network traffic analyzer to determine at which point your requests are being blocked.
- C. In a pre-production environment, disable all firewall rules individually to determine which one is blocking user traffic.
- D. Enable VPC Flow Logs in your VPC. Use Logs Explorer to analyze whether the rules are working correctly.

Answer: A

Explanation:

Enable Firewall Rules Logging on the latest rules that were changed. Use Logs Explorer to analyze whether the rules are working correctly:

Enable Firewall Rules Logging for the specific firewall rules in question through the Google Cloud Console.

Once logging is enabled, use Logs Explorer to filter and review the firewall logs.

Analyze the logs to determine if the rules are allowing or blocking traffic as intended, identifying any misconfigurations or issues.

Reference::

Firewall Rules Logging

Using Logs Explorer

Question: 129

You are a security administrator at your company. Per Google-recommended best practices, you implemented the domain restricted sharing organization policy to allow only required domains to access your projects. An engineering team is now reporting that users at an external partner outside your organization domain cannot be granted access to the resources in a project. How should you make an exception for your partner's domain while following the stated best practices?

- A. Turn off the domain restriction sharing organization policy. Set the policy value to "Allow All."
- B. Turn off the domain restricted sharing organization policy. Provide the external partners with the required permissions using Google's Identity and Access Management (IAM) service.
- C. Turn off the domain restricted sharing organization policy. Add each partner's Google Workspace customer ID to a Google group, add the Google group as an exception under the organization policy, and then turn the policy back on.
- D. Turn off the domain restricted sharing organization policy. Set the policy value to "Custom." Add each external partner's

Cloud Identity or Google Workspace customer ID as an exception under the organization policy, and then turn the policy back on.

Answer: D

Explanation:

https://cloud.google.com/resource-manager/docs/organization-policy/restricting-domains#setting_the_organization_policy

The domain restriction constraint is a type of list constraint. Google Workspace customer IDs can be added and removed from the allowed_values list of a domain restriction constraint. The domain restriction constraint does not support denying values, and an organization policy can't be saved with IDs in the denied_values list. All domains associated with a Google Workspace account listed in the allowed_values will be allowed by the organization policy. All other domains will be denied by the organization policy.

Question: 130

You plan to use a Google Cloud Armor policy to prevent common attacks such as cross-site scripting (XSS) and SQL injection (SQLi) from reaching your web application's backend. What are two requirements for using Google Cloud Armor security policies? (Choose two.)

- A. The load balancer must be an external SSL proxy load balancer.
- B. Google Cloud Armor Policy rules can only match on Layer 7 (L7) attributes.
- C. The load balancer must use the Premium Network Service Tier.
- D. The backend service's load balancing scheme must be EXTERNAL.
- E. The load balancer must be an external HTTP(S) load balancer.

Answer: D, E

Explanation:

Google Cloud Armor helps to protect applications from DDoS attacks and web application firewall (WAF) threats like XSS and SQLi. To use Google Cloud Armor security policies, certain requirements must be met:

External Load Balancers: Google Cloud Armor is specifically designed to work with external HTTP(S) load balancers, which handle traffic at the edge of the Google Cloud network. This type of load balancer provides a global frontend that can distribute traffic to various backends.

Load Balancing Scheme: The backend service associated with the Google Cloud Armor policy must have an EXTERNAL load balancing scheme. This scheme allows the service to accept traffic from outside the Google Cloud network, which is necessary for applying security policies effectively at the network edge.

These requirements ensure that Google Cloud Armor can inspect and filter incoming traffic before it reaches your web application's backend services, providing an additional layer of security against common web vulnerabilities.

Reference:

Google Cloud Armor Documentation

External HTTP(S) Load Balancing Overview

Question: 131

You perform a security assessment on a customer architecture and discover that multiple VMs have public IP addresses. After providing a recommendation to remove the public IP addresses, you are told those VMs need to communicate to external sites as part of the customer's typical operations. What should you recommend to reduce the need for public IP addresses in your customer's VMs?

- A. Google Cloud Armor
- B. Cloud NAT
- C. Cloud Router
- D. Cloud VPN

Answer: B

Explanation:

Cloud NAT (Network Address Translation) enables instances in a private network to connect to external services while not exposing their internal IP addresses to the public internet. This solution helps in situations where VMs need to initiate outbound connections without having a public IP address:

Cloud NAT Setup: Configure Cloud NAT for the subnet where your VMs are located. This allows these VMs to use the NAT gateway to communicate with external services securely.

Network Security: By using Cloud NAT, the internal IP addresses of VMs remain private, reducing the attack surface and enhancing security.

Operational Continuity: VMs can continue to communicate with external sites as needed for operations without requiring public IP addresses, meeting both security and functional requirements.

Reference:

Cloud NAT Documentation

Question: 132

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: BC

Explanation:

To meet the requirements for exporting and auditing security logs in near real-time for login activity events and API calls:

Organization-Level Log Sink with Pub/Sub: Create a log sink at the organization level to ensure logs from all projects are captured. Use the `includeChildren` parameter to include logs from all child resources (projects). Set the destination to a Pub/Sub topic to facilitate near real-time log export to an external SIEM.

Enable Data Access Audit Logs: Enable Data Access audit logs at the organization level to capture and export detailed information about API calls that modify configurations of Google Cloud resources across all projects.

These steps ensure comprehensive logging and real-time export capabilities, which are crucial for security auditing and monitoring.

Reference:

Audit Logs Overview

Exporting with Sinks

Question: 133

Your company's Chief Information Security Officer (CISO) creates a requirement that business data must be stored in specific locations due to regulatory requirements that affect the company's global expansion plans. After working on the details to implement this requirement, you determine the following:

The services in scope are included in the Google Cloud Data Residency Terms.

The business data remains within specific locations under the same organization.

The folder structure can contain multiple data residency locations.

You plan to use the Resource Location Restriction organization policy constraint. At which level in the resource hierarchy should you set the constraint?

- A. Folder
- B. Resource
- C. Project
- D. Organization

Answer: D

Explanation:

The Resource Location Restriction organization policy constraint ensures that business data is stored in specific geographic locations, which is critical for compliance with regulatory requirements. Organization Level: Setting the constraint at the organization level ensures that all resources within the organization, including those in different folders or projects, adhere to the location restrictions.

This provides a unified policy application across the entire organization, ensuring compliance with regulatory

requirements.

Policy Application: The policy will propagate down the resource hierarchy, ensuring that all relevant services within the organization comply with the specified data residency requirements.

This approach provides centralized control and simplifies the management of data residency constraints.

Reference:

Organization Policy Service Documentation

Question: 134

You need to set up a Cloud interconnect connection between your company's on-premises data center and VPC host network. You want to make sure that on-premises applications can only access Google APIs over the Cloud Interconnect and not through the public internet. You are required to only use APIs that are supported by VPC Service Controls to mitigate against exfiltration risk to nonsupported APIs. How should you configure the network?

- A. Enable Private Google Access on the regional subnets and global dynamic routing mode.
- B. Set up a Private Service Connect endpoint IP address with the API bundle of "all-apis", which is advertised as a route over the Cloud interconnect connection.
- C. Use private.googleapis.com to access Google APIs using a set of IP addresses only routable from within Google Cloud, which are advertised as routes over the connection.
- D. Use restricted.googleapis.com to access Google APIs using a set of IP addresses only routable from within Google Cloud, which are advertised as routes over the Cloud Interconnect connection.

Answer: D

Explanation:

<https://cloud.google.com/vpc/docs/private-service-connect>

An API bundle:

All APIs (all-apis): most Google APIs

(same as private.googleapis.com).

VPC-SC (vpc-sc): APIs that VPC Service Controls supports

(same as restricted.googleapis.com).

VMs in the same VPC network as the endpoint (all regions)

On-premises systems that are connected to the VPC network that contains the endpoint

Question: 135

You need to implement an encryption-at-rest strategy that protects sensitive data and reduces key management complexity for non-sensitive data

a. Your solution has the following requirements:

Schedule key rotation for sensitive data.

Control which region the encryption keys for sensitive data are stored in.

Minimize the latency to access encryption keys for both sensitive and non-sensitive data.

What should you do?

- A. Encrypt non-sensitive data and sensitive data with Cloud External Key Manager.
- B. Encrypt non-sensitive data and sensitive data with Cloud Key Management Service.
- C. Encrypt non-sensitive data with Google default encryption, and encrypt sensitive data with Cloud External Key Manager.
- D. Encrypt non-sensitive data with Google default encryption, and encrypt sensitive data with Cloud Key Management Service.

Answer: D

Explanation:

Encrypt non-sensitive data with Google default encryption:

Google Cloud automatically encrypts data at rest using AES-256 by default. This minimizes key management complexity for non-sensitive data as it is handled entirely by Google.

No additional setup is required for default encryption, ensuring low latency access to the encrypted data.

Encrypt sensitive data with Cloud Key Management Service (Cloud KMS):

Cloud KMS allows you to create and manage cryptographic keys in a centralized cloud service.

To meet the requirement of scheduling key rotation, configure Cloud KMS to automatically rotate keys on a regular schedule (e.g., every 90 days).

Control the region where the keys are stored by selecting the appropriate key ring location during key creation. This ensures compliance with data residency requirements.

Cloud KMS provides low-latency access to keys, ensuring minimal impact on data access performance.

Reference::

Cloud Key Management Service Documentation

Encryption at Rest in Google Cloud

Question: 136

Your security team uses encryption keys to ensure confidentiality of user data.

a. 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: BD

Explanation:

Enable automatic key version rotation on a regular schedule:

Regularly rotating keys reduces the impact of a potentially compromised key by limiting the amount of data encrypted with a single key version.

Set up automatic key rotation in Cloud KMS to ensure keys are rotated without manual intervention.

Limit the number of messages encrypted with each key version:

Reducing the number of messages encrypted with each key version minimizes the potential data exposure in case of a key compromise.

Implement policies to ensure that new key versions are used periodically to limit the usage of each key version.

Reference::

Cloud KMS Key Rotation

Best Practices for Using Cryptographic Keys

Question: 137

Your company's chief information security officer (CISO) is requiring business data to be stored in specific locations due to regulatory requirements that affect the company's global expansion plans. After working on a plan to implement this requirement, you determine the following:

The services in scope are included in the Google Cloud data residency requirements.

The business data remains within specific locations under the same organization.

The folder structure can contain multiple data residency locations.

The projects are aligned to specific locations.

You plan to use the Resource Location Restriction organization policy constraint with very granular control. At which level in the hierarchy should you set the constraint?

- A. Organization
- B. Resource
- C. Project
- D. Folder

Answer: D

Explanation:

Set the Resource Location Restriction organization policy constraint at the Folder level: Setting the constraint at the folder level allows for very granular control over data residency requirements.

Each folder can represent a specific geographical location, and projects within those folders will inherit the location constraint, ensuring compliance with regulatory requirements.

This approach provides flexibility to manage data residency across different regions within the same organization.

Reference::

Organization Policy Service: Resource Location Restriction

Question: 138

A database administrator notices malicious activities within their Cloud SQL instance. The database administrator wants to

monitor the API calls that read the configuration or metadata of resources. Which logs should the database administrator review?

- A. Admin Activity
- B. System Event
- C. Access Transparency
- D. Data Access

Answer: A

Explanation:

Review Admin Activity logs:

Admin Activity logs contain entries for API calls that modify or read the configuration or metadata of resources. These logs are useful for monitoring and auditing administrative actions, including those that could indicate malicious activity on a Cloud SQL instance.

Reference::

Audit Logs: Admin Activity

Question: 139

You are backing up application logs to a shared Cloud Storage bucket that is accessible to both the administrator and analysts. Analysts should not have access to logs that contain any personally identifiable information (PII). Log files containing PII should be stored in another bucket that is only accessible to the administrator. What should you do?

- A. Upload the logs to both the shared bucket and the bucket with PII that is only accessible to the administrator. Use the Cloud Data Loss Prevention API to create a job trigger. Configure the trigger to delete any files that contain PII from the shared bucket.
- B. On the shared bucket, configure Object Lifecycle Management to delete objects that contain PII.
- C. On the shared bucket, configure a Cloud Storage trigger that is only triggered when PII is uploaded. Use Cloud Functions to capture the trigger and delete the files that contain PII.
- D. Use Pub/Sub and Cloud Functions to trigger a Cloud Data Loss Prevention scan every time a file is uploaded to the administrator's bucket. If the scan does not detect PII, have the function move the objects into the shared Cloud Storage bucket.

Answer: D

Explanation:

Use Pub/Sub and Cloud Functions to trigger a Cloud Data Loss Prevention scan every time a file is uploaded to the administrator's bucket. If the scan does not detect PII, have the function move the objects into the shared Cloud Storage bucket:

Configure a Pub/Sub topic to publish notifications when new files are uploaded to the administrator's bucket.

Create a Cloud Function that is triggered by the Pub/Sub topic. This function uses the Cloud Data Loss Prevention (DLP) API to scan the uploaded files for PII.

If the scan does not detect PII, the function moves the file to the shared Cloud Storage bucket. This ensures that only non-sensitive data is accessible to analysts, while PII remains secure in the administrator's bucket.

Reference::

Using Pub/Sub with Cloud Functions

Cloud Data Loss Prevention API

Question: 140

You work for an organization in a regulated industry that has strict data protection requirements. The organization backs up their data in the cloud. To comply with data privacy regulations, this data can only be stored for a specific length of time and must be deleted after this specific period.

You want to automate the compliance with this regulation while minimizing storage costs. What should you do?

- A. Store the data in a persistent disk, and delete the disk at expiration time.
- B. Store the data in a Cloud Bigtable table, and set an expiration time on the column families.
- C. Store the data in a BigQuery table, and set the table's expiration time.
- D. Store the data in a Cloud Storage bucket, and configure the bucket's Object Lifecycle Management feature.

Answer: D

Explanation:

Google Cloud Storage provides an Object Lifecycle Management feature that can help automate data retention and deletion processes, ensuring compliance with data privacy regulations while

minimizing storage costs.

Lifecycle Management: Object Lifecycle Management allows you to define rules that automatically delete objects after a specific period. This ensures that data is only retained for the required amount of time and is deleted once it expires.

Configuration: You can configure lifecycle rules to delete objects based on conditions such as the age of the object, the creation date, or custom metadata. This allows for precise control over the retention period of your data.

Cost Efficiency: By using lifecycle policies to delete data automatically, you can reduce storage costs, as you only pay for the storage you actively use.

Reference:

Cloud Storage Object Lifecycle Management

Question: 141

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: AC

Explanation:

Security Command Center (SCC) in Google Cloud provides several features to help organizations detect and respond to security threats and misconfigurations.

Event Threat Detection: This feature continuously monitors and analyzes system logs to detect potential threats such as crypto mining. It uses machine learning and threat intelligence to identify suspicious activities and generate alerts.

Security Health Analytics: This feature helps identify common misconfigurations and compliance violations that could impact security. It provides visibility into security posture and helps remediate issues related to misconfigurations in your Google Cloud environment.

By using both Event Threat Detection and Security Health Analytics, you can effectively monitor for crypto mining activities and detect common misconfigurations that could compromise security.

Reference:

Security Command Center Documentation

Event Threat Detection

Security Health Analytics

Question: 142

You have noticed an increased number of phishing attacks across your enterprise user accounts. You want to implement the Google 2-Step Verification (2SV) option that uses a cryptographic signature to authenticate a user and verify the URL of the login page. Which Google 2SV option should you use?

- A. Titan Security Keys
- B. Google prompt
- C. Google Authenticator app
- D. Cloud HSM keys

Answer: A

Explanation:

Titan Security Keys are a physical form of two-step verification (2SV) that provide the highest level of account security by using cryptographic signatures to verify the user and the URL of the login page. Cryptographic Security: Titan Security Keys

use a hardware-based cryptographic method to authenticate users, which is resistant to phishing attacks. This ensures that the authentication process is secure and not susceptible to being intercepted or spoofed.

URL Verification: Titan Security Keys verify the URL of the login page during the authentication process, providing an additional layer of security against phishing attempts that may try to redirect users to malicious websites.

Ease of Use: These keys are easy to use and integrate with Google's 2SV process, providing a seamless and highly secure authentication method for users.

Reference:

Titan Security Keys

Question: 143

Your organization hosts a financial services application running on Compute Engine instances for a third-party company. The third-party company's servers that will consume the application also run on Compute Engine in a separate Google Cloud organization. You need to configure a secure network connection between the Compute Engine instances. You have the following requirements: The network connection must be encrypted.

The communication between servers must be over private IP addresses.

What should you do?

- A. Configure a Cloud VPN connection between your organization's VPC network and the third party's that is controlled by VPC firewall rules.
- B. Configure a VPC peering connection between your organization's VPC network and the third party's that is controlled by VPC firewall rules.
- C. Configure a VPC Service Controls perimeter around your Compute Engine instances, and provide access to the third party via an access level.
- D. Configure an Apigee proxy that exposes your Compute Engine-hosted application as an API, and is encrypted with TLS which allows access only to the third party.

Answer: A

Explanation:

To meet the requirements of encrypted communication over private IP addresses between Compute Engine instances in different Google Cloud organizations, a Cloud VPN connection is appropriate: Cloud VPN: Cloud VPN creates a secure, encrypted tunnel between your organization's VPC network and the third party's VPC network. This ensures that data transmitted over the network is encrypted and secure.

Private IP Communication: Cloud VPN allows communication over private IP addresses, which helps maintain security by keeping traffic within the Google Cloud network and not exposing it to the public internet.

Firewall Rules: VPC firewall rules can be configured to control the traffic that flows through the VPN, ensuring that only authorized traffic is allowed, further enhancing security.

By setting up a Cloud VPN connection, you can achieve secure, encrypted communication over private IP addresses between different Google Cloud organizations.

Reference:

Cloud VPN Overview

Question: 144

Your company's new CEO recently sold two of the company's divisions. Your Director asks you to help migrate the Google Cloud projects associated with those divisions to a new organization node. Which preparation steps are necessary before this migration occurs? (Choose two.)

- A. Remove all project-level custom Identity and Access Management (IAM) roles.
- B. Disallow inheritance of organization policies.
- C. Identify inherited Identity and Access Management (IAM) roles on projects to be migrated.
- D. Create a new folder for all projects to be migrated.
- E. Remove the specific migration projects from any VPC Service Controls perimeters and bridges.

Answer: C E

Explanation:

To prepare for migrating Google Cloud projects to a new organization node, it's crucial to ensure that the projects' current configurations and dependencies are appropriately managed. The two necessary preparation steps are:

Identify inherited Identity and Access Management (IAM) roles on projects to be migrated (C): Projects inherit IAM roles from their parent resources. Identifying these roles is essential to understand the permissions and access levels that users have on the projects. This will help in ensuring that after migration, the appropriate roles and permissions are applied correctly.

Remove the specific migration projects from any VPC Service Controls perimeters and bridges (E): VPC Service Controls provide security boundaries around your Google Cloud resources to mitigate data exfiltration risks. Before migrating the projects, they need to be removed from any existing VPC Service Controls perimeters and bridges to prevent any disruption in access or network communication. After migration, the projects can be added back to the necessary perimeters.

Reference:

Google Cloud IAM documentation
VPC Service Controls documentation

Question: 145

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:

To meet data residency requirements on Google Cloud, the recommended option is to use Organization Policy Service constraints. This service allows you to define and enforce specific constraints across your organization, including constraints related to the geographical location where data is stored.

Organization Policy Service constraints allow administrators to enforce policies that restrict resources to specific locations. For instance, you can set policies to ensure that all storage buckets, databases, and other data resources reside within specific geographic boundaries. This helps in complying with data residency requirements.

Reference:

Organization Policy Service documentation

Google Cloud Data Residency

Question: 146

Your security team wants to reduce the risk of user-managed keys being mismanaged and compromised. To achieve this, you need to prevent developers from creating user-managed service account keys for projects in their organization. How should you enforce this?

- A. Configure Secret Manager to manage service account keys.
- B. Enable an organization policy to disable service accounts from being created.
- C. Enable an organization policy to prevent service account keys from being created.
- D. Remove the iam.serviceAccounts.getAccessToken permission from users.

Answer: C

Explanation:

To prevent developers from creating user-managed service account keys and reduce the risk of key mismanagement, you should enable an organization policy that specifically prohibits the creation of these keys.

Enable an organization policy to prevent service account keys from being created (C):

Google Cloud provides the capability to enforce organizational policies that restrict various actions, including the creation of service account keys. By enabling this policy, you ensure that developers cannot create new user-managed service account keys, thus minimizing the risk of key mismanagement and potential security breaches.

Reference:

Service Accounts documentation

Organization Policy Service documentation

Question: 147

You are responsible for managing your company's identities in Google Cloud. Your company enforces 2-Step Verification (2SV) for all users. You need to reset a user's access, but the user lost their second factor for 2SV. You want to minimize risk. What should you do?

- A. On the Google Admin console, select the appropriate user account, and generate a backup code to allow the user to sign in. Ask the user to update their second factor.
- B. On the Google Admin console, temporarily disable the 2SV requirements for all users. Ask the user to log in and add their new second factor to their account. Re-enable the 2SV requirement for all users.
- C. On the Google Admin console, select the appropriate user account, and temporarily disable 2SV for this account. Ask the user to update their second factor, and then re-enable 2SV for this account.
- D. On the Google Admin console, use a super administrator account to reset the user account's credentials. Ask the user to update their credentials after their first login.

Answer: A

Explanation:

If a user loses their second factor for 2-Step Verification (2SV), you can help them regain access with minimal risk by generating a backup code.

Generate a Backup Code (A):

In the Google Admin console, navigate to the user's account settings.

Generate a backup code for the user. This code allows them to sign in despite not having access to their usual second factor.

Instruct the user to log in using the backup code and then update their second factor in their account settings.

This method ensures that only the affected user's access is temporarily adjusted, minimizing risk while maintaining overall security policies.

Reference:

Google Admin console 2-Step Verification documentation

Question: 148

Which Google Cloud service should you use to enforce access control policies for applications and resources?

- A. Identity-Aware Proxy
- B. Cloud NAT
- C. Google Cloud Armor
- D. Shielded VMs

Answer: A

Explanation:

To enforce access control policies for applications and resources in Google Cloud, the recommended service is Identity-Aware Proxy (IAP).

Identity-Aware Proxy (IAP):

IAP allows you to control access to your applications and resources based on the identity of the user and the context of the request. It integrates with IAM to provide fine-grained access control, ensuring that only authorized users can access specific resources.

IAP helps enforce security policies at the application layer, providing an additional layer of protection beyond traditional network-based security measures.

Reference:

Identity-Aware Proxy documentation

Question: 149

You want to update your existing VPC Service Controls perimeter with a new access level. You need to avoid breaking the existing perimeter with this change, and ensure the least disruptions to users while minimizing overhead. What should you do?

- A. Create an exact replica of your existing perimeter. Add your new access level to the replica. Update the original perimeter after the access level has been vetted.
- B. Update your perimeter with a new access level that never matches. Update the new access level to match your desired state one condition at a time to avoid being overly permissive.
- C. Enable the dry run mode on your perimeter. Add your new access level to the perimeter configuration. Update the perimeter configuration after the access level has been vetted.
- D. Enable the dry run mode on your perimeter. Add your new access level to the perimeter dry run configuration. Update the perimeter configuration after the access level has been vetted.

Answer: D

Explanation:

Enable Dry Run Mode: Start by enabling the dry run mode for your VPC Service Controls perimeter. This mode allows you to test changes without actually enforcing them, thus preventing any disruption to your current setup.

Add Access Level: Add your new access level to the dry run configuration. This way, you can monitor how the new access level would behave and interact with your existing setup without any real impact.

Vetting Process: Carefully vet the new access level by analyzing logs and monitoring the behavior in the dry run mode.

Ensure that the new configuration meets your security and operational requirements.

Update Perimeter: Once you are confident that the new access level will not disrupt existing services and meets all requirements, update the actual perimeter configuration with the new access level. This approach minimizes risk by allowing you to test changes before they take effect, ensuring seamless updates with minimal disruption. Reference::

Google Cloud - Configuring VPC Service Controls

Google Cloud - Using Dry Run Mode

Question: 150

Your organization's Google Cloud VMs are deployed via an instance template that configures them with a public IP address in order to host web services for external users. The VMs reside in a service project that is attached to a host (VPC) project containing one custom Shared VPC for the VMs. You have been asked to reduce the exposure of the VMs to the internet while continuing to service external users. You have already recreated the instance template without a public IP address configuration to launch the managed instance group (MIG). What should you do?

- A. Deploy a Cloud NAT Gateway in the service project for the MIG.
- B. Deploy a Cloud NAT Gateway in the host (VPC) project for the MIG.
- C. Deploy an external HTTP(S) load balancer in the service project with the MIG as a backend.
- D. Deploy an external HTTP(S) load balancer in the host (VPC) project with the MIG as a backend.

Answer: D

Explanation:

External HTTP(S) Load Balancer: Deploy an external HTTP(S) load balancer to manage traffic to your VMs. This load balancer will handle incoming traffic from the internet while the VMs themselves do not have public IP addresses.

Host (VPC) Project Deployment: Deploy the load balancer in the host (VPC) project. This allows for centralized management of network resources and maintains the integrity of your shared VPC configuration.

Backend Configuration: Configure the MIG as the backend for the load balancer. This setup ensures that the VMs can still serve external users while reducing their direct exposure to the internet. This solution provides the required access to external users through the load balancer, enhancing security by not exposing individual VM IP addresses.

Reference::

Google Cloud - External HTTP(S) Load Balancer Overview

[Google Cloud - Shared VPC Overview](#)

Question: 151

Your privacy team uses crypto-shredding (deleting encryption keys) as a strategy to delete personally identifiable information (PII). You need to implement this practice on Google Cloud while still utilizing the majority of the platform's services and minimizing operational overhead. What should you do?

- A. Use client-side encryption before sending data to Google Cloud, and delete encryption keys on premises
- B. Use Cloud External Key Manager to delete specific encryption keys.
- C. Use customer-managed encryption keys to delete specific encryption keys.
- D. Use Google default encryption to delete specific encryption keys.

Answer: C

Explanation:

<https://cloud.google.com/sql/docs/mysql/cmek>

"You might have situations where you want to permanently destroy data encrypted with CMEK. To do this, you destroy the customer-managed encryption key version. You can't destroy the keyring or key, but you can destroy key versions of the key."

Question: 152

You need to centralize your team's logs for production projects. You want your team to be able to search and analyze the logs using Logs Explorer. What should you do?

- A. Enable Cloud Monitoring workspace, and add the production projects to be monitored.
- B. Use Logs Explorer at the organization level and filter for production project logs.
- C. Create an aggregate org sink at the parent folder of the production projects, and set the destination to a Cloud Storage bucket.
- D. Create an aggregate org sink at the parent folder of the production projects, and set the destination to a logs bucket.

Answer: D

Explanation:

https://cloud.google.com/logging/docs/export/aggregated_sinks#supported-destinations

You can use aggregated sinks to route logs within or between the same organizations and folders to the following destinations: - Another Cloud Logging bucket: Log entries held in Cloud Logging log buckets.

Question: 153

You need to use Cloud External Key Manager to create an encryption key to encrypt specific BigQuery data at rest in Google Cloud. Which steps should you do first?

- A. 1. Create or use an existing key with a unique uniform resource identifier (URI) in your Google Cloud project.
2. Grant your Google Cloud project access to a supported external key management partner system.
- B. 1. Create or use an existing key with a unique uniform resource identifier (URI) in Cloud Key Management Service (Cloud KMS).
2. In Cloud KMS, grant your Google Cloud project access to use the key.
- C. 1. Create or use an existing key with a unique uniform resource identifier (URI) in a supported external key management partner system.
2. In the external key management partner system, grant access for this key to use your Google Cloud project.
- D. 1. Create an external key with a unique uniform resource identifier (URI) in Cloud Key Management Service (Cloud KMS).
2. In Cloud KMS, grant your Google Cloud project access to use the key.

Answer: C

Explanation:

https://cloud.google.com/kms/docs/ekm#how_it_works

- First, you create or use an existing key in a supported external key management partner system. This key has a unique URI or key path.
- Next, you grant your Google Cloud project access to use the key, in the external key management partner system.
- In your Google Cloud project, you create a Cloud EKM key, using the URI or key path for the externally-managed key.

Question: 154

Your company's cloud security policy dictates that VM instances should not have an external IP address. You need to identify the Google Cloud service that will allow VM instances without external IP addresses to connect to the internet to update the VMs. Which service should you use?

- A. Identity Aware-Proxy

- B. Cloud NAT
- C. TCP/UDP Load Balancing
- D. Cloud DNS

Answer: B

Explanation:

Cloud NAT Service: Use Cloud NAT (Network Address Translation) to allow VM instances without external IP addresses to access the internet securely.

Configuration: Configure Cloud NAT for the subnets containing your VM instances. This setup allows the VMs to initiate outbound connections to the internet for updates and other necessary communications.

Security Compliance: By using Cloud NAT, you adhere to the security policy of not assigning external IP addresses to VMs while still enabling necessary internet connectivity. Cloud NAT provides a secure method for outbound internet traffic without exposing VMs directly to the public internet.

Reference::

Google Cloud - Cloud NAT Overview

Google Cloud - Configuring Cloud NAT

Question: 155

You want to make sure that your organization's Cloud Storage buckets cannot have data publicly available to the internet. You want to enforce this across all Cloud Storage buckets. What should you do?

- A. Remove Owner roles from end users, and configure Cloud Data Loss Prevention.
- B. Remove Owner roles from end users, and enforce domain restricted sharing in an organization policy.
- C. Configure uniform bucket-level access, and enforce domain restricted sharing in an organization policy.
- D. Remove *.setIamPolicy permissions from all roles, and enforce domain restricted sharing in an organization policy.

Answer: C

Explanation:

Uniform Bucket-Level Access: Enable uniform bucket-level access for all your Cloud Storage buckets. This feature ensures that access control is applied consistently at the bucket level, simplifying management and improving security.

Domain Restricted Sharing: Enforce domain-restricted sharing through an organization policy. This policy ensures that only users within your organization's domain can access the data in the buckets, preventing public exposure.

Policy Enforcement: Apply the necessary IAM policies and ensure that no buckets are configured to allow public access. This combination of settings ensures that data in Cloud Storage buckets remains private and accessible only to authorized users within your organization. **Reference::**

Google Cloud - Uniform Bucket-Level Access

Google Cloud - Organization Policy Service

Question: 156

Your company plans to move most of its IT infrastructure to Google Cloud. They want to leverage their existing on-premises Active Directory as an identity provider for Google Cloud. Which two steps should you take to integrate the company's on-premises Active Directory with Google Cloud and configure access management? (Choose two.)

- A. Use Identity Platform to provision users and groups to Google Cloud.
- B. Use Cloud Identity SAML integration to provision users and groups to Google Cloud.
- C. Install Google Cloud Directory Sync and connect it to Active Directory and Cloud Identity.
- D. Create Identity and Access Management (IAM) roles with permissions corresponding to each Active Directory group.
- E. Create Identity and Access Management (IAM) groups with permissions corresponding to each Active Directory group.

Answer: CE

Explanation:

Google Cloud Directory Sync (GCDS): Install and configure GCDS to synchronize your on-premises Active Directory with Google Cloud Identity. This tool helps in maintaining consistency between your local directory and Google Cloud.

IAM Groups: Create IAM groups in Google Cloud with permissions that correspond to your Active Directory groups. This mapping ensures that users inherit the appropriate permissions based on their AD group membership.

Synchronization: Set up regular synchronization schedules to keep the user and group information up-to-date between your on-premises AD and Google Cloud.

Access Management: Use these IAM groups to manage access to Google Cloud resources, ensuring that permissions are applied consistently and securely. This approach leverages existing AD infrastructure for identity management, providing a seamless integration with Google Cloud. **Reference::**

Google Cloud - Google Cloud Directory Sync

Google Cloud - IAM Groups

Question: 157

You are in charge of creating a new Google Cloud organization for your company. Which two actions should you take when creating the super administrator accounts? (Choose two.)

- A. Create an access level in the Google Admin console to prevent super admin from logging in to Google Cloud.
- B. Disable any Identity and Access Management (IAM) roles for super admin at the organization level in the Google Cloud Console.
- C. Use a physical token to secure the super admin credentials with multi-factor authentication (MFA).
- D. Use a private connection to create the super admin accounts to avoid sending your credentials over the Internet.
- E. Provide non-privileged identities to the super admin users for their day-to-day activities.

Answer: CE

Explanation:

Physical Token for MFA: Implement multi-factor authentication (MFA) using physical tokens (such as security keys) for super admin accounts. This adds an extra layer of security to the highest privilege accounts.

Non-Privileged Identities: Provide super admins with separate non-privileged accounts for daily activities. This practice minimizes the risk associated with using highly privileged accounts for routine tasks.

Account Management: Ensure that super admin accounts are only used for tasks requiring elevated privileges, reducing exposure to potential security threats. These measures enhance the security of super admin accounts, protecting your

Google Cloud organization from unauthorized access. Reference::

Google Cloud - Best Practices for Securing Cloud Identity

Google Cloud - Using Security Keys

Question: 158

You are deploying a web application hosted on Compute Engine. A business requirement mandates that application logs are preserved for 12 years and data is kept within European boundaries. You want to implement a storage solution that minimizes overhead and is cost-effective. What should you do?

- A. Create a Cloud Storage bucket to store your logs in the EUROPE-WEST1 region. Modify your application code to ship logs directly to your bucket for increased efficiency.
- B. Configure your Compute Engine instances to use the Google Cloud's operations suite Cloud Logging agent to send application logs to a custom log bucket in the EUROPE-WEST1 region with a custom retention of 12 years.
- C. Use a Pub/Sub topic to forward your application logs to a Cloud Storage bucket in the EUROPE- WEST1 region.
- D. Configure a custom retention policy of 12 years on your Google Cloud's operations suite log bucket in the EUROPE-WEST1 region.

Answer: B

Explanation:

To fulfill the requirements of preserving logs for 12 years and ensuring data residency within European boundaries, the best approach is to use Google Cloud's operations suite (formerly Stackdriver) with a custom log bucket configured in the desired region.

Configure Cloud Logging Agent:

Install and configure the Cloud Logging agent on your Compute Engine instances. This agent collects logs from your application and system and sends them to Google Cloud's operations suite.

Create a Custom Log Bucket:

In the Cloud Logging interface, create a custom log bucket in the EUROPE-WEST1 region. This bucket will store your logs and can be configured with a custom retention period.

Set Custom Retention Policy:

Configure the retention policy for the custom log bucket to 12 years. This ensures that all logs are preserved for the required duration.

Ship Logs to the Custom Log Bucket:

Modify the logging configuration to direct logs from the Cloud Logging agent to the custom log bucket. This can be done through the logging configuration settings in the Cloud Console or by updating the agent configuration files.

This solution minimizes overhead by using managed services and ensures cost-effectiveness by leveraging Cloud Logging's

built-in capabilities for log storage and retention management.

Reference:

Cloud Logging Documentation
Creating and Managing Logs Buckets

Question: 159

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: C E

Explanation:

To handle PII data ingestion and ensure both redaction and re-identification for analytics purposes, you can use Cloud Data Loss Prevention (DLP) with appropriate techniques for masking and encryption.

Cloud Data Loss Prevention (DLP) with Cryptographic Hashing (C):

Use Cloud DLP to apply cryptographic hashing to PII data. Hashing transforms the data into a fixed-length string that is not directly readable, providing a layer of obfuscation. This helps in masking the PII while retaining the ability to verify data integrity.

Cloud Data Loss Prevention (DLP) with Deterministic Encryption using AES-SIV (E):

Apply deterministic encryption using AES-SIV through Cloud DLP. Deterministic encryption ensures that the same input will always produce the same encrypted output, allowing you to re-identify the PII when necessary. This method enables secure encryption while allowing data re-identification for analytics.

By combining these two approaches, you can effectively mask PII for privacy protection and later reidentify it when required for analysis.

Reference:

Cloud Data Loss Prevention Documentation
Data Redaction and Masking Techniques

Question: 160

You are working with a client that is concerned about control of their encryption keys for sensitive data.

F. The client does not want to store encryption keys at rest in the same cloud service provider (CSP) as the data that the keys are encrypting. Which Google Cloud encryption solutions should you recommend to this client? (Choose two.)

- A. Customer-supplied encryption keys.
- B. Google default encryption
- C. Secret Manager
- D. Cloud External Key Manager
- E. Customer-managed encryption keys

Answer: AD

Explanation:

For a client concerned about the control of their encryption keys and not wanting to store these keys within the same cloud service provider (CSP) as the data, the following solutions are suitable: Customer-supplied encryption keys (A):

With customer-supplied encryption keys, clients manage their own encryption keys outside of Google Cloud and supply them to encrypt and decrypt data. This ensures that the keys are not stored in Google Cloud, providing full control over the key management process.

Cloud External Key Manager (D):

Cloud External Key Manager (EKM) allows clients to integrate an external key management system (KMS) with Google Cloud services. This setup enables the client to keep their encryption keys outside Google Cloud while still allowing the data to be encrypted and decrypted within Google Cloud services. This method offers an additional layer of security and control over the encryption keys. These options provide robust solutions for clients requiring external key management and enhanced control over their encryption processes.

Reference:

Customer-Supplied Encryption Keys

Cloud External Key Manager

Question: 161

You are implementing data protection by design and in accordance with GDPR requirements. As part of design reviews, you are told that you need to manage the encryption key for a solution that includes workloads for Compute Engine, Google Kubernetes Engine, Cloud Storage, BigQuery, and Pub/Sub. Which option should you choose for this implementation?

- A. Cloud External Key Manager
- B. Customer-managed encryption keys
- C. Customer-supplied encryption keys
- D. Google default encryption

Answer: B

Explanation:

To comply with GDPR requirements and manage encryption keys for workloads across multiple Google Cloud services, customer-managed encryption keys (CMEK) offer a suitable solution.

Customer-managed encryption keys (B):

CMEK allows you to create and manage encryption keys using Google Cloud Key Management Service (KMS). You maintain full control over the key lifecycle, including key rotation and destruction. CMEK can be used with various Google Cloud services, such as Compute Engine, Google Kubernetes Engine, Cloud Storage, BigQuery, and Pub/Sub, ensuring consistent

and compliant encryption across your environment.

Using CMEK, you can implement data protection by design, aligning with GDPR requirements by ensuring that encryption keys are appropriately managed and secured.

Reference:

Customer-Managed Encryption Keys Documentation

Encryption at Rest in Google Cloud

Question: 162

Which Identity-Aware Proxy role should you grant to an Identity and Access Management (IAM) user to access HTTPS resources?

- A. Security Reviewer
- B. IAP-Secured Tunnel User
- C. IAP-Secured Web App User
- D. Service Broker Operator

Answer: C

Explanation:

To grant an IAM user access to HTTPS resources protected by Identity-Aware Proxy (IAP), you should assign the IAP-Secured Web App User role.

IAP-Secured Web App User (C):

This role grants the necessary permissions for a user to access web applications secured by IAP. It includes permissions to access the IAP-secured resources, ensuring that users can authenticate and gain the appropriate access based on the policies defined in IAP.

Assigning this role ensures that users have the right level of access to protected web resources, aligned with the security controls enforced by IAP.

Reference:

Identity-Aware Proxy Documentation

IAP Roles and Permissions

Question: 163

You need to audit the network segmentation for your Google Cloud footprint. You currently operate Production and Non-Production infrastructure-as-a-service (IaaS) environments. All your VM instances are deployed without any service account customization.

After observing the traffic in your custom network, you notice that all instances can communicate freely – despite tag-based VPC firewall rules in place to segment traffic properly – with a priority of 1000. What are the most likely reasons for this behavior?

- A. All VM instances are missing the respective network tags.
- B. All VM instances are residing in the same network subnet.
- C. All VM instances are configured with the same network route.
- D. A VPC firewall rule is allowing traffic between source/targets based on the same service account with priority 999.
- E. A VPC firewall rule is allowing traffic between source/targets based on the same service account with priority 1001.

Answer: D

Explanation:

Firewall Rule Analysis: Analyze the existing VPC firewall rules to identify any rules that might allow traffic between VM instances based on the same service account.

Priority Check: Check the priority of these rules. A rule with a priority lower than 1000 (such as 999) will take precedence over your tag-based rules.

Service Account Configuration: Since your VM instances are deployed without any service account customization, they are likely using the default service account. A firewall rule allowing traffic between instances using this default service account will override the tag-based rules if it has a higher priority.

Testing and Validation: Disable or adjust the priority of the rule with priority 999 to test if the tagbased segmentation works correctly. Validate that the traffic is segmented according to your intended configuration. Reference::

Google Cloud - VPC Firewall Rules

Google Cloud - Service Accounts

Question: 164

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 1AM Owner role at the organization level. Create a VPC Service Controls

perimeter around the project that restricts the [compute.googleapis.com](https://cloud.google.com/compute/docs/api/) 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:

Organization Policy: Use the constraints/compute.skipDefaultNetworkCreation organization policy constraint to disable the creation of default networks in new projects.

Policy Application: Apply this constraint at the organization level to ensure it affects all projects within your organization, preventing the creation of default networks.

Best Practices Compliance: Following this best practice helps maintain a clean and secure network configuration by avoiding the use of default networks, which may not be properly segmented or secured.

Verification: Verify the policy application by creating new projects and ensuring that default networks are not created.

Reference::

Google Cloud - Organization Policy Constraints

Google Cloud - Best Practices for Enterprise Organizations

Question: 165

You are a security administrator at your company and are responsible for managing access controls (identification, authentication, and authorization) on Google Cloud. Which Google-recommended best practices should you follow when configuring authentication and authorization? (Choose two.)

- A. Use Google default encryption.
- B. Manually add users to Google Cloud.
- C. Provision users with basic roles using Google's Identity and Access Management (IAM) service.
- D. Use SSO/SAML integration with Cloud Identity for user authentication and user lifecycle management.
- E. Provide granular access with predefined roles.

Answer: DE

Explanation:

SSO/SAML Integration: Implement SSO (Single Sign-On) with SAML integration through Cloud Identity to streamline user authentication and lifecycle management. This ensures centralized management of user identities and access.

Predefined Roles: Use predefined roles to provide granular access control. These roles are designed to follow the principle of least privilege, ensuring that users have the minimum necessary permissions to perform their tasks.

User Management: By leveraging SSO/SAML, user provisioning and de-provisioning become more efficient and secure. This integration helps maintain consistent access policies across your organization.

Access Control: Predefined roles reduce the risk of over-permission by offering well-defined access levels, enhancing security and compliance. Reference:: Google Cloud - SSO with SAML

Google Cloud - IAM Best Practices

Question: 166

You have been tasked with inspecting IP packet data for invalid or malicious content. What should you do?

- A. Use Packet Mirroring to mirror traffic to and from particular VM instances. Perform inspection using security software that analyzes the mirrored traffic.
- B. Enable VPC Flow Logs for all subnets in the VPC. Perform inspection on the Flow Logs data using Cloud Logging.
- C. Configure the Fluentd agent on each VM Instance within the VPC. Perform inspection on the log data using Cloud Logging.
- D. Configure Google Cloud Armor access logs to perform inspection on the log data.

Answer: A

Explanation:

Packet Mirroring Setup: Configure Packet Mirroring in your Google Cloud VPC to capture traffic to and from specific VM instances. This allows you to analyze the traffic for security and compliance purposes.

Security Software: Use specialized security software to inspect the mirrored traffic. This software can detect invalid or malicious content in the IP packets.

Mirroring Configuration: Specify the instances, network, and traffic direction (ingress, egress, or both) to be mirrored. Ensure that the mirrored traffic is directed to an appropriate analysis destination.

Traffic Analysis: Continuously monitor and analyze the mirrored traffic for any signs of malicious activity or anomalies. Use the findings to enhance your security posture and respond to potential threats. **Reference::**


Google Cloud - Packet Mirroring

Google Cloud - Packet Mirroring Best Practices

Question: 167

You have the following resource hierarchy. There is an organization policy at each node in the

hierarchy as shown. Which load balancer types are denied in VPC A?



- A. All load balancer types are denied in accordance with the global node's policy.
- B. INTERNAL_TCP_UDP, INTERNAL_HTTP_HTTPS is denied in accordance with the folder's policy.
- C. EXTERNAL_TCP_PROXY, EXTERNAL_SSL_PROXY are denied in accordance with the project's policy.
- D. EXTERNAL_TCP_PROXY, EXTERNAL_SSL_PROXY, INTERNAL_TCP_UDP, and INTERNAL_HTTP_HTTPS are denied in accordance with the folder and project's policies.

Answer: D

Explanation:

Understanding Organization Policies:

Organization policies are rules that can be set at different levels of the resource hierarchy in GCP to enforce governance and compliance.

These policies can be set at the organization node, folders, and projects, and they are inherited down the hierarchy unless explicitly overridden.

Hierarchy and Policy Inheritance:

The provided resource hierarchy has an organization node (Example.com), folders (Folder 1 and

Folder 2), and a project (Project 2) under Folder 2 with a specific VPC (VPC A).

Each node in the hierarchy can have its own policies, and these policies are inherited by child nodes unless overridden.

Analyzing the Policies in the Hierarchy:

Organization Node Policy:

json

Copy code

```
{ "constraint": "constraints/compute.restrictLoadBalancerCreationForTypes", "listPolicy": { "allValues": "DENY" } }
```

This policy at the organization node denies all load balancer types.

Folder 2 Policy:

json

Copy code

```
{ "constraint": "constraints/compute.restrictLoadBalancerCreationForTypes", "listPolicy": { "deniedValues":  
  ["INTERNAL_TCP_UDP", "INTERNAL_HTTP_HTTPS"] } }
```

This policy at Folder 2 denies the creation of INTERNAL_TCP_UDP and INTERNAL_HTTP_HTTPS load balancers.

Project 2 Policy:

json

Copy code

```
{ "constraint": "constraints/compute.restrictLoadBalancerCreationForTypes", "listPolicy": { "deniedValues":  
  ["EXTERNAL_TCP_PROXY", "EXTERNAL_SSL_PROXY"] } }
```

This policy at Project 2 denies the creation of EXTERNAL_TCP_PROXY and EXTERNAL_SSL_PROXY load balancers.

Policy Application to VPC A:

Since policies are inherited, VPC A (which is within Project 2 under Folder 2) will be affected by the policies of both Folder 2 and Project 2.

Combining the denied values from both Folder 2 and Project 2:

From Folder 2: INTERNAL_TCP_UDP, INTERNAL_HTTP_HTTPS

From Project 2: EXTERNAL_TCP_PROXY, EXTERNAL_SSL_PROXY

Conclusion:

VPC A will have the following load balancer types denied: INTERNAL_TCP_UDP, INTERNAL_HTTP_HTTPS, EXTERNAL_TCP_PROXY, EXTERNAL_SSL_PROXY.

Reference::

GCP Documentation on Organization Policies

GCP Documentation on Constraints and List Policies

Question: 168

Your security team wants to implement a defense-in-depth approach to protect sensitive data stored in a Cloud Storage bucket. Your team has the following requirements:

The Cloud Storage bucket in Project A can only be readable from Project B.

The Cloud Storage bucket in Project A cannot be accessed from outside the network.

Data in the Cloud Storage bucket cannot be copied to an external Cloud Storage bucket.

What should the security team do?

A. Enable domain restricted sharing in an organization policy, and enable uniform bucket-level access on the Cloud Storage bucket.

B. Enable VPC Service Controls, create a perimeter around Projects A and B. and include the Cloud Storage API in the Service Perimeter configuration.

- C. Enable Private Access in both Project A and B's networks with strict firewall rules that allow communication between the networks.
- D. Enable VPC Peering between Project A and B's networks with strict firewall rules that allow communication between the networks.

Answer: B

Explanation:

VPC Peering is between organizations not between Projects in an organization. That is Shared VPC. In this case, both projects are in same organization so having VPC Service Controls around both projects with necessary rules should be fine.

<https://cloud.google.com/vpc-service-controls/docs/overview>

Question: 169

You need to create a VPC that enables your security team to control network resources such as firewall rules. How should you configure the network to allow for separation of duties for network resources?

- A. Set up multiple VPC networks, and set up multi-NIC virtual appliances to connect the networks.
- B. Set up VPC Network Peering, and allow developers to peer their network with a Shared VPC.
- C. Set up a VPC in a project. Assign the Compute Network Admin role to the security team, and assign the Compute Admin role to the developers.
- D. Set up a Shared VPC where the security team manages the firewall rules, and share the network with developers via service projects.

Answer: D

Explanation:

Setting up a Shared VPC allows you to create a centrally managed network that spans multiple projects. Here's how you can achieve this while ensuring separation of duties:

Create a Host Project:

Create a project that will act as the host project for your Shared VPC.

Configure Shared VPC:

In the host project, enable the Shared VPC feature.

Create Service Projects:

Create separate service projects for different teams, such as developers and other stakeholders.

Assign Roles:

Security Team: Grant the Compute Network Admin role. This allows the security team to manage network resources, such as firewall rules, subnets, and routes.

Developers: Share the host project's network with the service projects. Assign roles like Compute Instance Admin to developers in the service projects, enabling them to create and manage VM instances without altering network configurations.

Firewall Management:

The security team will define and manage firewall rules within the host project, ensuring consistent and secure network policies.

Benefits:

Separation of Duties: Security teams handle networking, and developers focus on application deployment and management.

Centralized Control: Network policies are centrally managed, ensuring compliance and security.

Scalability: Easy to add new projects and teams without compromising the overall network security. **Reference:** Google Cloud VPC Documentation

Managing Resources with Shared VPC

Question: 170

You are onboarding new users into Cloud Identity and discover that some users have created consumer user accounts using the corporate domain name. How should you manage these consumer user accounts with Cloud Identity?

- A. Use Google Cloud Directory Sync to convert the unmanaged user accounts.
- B. Create a new managed user account for each consumer user account.
- C. Use the transfer tool for unmanaged user accounts.
- D. Configure single sign-on using a customer's third-party provider.

Answer: C

Explanation:

To manage consumer user accounts created using the corporate domain name, you can use the transfer tool for unmanaged user accounts provided by Google Cloud Identity. Here's how you can proceed:

Identify Unmanaged Accounts:

Use the Cloud Identity interface to identify consumer (unmanaged) accounts that exist with your corporate domain.

Initiate Transfer Process:

Use the transfer tool for unmanaged user accounts to initiate the transfer. This tool helps in converting unmanaged accounts (consumer accounts) into managed accounts.

User Notification:

Users with unmanaged accounts will receive an email notification prompting them to accept the transfer to the organization's managed account system.

Accept Transfer:

Users need to follow the instructions in the email to accept the transfer. Once accepted, their accounts will be managed under your organization's Cloud Identity setup.

Benefits:

Centralized Management: All user accounts under your corporate domain are managed centrally, ensuring compliance and security.

Enhanced Security: Managed accounts provide better control over security policies and access management.

Reference:

Transfer tool for unmanaged users

Question: 171

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: BD

Explanation:

<https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints-constraints/compute.trustedImageProjects>

This list constraint defines the set of projects that can be used for image storage and disk instantiation for Compute Engine. If this constraint is active, only images from trusted projects will be allowed as the source for boot disks for new instances.

Question: 172

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 (IAM) role to the DevOps team.
- B. Create a custom IAM role with limited list/view permissions, and assign it to the DevOps team.
- C. Create a service account, and grant it the Project Owner IAM 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.

Answer: D

Explanation:

To ensure least-privilege access and provide necessary permissions to the DevOps team only during a deployment issue, follow these steps:

Create a Service Account:

In your Google Cloud project, create a new service account specifically for the DevOps team.

Assign Limited Permissions:

Grant the service account permissions with only the necessary list/view roles. For instance, you can create a custom IAM role with `compute.instances.list` and `compute.instances.get` permissions.

Grant Service Account User Role:

Assign the Service Account User role to the DevOps team members for the created service account.

This allows them to act as the service account and use its permissions.

Access Control During Incidents:

During a deployment issue, the DevOps team can temporarily use the service account to access the resources. This ensures they have the least-privilege access required to investigate and resolve the issue.

Automation and Monitoring:

Implement automation to enable and disable the service account access as needed and monitor the usage to ensure compliance with the least-privilege principle.

Benefits:

Security: Limits access to only what is necessary, reducing the risk of unauthorized changes.

Flexibility: Provides necessary access during incidents without granting permanent elevated permissions.

Reference:

Creating and Managing Service Accounts

Service Account User Role

Question: 173

You are working with a client who plans to migrate their data to Google Cloud. You are responsible for recommending an encryption service to manage their encrypted keys. You have the following requirements:

The master key must be rotated at least once every 45 days.

The solution that stores the master key must be FIPS 140-2 Level 3 validated.

The master key must be stored in multiple regions within the US for redundancy.

Which solution meets these requirements?

- A. Customer-managed encryption keys with Cloud Key Management Service
- B. Customer-managed encryption keys with Cloud HSM
- C. Customer-supplied encryption keys
- D. Google-managed encryption keys

Answer: B

Explanation:

To meet the requirements of rotating the master key every 45 days, achieving FIPS 140-2 Level 3 validation, and ensuring the master key is stored redundantly in multiple US regions, you should use Customer-managed encryption keys with Cloud HSM. Here's how:

Set Up Cloud HSM:

Deploy Cloud HSM in your Google Cloud environment. Cloud HSM provides a hardware-based key management solution that meets FIPS 140-2 Level 3 compliance.

Create and Manage Keys:

Create your encryption keys in Cloud HSM. These keys can be managed and rotated per your policy requirements.

Key Rotation:

Set up a key rotation schedule to rotate the master key every 45 days. Cloud HSM allows you to automate this process.

Geographic Redundancy:

Ensure that your Cloud HSM configuration spans multiple regions within the US to achieve redundancy. This will ensure that your keys are available even if a particular region experiences an outage.

Compliance:

Cloud HSM's FIPS 140-2 Level 3 validation ensures that your encryption keys are managed in a secure and compliant manner.

Benefits:

Security and Compliance: Meets stringent compliance requirements.

Automated Management: Simplifies key management and rotation.

Redundancy: Ensures high availability of keys across multiple regions.

Reference:

Cloud HSM Documentation

Key Management with Cloud KMS and Cloud HSM

Question: 174

You manage your organization's Security Operations Center (SOC). You currently monitor and detect network traffic anomalies in your VPCs based on network logs. However, you want to explore your environment using network payloads and headers. Which Google Cloud product should you use?

- A. Cloud IDS
- B. VPC Service Controls logs
- C. VPC Flow Logs
- D. Google Cloud Armor
- E. Packet Mirroring

Answer: E

Explanation:

<https://cloud.google.com/vpc/docs/packet-mirroring>

Packet Mirroring clones the traffic of specified instances in your Virtual Private Cloud (VPC) network and forwards it for examination. Packet Mirroring captures all traffic and packet data, including payloads and headers.

Question: 175

You are consulting with a client that requires end-to-end encryption of application data (including data in transit, data in use, and data at rest) within Google Cloud. Which options should you utilize to accomplish this? (Choose two.)

- A. External Key Manager
- B. Customer-supplied encryption keys
- C. Hardware Security Module
- D. Confidential Computing and Istio
- E. Client-side encryption

Answer: DE

Explanation:

Google Cloud customers with additional requirements for encryption of data over WAN can choose to implement further protections for data as it moves from a user to an application, or virtual machine to virtual machine. These protections include IPSec tunnels, Gmail S/MIME, managed SSL certificates, and Istio.

<https://cloud.google.com/docs/security/encryption-in-transit>

Question: 176

You need to enforce a security policy in your Google Cloud organization that prevents users from exposing objects in their buckets externally. There are currently no buckets in your organization. Which solution should you implement proactively to achieve this goal with the least operational overhead?

- A. Create an hourly cron job to run a Cloud Function that finds public buckets and makes them private.
- B. Enable the constraints/storage.publicAccessPrevention constraint at the organization level.
- C. Enable the constraints/storage.uniformBucketLevelAccess constraint at the organization level.
- D. Create a VPC Service Controls perimeter that protects the storage.googleapis.com service in your projects that contains buckets. Add any new project that contains a bucket to the perimeter.

Answer: B

Explanation:

<https://cloud.google.com/storage/docs/public-access-prevention>

Public access prevention protects Cloud Storage buckets and objects from being accidentally exposed to the public. If your bucket is contained within an organization, you can enforce public access prevention by using the organization policy constraint storage.publicAccessPrevention at the project, folder, or organization level.

Question: 177

Your company requires the security and network engineering teams to identify all network anomalies and be able to capture payloads within VPCs. Which method should you use?

- A. Define an organization policy constraint.
- B. Configure packet mirroring policies.
- C. Enable VPC Flow Logs on the subnet.
- D. Monitor and analyze Cloud Audit Logs.

Answer: B

Explanation:

<https://cloud.google.com/vpc/docs/packet-mirroring>

Packet Mirroring clones the traffic of specified instances in your Virtual Private Cloud (VPC) network and forwards it for examination. Packet Mirroring captures all traffic and packet data, including payloads and headers.

Question: 178

Your company is moving to Google Cloud. You plan to sync your users first by using Google Cloud

Directory Sync (GCDS). Some employees have already created Google Cloud accounts by using their company email addresses that were created outside of GCDS. You must create your users on Cloud Identity.

What should you do?

- A. Configure GCDS and use GCDS search rules to sync these users.
- B. Use the transfer tool to migrate unmanaged users.
- C. Write a custom script to identify existing Google Cloud users and call the Admin SDK Directory API to transfer their account.
- D. Configure GCDS and use GCDS exclusion rules to ensure users are not suspended.

Answer: B

Explanation:

Understanding Unmanaged Users:

Unmanaged users are those who have created Google Cloud accounts using their company email addresses outside of the organization's management (e.g., through GCDS).

Challenge:

The goal is to bring these unmanaged accounts under your company's control without disrupting their existing accounts and access.

Using the Transfer Tool:

Google provides a transfer tool specifically designed to migrate unmanaged users to a managed state.

This tool allows administrators to invite unmanaged users to join the organization's Google Cloud Identity or Google

Workspace account.

Steps to Use the Transfer Tool:

Step 1: Access the transfer tool from the Google Admin console.

Step 2: Identify the unmanaged users using their email addresses.

Step 3: Send invitations to these users to transfer their accounts.

Step 4: Users accept the invitations, allowing their accounts to be managed under the organization's domain.

Benefits:

This method ensures a smooth transition for users without losing access to their existing data and services.

It aligns with best practices for managing user accounts in a corporate environment.

Reference::

Migrate unmanaged users

Google Cloud Directory Sync

Question: 179

Your organization is using GitHub Actions as a continuous integration and delivery (CI/CD) platform. You must enable access to Google Cloud resources from the CI/CD pipelines in the most secure way. What should you do?

- A. Create a service account key and add it to the GitHub pipeline configuration file.
- B. Create a service account key and add it to the GitHub repository content.
- C. Configure a Google Kubernetes Engine cluster that uses Workload Identity to supply credentials to GitHub.
- D. Configure workload identity federation to use GitHub as an identity pool provider.

Answer: D

Explanation:

Challenge:

Ensuring secure access to Google Cloud resources from GitHub Actions CI/CD pipelines without directly managing service account keys.

Workload Identity Federation:

Allows for the delegation of access to Google Cloud resources based on federated identities, such as those from GitHub.

Benefits:

This approach eliminates the need to manage service account keys, reducing the risk of key leakage.

It leverages GitHub's identity provider capabilities to authenticate and authorize access.

Steps to Configure Workload Identity Federation:

Step 1: Create a workload identity pool in Google Cloud.

Step 2: Add GitHub as an identity provider within the pool.

Step 3: Configure the necessary permissions and bindings for the identity pool to allow GitHub

Actions to access Google Cloud resources.

Step 4: Update the GitHub Actions workflow to use the identity federation for authentication.

Reference::

Workload Identity Federation

Configuring Workload Identity Federation with GitHub

Question: 180

Your company must follow industry specific regulations. Therefore, you need to enforce customer managed encryption keys (CMEK) for all new Cloud Storage resources in the organization called org1. What command should you execute?

- A. • organization policy: constraints/gcp.restrictStorageNonCmekServices
 - binding at: org1
 - policy type: deny
 - policy value: storage.googleapis.com
- B. • organization policy: constraints/gcp.restrictNonCmekServices
 - binding at: org1
 - policy type: deny
 - policy value: storage.googleapis.com
- C. • organization policy: constraints/gcp.restrictStorageNonCmekServices
 - binding at: org1
 - policy type: allow
 - policy value: all supported services
- D. • organization policy: constraints/gcp.restrictNonCmekServices
 - binding at: org1
 - policy type: allow
 - policy value: storage.googleapis.com

Answer: D

Explanation:

Requirement:

Enforce the use of Customer-Managed Encryption Keys (CMEK) for all new Cloud Storage resources in the organization.

Policy Constraint:

Use the constraints/gcp.restrictNonCmekServices constraint to enforce CMEK usage.

Policy Type and Value:

Set the policy type to allow to specify which services must use CMEK.

In this case, the policy value should be storage.googleapis.com to target Cloud Storage.

Command:

Applying the organization policy with the appropriate binding ensures that all new Cloud Storage resources under the organization will require CMEK.

Steps:

Step 1: Go to the Google Cloud Console.

Step 2: Navigate to the Organization Policies page.

Step 3: Apply the policy constraint constraints/gcp.restrictNonCmekServices with the allow policy

type and storage.googleapis.com as the policy value.

Reference::

Organization Policy Constraints

Customer-Managed Encryption Keys (CMEK)

Question: 181

Your Google Cloud organization allows for administrative capabilities to be distributed to each team through provision of a

Google Cloud project with Owner role (roles/owner). The organization contains thousands of Google Cloud Projects. Security Command Center Premium has surfaced multiple open_mysql_port findings. You are enforcing the guardrails and need to prevent these types of common misconfigurations.

What should you do?

- A. Create a firewall rule for each virtual private cloud (VPC) to deny traffic from 0.0.0.0/0 with priority 0.
- B. Create a hierarchical firewall policy configured at the organization to deny all connections from 0.0.0.0/0.
- C. Create a Google Cloud Armor security policy to deny traffic from 0.0.0.0/0.
- D. Create a hierarchical firewall policy configured at the organization to allow connections only from internal IP ranges.

Answer: D

Explanation:

Challenge:

Prevent common misconfigurations that expose services (e.g., MySQL) to the public internet.

Hierarchical Firewall Policies:

These policies can be applied at the organization level to enforce consistent network security rules across all projects.

Solution:

Create a hierarchical firewall policy that allows connections only from internal IP ranges.

This policy ensures that services like MySQL are not exposed to 0.0.0.0/0 (the entire internet).

Steps:

Step 1: Define the hierarchical firewall policy at the organization level.

Step 2: Set the rule to allow traffic only from internal IP ranges.

Step 3: Apply the policy to all projects under the organization.

Benefits:

Centralized management of network security.

Prevents accidental exposure of services to the public internet, enhancing security.

Reference::

Hierarchical Firewall Policies

Securing MySQL on GCP

Question: 182

Your organization must comply with the regulation to keep instance logging data within Europe. Your workloads will be hosted in the Netherlands in region europe-west4 in a new project. You must configure Cloud Logging to keep your data in the country.

What should you do?

- A. Configure the organization policy constraint gcp.resourceLocations to europe-west4.
- B. Set the logging storage region to europe-west4 by using the gcloud CLI logging settings update.
- C. Create a new log bucket in europe-west4, and redirect the default bucket to the new bucket.
- D. Configure log sink to export all logs into a Cloud Storage bucket in europe-west4.

Answer: D

Explanation:

To comply with the regulation to keep instance logging data within Europe, specifically in the Netherlands (region europe-west4), you need to configure Cloud Logging to ensure that all logs are stored in a specified region. Here are the steps to achieve this:

Create a Cloud Storage Bucket:

Go to the Google Cloud Console.

Navigate to Cloud Storage and create a new bucket.

Set the location type to "Region" and select "europe-west4" as the region.

Configure Log Sink:

Go to the Logging section in the Google Cloud Console.

Create a new log sink and configure it to export logs to the Cloud Storage bucket you just created.

Ensure the sink includes all logs you want to keep within europe-west4.

Verify Configuration:

Ensure that the logs are being exported to the Cloud Storage bucket by checking the contents of the bucket.

This setup ensures that all your logging data remains within the specified region, adhering to compliance requirements.

Reference:

Creating and Managing Log Sinks

Cloud Storage Locations

Question: 183

Your organization is rolling out a new continuous integration and delivery (CI/CD) process to deploy infrastructure and applications in Google Cloud. Many teams will use their own instances of the CI/CD workflow. It will run on Google Kubernetes Engine (GKE). The CI/CD pipelines must be designed to securely access Google Cloud APIs.

What should you do?

- A. • 1 Create a dedicated service account for the CI/CD pipelines
- 2 Run the deployment pipelines in a dedicated nodes pool in the GKE cluster
- 3 Use the service account that you created as identity for the nodes in the pool to authenticate to the Google

Cloud APIs

- B. • 1 Create service accounts for each deployment pipeline
- 2 Generate private keys for the service accounts
- 3 Securely store the private keys as Kubernetes secrets accessible only by the pods that run the

specific deploy pipeline

- C. * 1 Create individual service accounts (or each deployment pipeline
- 2 Add an identifier for the pipeline in the service account naming convention
- 3 Ensure each pipeline runs on dedicated pods
- 4 Use workload identity to map a deployment pipeline pod with a service account

- D. • 1 Create two service accounts one for the infrastructure and one for the application

deployment

- 2 Use workload identities to let the pods run the two pipelines and authenticate with the

service accounts

- 3 Run the infrastructure and application pipelines in separate namespaces

Answer: C

Explanation:

To securely access Google Cloud APIs from CI/CD pipelines running on Google Kubernetes Engine (GKE), follow these

steps:

Create Service Accounts:

Create individual service accounts for each CI/CD pipeline. This ensures isolation and minimal permissions per pipeline.

Use a naming convention that includes an identifier for each pipeline, such as pipeline-a-sa, pipeline-b-sa, etc.

Configure Kubernetes Service Accounts:

Create Kubernetes service accounts for each CI/CD pipeline pod.

Map Kubernetes Service Accounts to Google Service Accounts:

Use Workload Identity to associate Kubernetes service accounts with the corresponding Google service accounts. This allows the pods to authenticate to Google Cloud APIs securely.

Example command to bind the Kubernetes service account to the Google service account: `gcloud iam service-accounts add-iam-policy-binding \ --role roles/iam.workloadIdentityUser \ -member`

```
"serviceAccount:<PROJECT_ID>.svc.id.goog[<NAMESPACE>/<KSA_NAME>]" \
<GSA_NAME>@<PROJECT_ID>.iam.gserviceaccount.com
```

Deploy CI/CD Pipelines:

Ensure each pipeline runs in dedicated pods that use the specific Kubernetes service accounts configured earlier.

This setup ensures that each pipeline has the necessary permissions to interact with Google Cloud APIs securely, adhering to the principle of least privilege.

Reference:

Using Workload Identity

Managing Service Accounts

Question: 184

Your organization processes sensitive health information. You want to ensure that data is encrypted while in use by the virtual machines (VMs). You must create a policy that is enforced across the entire organization.

What should you do?

- A. Implement an organization policy that ensures that all VM resources created across your organization use customer-managed encryption keys (CMEK) protection.
- B. Implement an organization policy that ensures all VM resources created across your organization are Confidential VM instances.
- C. Implement an organization policy that ensures that all VM resources created across your organization use Cloud External Key Manager (EKM) protection.
- D. No action is necessary because Google encrypts data while it is in use by default.

Answer: B

Explanation:

To ensure that data is encrypted while in use by the virtual machines (VMs) and enforce this policy across your organization, you should use Confidential VM instances. Here are the steps: Enable Confidential VM:

Ensure that Confidential VMs are available in your selected regions and enabled for your project.

Set Organization Policy:

Implement an organization policy to enforce the use of Confidential VM instances for all VMs across your organization.

Use the Google Cloud Console or the `gcloud` command-line tool to set this policy. Example

command:

```
gcloud resource-manager org-policies set-policy my_policy.yaml
```

Example my_policy.yaml:

```
name: organizations/1234567890/policies/compute.requireConfidentialCompute spec: rules: - enforce: true
```

Verify and Monitor:

Ensure that all newly created VMs across your organization are Confidential VMs.

Regularly monitor compliance through the Google Cloud Console and set up alerts if non-compliant VMs are created.

Benefits:

Data Encryption in Use: Confidential VMs ensure that data is encrypted not just at rest and in transit but also while in use.

Policy Enforcement: Organization policies provide a way to enforce security configurations across all projects under your organization.

Reference:

Confidential Computing Documentation

Creating and Managing Organization Policies

Question: 185

You are a Cloud Identity administrator for your organization. In your Google Cloud environment groups are used to manage user permissions. Each application team has a dedicated group Your team is responsible for creating these groups and the application teams can manage the team members on their own through the Google Cloud console. You must ensure that the application teams can only add users from within your organization to their groups.

What should you do?

- A. Change the configuration of the relevant groups in the Google Workspace Admin console to prevent external users from being added to the group.
- B. Set an Identity and Access Management (IAM) policy that includes a condition that restricts group membership to user principals that belong to your organization.
- C. Define an Identity and Access Management (IAM) deny policy that denies the assignment of principals that are outside your organization to the groups in scope.
- D. Export the Cloud Identity logs to BigQuery Configure an alert for external members added to groups Have the alert trigger a Cloud Function instance that removes the external members from the group.

Answer: A

Explanation:

To ensure that application teams can only add users from within your organization to their groups, you need to configure the group settings in the Google Workspace Admin console. Here are the steps:

Access Google Workspace Admin Console:

Go to the Google Workspace Admin console.

Navigate to the Groups section.

Configure Group Settings:

Select the relevant groups used by the application teams.

Go to the group settings and set the group to only allow members from your domain.

This prevents external users from being added to these groups.

Apply and Save Changes:

Apply the changes to ensure that only users from within your organization can be added to these groups.

Verify Configuration:

Test the configuration by attempting to add an external user to the group and verifying that it is not allowed.

Benefits:

Security: Prevents unauthorized access by ensuring that only internal users can be added to groups.

Compliance: Helps in adhering to organizational policies regarding user access and management.

Reference:

Google Workspace Admin Help: Groups

Question: 186

Your organization wants to be continuously evaluated against CIS Google Cloud Computing Foundations Benchmark v1 3 0 (CIS Google Cloud Foundation 1 3). Some of the controls are irrelevant to your organization and must be disregarded in evaluation. You need to create an automated system or process to ensure that only the relevant controls are evaluated. What should you do?

- A. Mark all security findings that are irrelevant with a tag and a value that indicates a security exception Select all marked findings and mute them on the console every time they appear Activate Security Command Center (SCC) Premium.
- B. Activate Security Command Center (SCC) Premium Create a rule to mute the security findings in SCC so they are not evaluated.
- C. Download all findings from Security Command Center (SCC) to a CSV file Mark the findings that are part of CIS Google Cloud Foundation 1 3 in the file Ignore the entries that are irrelevant and out of scope for the company.
- D. Ask an external audit company to provide independent reports including needed CIS benchmarks. In the scope of the audit clarify that some of the controls are not needed and must be disregarded.

Answer: B

Explanation:

Activate Security Command Center (SCC) Premium: Security Command Center (SCC) Premium provides advanced security analytics and best practice recommendations for your Google Cloud environment. It includes functionalities such as asset discovery, vulnerability scanning, and security findings.

Create a Custom Rule to Mute Irrelevant Security Findings:

Navigate to the Security Command Center (SCC) in the Google Cloud Console.

Go to the "Settings" tab and find the "Mute findings" section.

Create a new mute rule by specifying the conditions that match the irrelevant controls you want to disregard. These conditions can be based on attributes such as resource type, finding type, and other metadata.

Apply this mute rule, which will ensure that the specified findings are not evaluated in your security posture assessments.

Ensure Continuous Compliance Monitoring:

The mute rules will automatically filter out the irrelevant findings, ensuring that only relevant controls from the CIS Google Cloud Computing Foundations Benchmark v1.3.0 are evaluated. Regularly review and update the mute rules to adapt to any changes in your compliance requirements or security posture.

Reference::

Security Command Center Documentation

Creating and Managing Mute Rules

Question: 187

You are routing all your internet facing traffic from Google Cloud through your on-premises internet connection. You want to accomplish this goal securely and with the highest bandwidth possible.

What should you do?

- A. Create an HA VPN connection to Google Cloud Replace the default 0 0 0 0/0 route.
- B. Create a routing VM in Compute Engine Configure the default route with the VM as the next hop.
- C. Configure Cloud Interconnect with HA VPN Replace the default 0 0 0 0/0 route to an on-premises destination.
- D. Configure Cloud Interconnect and route traffic through an on-premises firewall.

Answer: C

Explanation:

Configure Cloud Interconnect:

Cloud Interconnect provides high-bandwidth, low-latency connectivity between your on-premises network and Google Cloud. You can choose between Dedicated Interconnect or Partner Interconnect depending on your requirements.

Set up the physical connection and establish the interconnect through the Google Cloud Console or by working with a partner.

Set Up HA VPN:

High Availability (HA) VPN provides a robust, reliable VPN connection to Google Cloud with SLA guarantees. This is crucial for ensuring secure and high-bandwidth connectivity.

Configure two VPN tunnels to ensure redundancy and failover capabilities.

Update Default Route:

Replace the default 0.0.0.0/0 route in your Google Cloud VPC to direct traffic to your on-premises network via the Cloud Interconnect and HA VPN setup.

This ensures all internet-facing traffic is securely routed through your on-premises internet connection.

Ensure Proper Security and Routing Policies:

Implement appropriate firewall rules and security policies on both Google Cloud and on-premises environments to control traffic and ensure secure communication.

Monitor the traffic and connectivity status to maintain optimal performance and security. Reference::

Cloud Interconnect Documentation

HA VPN Documentation

Routing Traffic

Question: 188

Your organization is transitioning to Google Cloud You want to ensure that only trusted container images are deployed on Google Kubernetes Engine (GKE) clusters in a project. The containers must be deployed from a centrally managed. Container Registry and signed by a trusted authority.

What should you do?

Choose 2 answers

- A. Configure the Binary Authorization policy with respective attestations for the project.
- B. Create a custom organization policy constraint to enforce Binary Authorization for Google Kubernetes Engine (GKE).
- C. Enable Container Threat Detection in the Security Command Center (SCC) for the project.

- D. Configure the trusted image organization policy constraint for the project.
- E. Enable Pod Security standards and set them to Restricted.

Answer: A, B

Explanation:

Configure Binary Authorization:

Binary Authorization is a deploy-time security control that ensures only trusted container images are deployed on GKE. It uses attestations to verify the authenticity and integrity of the images.

Enable Binary Authorization in your project through the Google Cloud Console or using the `gcloud` command-line tool. Define attestation policies that specify which attestors (trusted entities) must sign off on container images before deployment.

Set Up Attestors:

Create and configure attestors that will sign the container images. This involves generating cryptographic keys and setting up trusted authorities.

Attestors can be configured to sign images based on criteria such as vulnerability scanning results, compliance checks, and other security policies.

Create a Custom Organization Policy Constraint:

Define an organization policy constraint that enforces Binary Authorization across your GKE clusters. This custom constraint ensures that all clusters in the organization must adhere to the Binary Authorization policy, preventing the deployment of unsigned or unauthorized container images.

Implement and Enforce the Policies:

Apply the Binary Authorization policy and the organization policy constraint to your GKE clusters.

Regularly review and update the policies and attestation rules to align with your security and compliance requirements.

Reference::

Binary Authorization Documentation

Creating Attestors

Organization Policy Constraints

Question: 189

Your organization uses Google Workspace Enterprise Edition for authentication. You are concerned about employees leaving their laptops unattended for extended periods of time after authenticating into Google Cloud. You must prevent malicious people from using an employee's unattended laptop to modify their environment.

What should you do?

- A. Create a policy that requires employees to not leave their sessions open for long durations.
- B. Review and disable unnecessary Google Cloud APIs.
- C. Require strong passwords and 2SV through a security token or Google authenticator.
- D. Set the session length timeout for Google Cloud services to a shorter duration.

Answer: D

Explanation:

Access Google Cloud Console:

Log in to the Google Cloud Console with administrative privileges.

Navigate to the "IAM & Admin" section.

Set Session Length Timeout:

Go to the "Settings" page within IAM & Admin.

Locate the "Session control" settings.

Configure the session length timeout to a shorter duration, such as 15 or 30 minutes. This ensures that user sessions expire automatically after the specified time of inactivity.

Apply and Enforce the Policy:

Save the changes and ensure the new session timeout policy is applied across all users and services. Communicate the new policy to employees, highlighting the importance of session security and the rationale behind the change.

Additional Security Measures:

Consider implementing additional measures such as automatic screen locks and secure session management practices.

Educate employees on the importance of logging out of their sessions and securing their devices when not in use.

Reference::

Google Cloud IAM Documentation

Session Management Best Practices

Question: 190

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: A E

Explanation:

Customer-Managed Encryption Keys (CMEK):

CMEK allows you to manage encryption keys using Cloud Key Management Service (KMS). This gives you control over the lifecycle of the keys, including rotation, destruction, and auditing.

Set up a Cloud KMS key ring and create encryption keys that will be used to protect your data in BigQuery, Cloud SQL, and Cloud Storage.

Configure the services to use CMEK for encrypting data at rest, ensuring compliance with your organization's security policies.

Cloud External Key Manager (EKM):

Cloud EKM allows you to use keys managed by an external key management provider to encrypt data in Google Cloud services.

Integrate your external key management system with Google Cloud using supported protocols and APIs.

Configure your data warehouse services to use the external keys for encryption, ensuring that key management is handled outside of the Google Cloud environment.

Key Access Justifications:

Enable Key Access Justifications to provide visibility into why encryption keys are being accessed.

This helps in monitoring and auditing key usage to ensure compliance and security.

Set up policies and logging to capture and review key access requests, providing insights into how and why keys are used.

Access Transparency and Approval:

Implement Access Transparency to gain visibility into Google's access to your data and encryption keys.

Configure Access Approval to require explicit approval for Google support or engineering access to your data, adding an additional layer of security and control.

Reference::

Customer-Managed Encryption Keys (CMEK)

Cloud External Key Manager (EKM)

Key Access Justifications

Access Transparency

Access Approval

Question: 191

You are auditing all your Google Cloud resources in the production project. You want to identify all principals who can change firewall rules.

What should you do?

- A. Use Policy Analyzer to query the permissions compute, firewalls, create of compute, firewalls. Create of compute, firewalls.delete.
- B. Reference: the Security Health Analytics - Firewall Vulnerability Findings in the Security Command Center.
- C. Use Policy Analyzer to query the permissions compute, firewalls, get of compute, firewalls, list.
- D. Use Firewall Insights to understand your firewall rules usage patterns.

Answer: A

Explanation:

To identify all principals who can change firewall rules, you need to determine which users or service accounts have permissions that allow them to modify firewall rules in your Google Cloud project. The correct permissions to check for this are compute.firewalls.create and compute.firewalls.delete.

These permissions enable a user to create and delete firewall rules, respectively.

The Policy Analyzer tool in Google Cloud allows you to query and analyze IAM policies to identify which principals have specific permissions. By using Policy Analyzer, you can effectively identify all principals with the compute.firewalls.create and compute.firewalls.delete permissions.

Open Policy Analyzer: Go to the Google Cloud Console, navigate to IAM & Admin, and select Policy Analyzer.

Set Up Query: Create a new query specifying the permissions compute.firewalls.create and compute.firewalls.delete.

Run Query: Execute the query to retrieve a list of principals who have these permissions.

Review Results: Analyze the results to identify all users and service accounts with the capability to modify firewall rules. This method ensures you have a comprehensive list of all principals who can change firewall rules, enhancing your audit and security posture.

Reference::

Google Cloud Policy Analyzer Documentation

Google Cloud IAM Documentation

Question: 192

You manage one of your organization's Google Cloud projects (Project A). AVPC Service Control (SC) perimeter is blocking API access requests to this project including Pub/Sub. A resource running under a service account in another project (Project B) needs to collect messages from a Pub/Sub topic in your project Project B is not included in a VPC SC perimeter. You need to provide access from Project B to the Pub/Sub topic in Project A using the principle of least Privilege. What should you do?

- A. Configure an ingress policy for the perimeter in Project A and allow access for the service account in Project B to collect messages.
- B. Create an access level that allows a developer in Project B to subscribe to the Pub/Sub topic that is located in Project A.
- C. Create a perimeter bridge between Project A and Project B to allow the required communication between both projects.
- D. Remove the Pub/Sub API from the list of restricted services in the perimeter configuration for Project A.

Answer: A

Explanation:

When dealing with VPC Service Controls (VPC SC), it's important to ensure that only authorized resources can access sensitive data and services. To allow a resource in Project B to access Pub/Sub in Project A without compromising security, you should configure an ingress policy for the service perimeter in Project A.

Identify the Service Account: Determine the service account in Project B that requires access to the Pub/Sub topic in Project A.

Configure Ingress Policy:

Go to the Google Cloud Console.

Navigate to Security > VPC Service Controls.

Select the service perimeter for Project A.

Add an ingress rule specifying the service account from Project B and allowing it access to the necessary Pub/Sub resources.

Define Conditions: Ensure that the ingress policy adheres to the principle of least privilege, granting only the necessary permissions to collect messages from the Pub/Sub topic.

Save and Apply: Save the policy and apply the changes to enforce the new access controls.

This approach maintains the security boundaries set by VPC SC while enabling the required access from Project B to Project A.

Reference::

VPC Service Controls Documentation

Configuring Ingress Policies

Question: 193

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.allowedBinaryAuthorizationPolicies to the list of allowed Binary Authorization policy names.
- C. Set the organization policy constraint constraints/compute.trustedimageProjects to the list of projects 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:

To ensure that only trusted container images are deployed on Cloud Run, you can implement Binary Authorization, which is a deploy-time security control that ensures only trusted images are used. Set Up Binary Authorization:

Navigate to the Google Cloud Console.

Go to Security > Binary Authorization.

Configure the policy to include attestors that verify your trusted images.

Enable Binary Authorization on Cloud Run:

Go to the Cloud Run service.

Enable Binary Authorization on your existing Cloud Run services by selecting the appropriate Binary Authorization policy.

Set Organization Policy:

Go to the Organization Policies page in the Google Cloud Console.

Add a constraint for constraints/run.allowedBinaryAuthorizationPolicies.

Specify the list of allowed Binary Authorization policy names to enforce across your organization.

These steps ensure that any container image deployed on Cloud Run is validated against the specified Binary Authorization policies, preventing untrusted images from being deployed. Reference::

Binary Authorization Documentation

Enabling Binary Authorization on Cloud Run

Question: 194

You have a highly sensitive BigQuery workload that contains personally identifiable information (PII) that you want to ensure is not accessible from the internet. To prevent data exfiltration only requests from authorized IP addresses are allowed to query your BigQuery tables.

What should you do?

- A. Use service perimeter and create an access level based on the authorized source IP address as the condition.
- B. Use Google Cloud Armor security policies defining an allowlist of authorized IP addresses at the global HTTPS load

balancer.

- C. Use the Restrict allowed Google Cloud APIs and services organization policy constraint along with Cloud Data Loss Prevention (DLP).
- D. Use the Restrict Resource service usage organization policy constraint along with Cloud Data Loss Prevention (DLP).

Answer: A

Explanation:

Enable VPC Service Controls:

VPC Service Controls help mitigate the risk of data exfiltration by allowing you to define a security perimeter around GCP resources.

Set up a service perimeter around your BigQuery project to restrict data access to within the defined perimeter.

Create Access Levels:

In the Google Cloud Console, navigate to the Access Context Manager.

Define access levels based on IP address conditions, specifying the authorized source IP addresses that are allowed to access your BigQuery resources.

These access levels are used to enforce policies that restrict who can access your sensitive data based on their IP addresses.

Apply Service Perimeter with Access Levels:

Apply the created access levels to the service perimeter to ensure that only requests originating from the specified IP addresses are able to access BigQuery tables.

This setup ensures that the sensitive PII data is not accessible from unauthorized IP addresses, reducing the risk of data exfiltration.

Reference::

VPC Service Controls

Access Context Manager

Defining Access Levels

Question: 195

Your organization is moving virtual machines (VMs) to Google Cloud. You must ensure that operating system images that are used across your projects are trusted and meet your security requirements. What should you do?

- A. Implement an organization policy to enforce that boot disks can only be created from images that come from the trusted image project.
- B. Create a Cloud Function that is automatically triggered when a new virtual machine is created from the trusted image repository. Verify that the image is not deprecated.
- C. Implement an organization policy constraint that enables the Shielded VM service on all projects to enforce the trusted image repository usage.
- D. Automate a security scanner that verifies that no common vulnerabilities and exposures (CVEs) are present in your trusted image repository.

Answer: A

Explanation:

Define Trusted Image Projects:

Identify the project or projects where your trusted operating system images are stored.

Ensure these images meet your organization's security requirements and are regularly updated to mitigate vulnerabilities.

Create an Organization Policy:

Navigate to the Organization Policies page in the Google Cloud Console.

Create a policy constraint that restricts the creation of boot disks to only those images stored in your trusted image project(s).

The policy constraint to use is constraints/compute.trustedImageProjects.

Apply the Policy:

Apply this organization policy at the appropriate level (organization, folder, or project) to enforce that all new VM instances use images from the trusted repository.

This ensures consistency in the security posture across all projects within the organization.

Monitor Compliance:

Regularly monitor the compliance with this policy using audit logs and other monitoring tools.

Update the trusted images as necessary to ensure they remain secure and compliant with your security standards.

Reference::

Organization Policy Service

Trusted Image Projects Constraint

Question: 196

You define central security controls in your Google Cloud environment for one of the folders in your organization you set an organizational policy to deny the assignment of external IP addresses to VMs. Two days later you receive an alert about a new VM with an external IP address under that folder. What could have caused this alert?

- A. The VM was created with a static external IP address that was reserved in the project before the organizational policy rule was set.
- B. The organizational policy constraint wasn't properly enforced and is running in "dry run mode."
- C. At project level, the organizational policy control has been overwritten with an 'allow' value.
- D. The policy constraint on the folder level does not have any effect because of an 'allow' value for that constraint on the organizational level.

Answer: C

Explanation:

Understand Organization Policies:

Organization policies allow you to enforce restrictions on Google Cloud resources to adhere to your organization's security and compliance requirements.

Policies can be set at the organization, folder, or project level, with project-level policies able to override higher-level policies unless explicitly prevented.

Identify the Policy Constraint:

The specific constraint in question is likely constraints/compute.vmExternalIpAccess, which controls whether VMs can have external IP addresses.

Check Policy Overwrites:

Navigate to the Organization Policies page in the Google Cloud Console.

Check the policy settings at the project level under the affected folder to see if there is an override in place with an 'allow' value.

This override would permit the creation of VMs with external IP addresses despite the higher-level restriction.

Resolve the Policy Conflict:

If an override is found, remove or modify the project-level policy to align with the organizational policy denying external IP addresses.

Communicate with project administrators to ensure they understand and comply with the overarching security policies.

Reference::

Organization Policy Best Practices

Managing Policy Constraints

Question: 197

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. Container Threat Detection
- B. Web Security Scanner
- C. Rapid Vulnerability Detection
- D. Virtual Machine Threat Detection

Answer: D

Explanation:

Enable Security Command Center (SCC):

SCC provides centralized visibility and control over your cloud resources' security status.

Ensure that SCC is enabled in your Google Cloud environment.

Configure Virtual Machine Threat Detection (VMTD):

VMTD is part of SCC and specializes in detecting threats within VM instances, such as cryptocurrency mining malware.

Navigate to the SCC settings in the Google Cloud Console.

Activate VMTD:

Enable VMTD for the projects or resources where you want to monitor and detect potential threats.

VMTD uses behavioral analysis to identify anomalies indicative of unauthorized mining activities.

Monitor and Respond to Alerts:

VMTD generates alerts when it detects suspicious activities, such as unauthorized cryptocurrency mining.

Set up appropriate response actions, such as notifications, automatic remediation, or manual investigation, to handle these alerts.

Reference::

Question: 198

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, `constraints/compute.vnExternallpAccess`, 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:

Cloud NAT (Network Address Translation) allows instances without external IP addresses to access the internet securely.

In the Google Cloud Console, navigate to the VPC Network section and select Cloud NAT.

Create a new Cloud NAT configuration, specifying the VPC and region where your Compute Engine VMs are deployed.

Configure Cloud NAT:

Ensure that the Cloud NAT instance is configured to provide outbound internet connectivity for the VMs in your specified subnet.

This setup allows the VMs to access the internet for package updates and external installations without requiring public IP addresses.

Enable Private Google Access:

Private Google Access allows VMs in a subnet to reach Google APIs and services using internal IP addresses.

In the Google Cloud Console, navigate to the VPC Network section and select Subnets.

Edit the subnet used by your Compute Engine VMs and enable Private Google Access.

Update DevOps Scripts:

Ensure that your DevOps scripts are updated to work with the new network configuration.

Test the build process to confirm that the VMs can access necessary resources and complete the build pipeline successfully.

Reference::

Question: 199

Your company recently published a security policy to minimize the usage of service account keys. On-premises Windows-based applications are interacting with Google Cloud APIs. You need to implement Workload Identity Federation (WIF) with your identity provider on-premises. What should you do?

- A. Set up a workload identity pool with your corporate Active Directory Federation Service (ADFS) Configure a rule to let principals in the pool impersonate the Google Cloud service account.
- B. Set up a workload identity pool with your corporate Active Directory Federation Service (ADFS) Let all principals in the pool impersonate the Google Cloud service account.
- C. Set up a workload identity pool with an OpenID Connect (OIDC) service on the name machine Configure a rule to let principals in the pool impersonate the Google Cloud service account.
- D. Set up a workload identity pool with an OpenID Connect (OIDC) service on the same machine Let all principals in the pool impersonate the Google Cloud service account.

Answer: A

Explanation:

To minimize the usage of service account keys and implement Workload Identity Federation (WIF) with your on-premises identity provider, you can use a workload identity pool integrated with your corporate Active Directory Federation Service (ADFS). This setup allows your on-premises Windowsbased applications to authenticate to Google Cloud APIs without using long-lived service account keys.

Set Up a Workload Identity Pool:

In the Google Cloud Console, go to IAM & Admin > Workload Identity Federation.

Create a new workload identity pool.

Configure the pool to trust your corporate ADFS by specifying the federation provider details.

Create a Workload Identity Provider:

Within the created pool, set up a new provider for ADFS.

Configure the provider with the necessary details such as the issuer URL and credentials.

Configure Impersonation Rules:

Set up rules to allow principals in the workload identity pool to impersonate specific Google Cloud service accounts.

This is done by specifying the identity provider and the conditions under which the service accounts can be impersonated.

Update Applications:

Modify your on-premises applications to use the configured ADFS authentication to obtain tokens. These tokens can then be exchanged for Google Cloud access tokens to interact with Google Cloud APIs securely.

By setting up the workload identity pool and configuring impersonation rules, you achieve secure authentication without needing to distribute and manage service account keys.

Reference::

Workload Identity Federation Documentation

Federating On-Premises Identities to Workload Identity Federation

Question: 200

You have stored company approved compute images in a single Google Cloud project that is used as an image repository. This project is protected with VPC Service Controls and exists in the perimeter along with other projects in your organization. This lets other projects deploy images from the image repository project. A team requires deploying a third-party disk image that is stored in an external Google Cloud organization. You need to grant read access to the disk image so that it can be deployed into the perimeter.

What should you do?

- A.
 - 1 Update the perimeter
 - 2 Configure the egressTo field to set identity Type to any_identity.
 - 3 Configure the egressFrom field to include the external Google Cloud project number as an allowed resource and the serviceName to compute.googleapis.com.
- B. * Allow the external project by using the organizational policy constraints/compute.trustedImageProjects.
- C.
 - 1 Update the perimeter
 - 2 Configure the egressTo field to include the external Google Cloud project number as an allowed resource and the serviceName to compute.googleapis.com.
 - 3 Configure the egressFrom field to set identity Type to any_identity.
- D.
 - 1 Update the perimeter
 - 2 Configure the ingressFrom field to set identityType to any_identity.
 - 3 Configure the ingressTo field to include the external Google Cloud project number as an allowed resource and the serviceName to compute.googleapis.com.

Answer: A

Explanation:

To grant read access to a third-party disk image stored in an external Google Cloud organization so it can be deployed into a VPC Service Controls perimeter, you need to update the service perimeter to allow egress traffic from your projects to the external project.

Update the Service Perimeter:

Go to the Google Cloud Console, navigate to Security > VPC Service Controls.

Select the appropriate service perimeter that includes your image repository project.

Configure Egress Policy:

Within the perimeter settings, configure the egressTo field to allow traffic to the external project.

Set the identityType to ANY_IDENTITY to permit any principal to access the external project for this specific egress rule.

Specify External Project and Service:

In the egressFrom field, include the external Google Cloud project number as an allowed resource.

Set the serviceName to compute.googleapis.com to specifically allow access to the Compute Engine service in the external project.

This configuration permits your internal projects to read the disk image from the external project while maintaining the security boundaries established by the service perimeter.

Reference::

VPC Service Controls Documentation

Configuring Service Perimeters

Question: 201

Your organization recently activated the Security Command Center (SCC) standard tier. There are a few Cloud Storage buckets that were accidentally made accessible to the public. You need to investigate the impact of the incident and remediate it.

What should you do?

- A.
 - 1 Remove the Identity and Access Management (IAM) granting access to all users from the buckets
 - 2 Apply the organization policy storage.uniformBucketLevelAccess to prevent regressions
 - 3 Query the data access logs to report on unauthorized access
- B.
 - 1 Change bucket permissions to limit access
 - 2 Query the data access audit logs for any unauthorized access to the buckets
 - 3 After the misconfiguration is corrected, mute the finding in the Security Command Center
- C.
 - 1 Change permissions to limit access for authorized users
 - 2 Enforce a VPC Service Controls perimeter around all the production projects to immediately stop any unauthorized access
 - 3 Review the administrator activity audit logs to report on any unauthorized access
- D.
 - 1 Change the bucket permissions to limit access
 - 2 Query the buckets usage logs to report on unauthorized access to the data
 - 3 Enforce the organization policy storage.publicAccessPrevention to avoid regressions

■ ■

To investigate and remediate the issue of public access to Cloud Storage buckets, you can follow these steps:

Change Bucket Permissions:

Navigate to the Cloud Storage section in the Google Cloud Console.

For each affected bucket, remove any public access permissions (e.g., removing allUsers or allAuthenticatedUsers from the IAM policy).

Ensure that only authorized users have the necessary permissions to access the buckets.

Query Data Access Audit Logs:

Go to the Logging section in the Google Cloud Console.

Query the audit logs for the affected buckets to identify any unauthorized access. You can use filters to search for access by unauthorized users.

Correct the Misconfiguration:

After correcting the permissions, mute the relevant findings in the Security Command Center to indicate that the issue has been resolved.

This helps in maintaining a clear view of ongoing security issues and ensures the findings are not flagged again unless there's a new occurrence.

By following these steps, you ensure that the buckets are no longer publicly accessible, investigate any potential unauthorized access, and update the Security Command Center status to reflect the resolution of the issue.

Reference::

Cloud Storage IAM Permissions

Viewing Audit Logs

Security Command Center Documentation

Answer: B

Explanation:

Question: 202

Your organization uses the top-tier folder to separate application environments (prod and dev). The developers need to see all application development audit logs but they are not permitted to review production logs. Your security team can review all logs in production and development environments. You must grant Identity and Access Management (IAM) roles at the right resource level for the developers and security team while you ensure least privilege.

What should you do?

- A.
 - 1 Grant logging, viewer role to the security team at the organization resource level.
 - 2 Grant logging, viewer role to the developer team at the folder resource level that contains all the dev projects.
- B.
 - 1 Grant logging, viewer role to the security team at the organization resource level.
 - 2 Grant logging, admin role to the developer team at the organization resource level.
- C.
 - 1 Grant logging, admin role to the security team at the organization resource level.
 - 2 Grant logging, viewer role to the developer team at the folder resource level that contains all the dev projects.
- D.
 - 1 Grant logging, admin role to the security team at the organization resource level.
 - 2 Grant logging, admin role to the developer team at the organization resource level.

Answer: C

Explanation:

To ensure that the developers can view audit logs for the development environment and the security team can review all logs, you should grant IAM roles at the appropriate resource levels:

Grant logging.admin Role to the Security Team:

Assign the logging.admin role to the security team at the organization resource level.

This grants the security team full access to all logging data across the organization, including both production and development environments.

Grant logging.viewer Role to the Developer Team:

Assign the logging.viewer role to the developer team at the folder resource level that contains all the development projects.

This restricts the developers' access to only view logs in the development environment, ensuring they do not have access to production logs.

By using these roles and assigning them at the appropriate levels, you ensure that each team has the access they need while adhering to the principle of least privilege.

Reference::

IAM Roles for Cloud Logging

Resource Hierarchy in Google Cloud

Question: 203

Your organization's customers must scan and upload the contract and their driver license into a web portal in Cloud Storage. You must remove all personally identifiable information (PII) from files that are older than 12 months. Also you must archive the anonymized files for retention purposes.

What should you do?

- A. Set a time to live (TTL) of 12 months for the files in the Cloud Storage bucket that removes PII and moves the files to the archive storage class.
- B. Create a Cloud Data Loss Prevention (DLP) inspection job that de-identifies PII in files created more than 12 months ago and archives them to another Cloud Storage bucket. Delete the original files.
- C. Schedule a Cloud Key Management Service (KMS) rotation period of 12 months for the encryption keys of the Cloud Storage files containing PII to de-identify them. Delete the original keys.
- D. Configure the Autoclass feature of the Cloud Storage bucket to de-identify PII. Archive the files that are older than 12 months. Delete the original files.

Answer: B

Explanation:

To remove personally identifiable information (PII) from files older than 12 months and archive the anonymized files for retention purposes, you can use Google Cloud Data Loss Prevention (DLP). Create a Cloud DLP Inspection Job:

Go to the Cloud DLP section in the Google Cloud Console.

Create an inspection job that scans files in your Cloud Storage bucket for PII.

Configure the job to only target files that are older than 12 months.

Configure De-identification:

In the inspection job settings, configure de-identification actions to remove or obfuscate PII in the files.

Specify the transformation techniques appropriate for your data, such as masking or tokenization.

Archive Anonymized Files:

Set up the job to move the de-identified files to another Cloud Storage bucket designated for archival.

Ensure this bucket has the appropriate retention policies and access controls in place.

Delete Original Files:

After de-identification and archiving, configure the job to delete the original files from the source bucket.

This approach ensures that PII is effectively removed from old files and that the anonymized data is securely archived, maintaining compliance with data retention and privacy policies.

Reference::

Cloud Data Loss Prevention Documentation

Setting Up DLP Jobs

Cloud Storage Documentation

Question: 204

After completing a security vulnerability assessment, you learned that cloud administrators leave Google Cloud CLI sessions open for days. You need to reduce the risk of attackers who might exploit these open sessions by setting these sessions to the minimum duration.

What should you do?

- A. Set the session duration for the Google session control to one hour.
- B. Set the reauthentication frequency (or the Google Cloud Session Control) to one hour.
- C. Set the organization policy constraint `constraints/iam.allowServiceAccountCredentialLifetimeExtension` to one hour.
- D. Set the organization policy constraint `constraints/iam.serviceAccountKeyExpiryHours` to one hour and `inheritFromParent` to false.

Answer: B

Explanation:

To mitigate the risk posed by long-running Google Cloud CLI sessions, it is essential to enforce a reauthentication frequency. This ensures that users must periodically reauthenticate, reducing the window of opportunity for an attacker to exploit an open session. Setting the reauthentication frequency to one hour forces users to reauthenticate after this period, thereby limiting the duration an attacker can use a compromised session.

Access Google Cloud Console: Log in to your Google Cloud Console using your admin credentials. Navigate to Security Settings: Go to the "Security" section of the Cloud Console.

Set Session Control: Under the session management settings, locate the "Reauthentication frequency" setting. This controls how often users must reauthenticate.

Configure Reauthentication Frequency: Set the reauthentication frequency to "1 hour". This configuration will force users to reauthenticate every hour, thus limiting the duration of each session. **Save Changes:** Confirm and save your changes.

This setting will now apply to all users, ensuring that open sessions are minimized to a duration of one hour.

Reference::

Google Cloud IAM Documentation

Google Cloud Security Best Practices

Question: 205

Your application is deployed as a highly available cross-region solution behind a global external HTTP(S) load balancer.

You notice significant spikes in traffic from multiple IP addresses but it is unknown whether the IPs are malicious. You are concerned about your application's availability. You want to limit traffic from these clients over a specified time interval.

What should you do?

- A. Configure a rate_based_ban action by using Google Cloud Armor and set the ban_duration_sec parameter to the specified time interval.
- B. Configure a deny action by using Google Cloud Armor to deny the clients that issued too many requests over the specified time interval.
- C. Configure a throttle action by using Google Cloud Armor to limit the number of requests per client over a specified time interval.
- D. Configure a firewall rule in your VPC to throttle traffic from the identified IP addresses.

Answer: A

Explanation:

To handle significant traffic spikes and potentially malicious IPs, you can use Google Cloud Armor to configure rate-based bans. This approach allows you to automatically ban clients that exceed a predefined request rate, protecting your application from potential denial-of-service attacks.

Access Google Cloud Console: Log in to your Google Cloud Console.

Navigate to Google Cloud Armor: Go to the "Security" section and select "Google Cloud Armor". Create Security Policy: Create a new security policy or edit an existing one. Add a new rule to the policy.

Configure Rate-Based Ban: Set the action to rate_based_ban. Define the rate limit (e.g., requests per second) and set the ban_duration_sec parameter to the desired time interval.

Apply the Policy: Apply the security policy to your backend service or load balancer.

Monitor and Adjust: Monitor the traffic patterns and adjust the rate limits and ban durations as necessary to balance security and availability.

Reference::

Google Cloud Armor Documentation

Rate Limiting with Cloud Armor

Question: 206

You have numerous private virtual machines on Google Cloud. You occasionally need to manage the servers through Secure Socket Shell (SSH) from a remote location. You want to configure remote access to the servers in a manner that optimizes security and cost efficiency.

What should you do?

- A. Create a site-to-site VPN from your corporate network to Google Cloud.
- B. Configure server instances with public IP addresses Create a firewall rule to only allow traffic from your corporate IPs.
- C. Create a firewall rule to allow access from the Identity-Aware Proxy (IAP) IP range Grant the role of an IAP-secured Tunnel User to the administrators.
- D. Create a jump host instance with public IP Manage the instances by connecting through the jump host.

Answer: C

Explanation:

Using Identity-Aware Proxy (IAP) for managing SSH access to private VMs ensures secure access control and avoids the need for public IPs. IAP allows you to enforce identity-based access control policies.

Enable IAP: Ensure that IAP is enabled for your project. This can be done via the Google Cloud Console under "Security" -> "Identity-Aware Proxy".

Set Up Firewall Rule: Create a firewall rule to allow SSH traffic from the IAP IP ranges.

Navigate to "VPC network" -> "Firewall".

Create a new rule allowing ingress traffic on port 22 (SSH) from the IAP IP ranges.

Assign IAP-Secured Tunnel User Role: Grant the roles/iap.tunnelResourceAccessor role to the administrators who need SSH access.

Go to "IAM & Admin" -> "IAM".

Assign the IAP-Secured Tunnel User role to the relevant users or groups.

SSH Using IAP: Administrators can now use IAP to SSH into the instances. This can be done using the `gcloud` command:

```
gcloud compute ssh [INSTANCE_NAME] --tunnel-through-iap
```

Reference::

Using Identity-Aware Proxy for TCP forwarding

Google Cloud Firewall Rules

Question: 207

Your organization has on-premises hosts that need to access Google Cloud APIs You must enforce private connectivity between these hosts minimize costs and optimize for operational efficiency What should you do?

- A. Route all on-premises traffic to Google Cloud through an IPsec VPN tunnel to a VPC with Private Google Access enabled.
- B. Set up VPC peering between the hosts on-premises and the VPC through the internet.
- C. Enforce a security policy that mandates all applications to encrypt data with a Cloud Key Management. Service (KMS) key before you send it over the network.
- D. Route all on-premises traffic to Google Cloud through a dedicated or Partner interconnect to a VPC with Private Google Access enabled.

Answer: D

Explanation:

To enforce private connectivity between on-premises hosts and Google Cloud APIs while optimizing for cost and operational efficiency, using a dedicated or Partner Interconnect is the best solution. This setup ensures a reliable, high-bandwidth connection with private IP addressing.

Choose Interconnect Type: Decide between Dedicated Interconnect and Partner Interconnect based on your bandwidth needs and proximity to Google Cloud locations.

Set Up Interconnect:

For Dedicated Interconnect, order circuits through the Google Cloud Console.

For Partner Interconnect, select a supported service provider and order the connection through them.

Configure VPC and Private Google Access:

In your VPC, enable Private Google Access to allow on-premises hosts to access Google APIs privately.

Go to "VPC network" -> "Private Google Access" and enable it for your subnets.

Establish Connectivity: Work with your network team and (if applicable) your Partner Interconnect provider to set up the physical and logical connections.

Test Connectivity: Verify that on-premises hosts can reach Google Cloud services using private IP addresses.

Reference::

Google Cloud Interconnect Overview

Configuring Private Google Access

Question: 208

Your organization's record data exists in Cloud Storage. You must retain all record data for at least seven years. This policy must be permanent.

What should you do?

- A.
 - 1 Identify buckets with record data
 - 2 Apply a retention policy and set it to retain for seven years
 - 3 Monitor the bucket by using log-based alerts to ensure that no modifications to the retention policy occurs
- B.
 - 1 Identify buckets with record data
 - 2 Apply a retention policy and set it to retain for seven years
 - 3 Remove any Identity and Access Management (IAM) roles that contain the storage buckets
- C.
 - 1 Identify buckets with record data
 - 2 Enable the bucket policy only to ensure that data is retained

- 3 Enable bucket lock
- D. * 1 Identify buckets with record data
- 2 Apply a retention policy and set it to retain for seven years
- 3 Enable bucket lock

Answer: D

Explanation:

To ensure that your organization's record data is retained for at least seven years in Cloud Storage, you need to apply a retention policy and enable bucket lock. This prevents the policy from being altered or the data from being deleted before the retention period ends.

Identify Buckets: Determine which Cloud Storage buckets contain the record data that needs to be retained.

Apply Retention Policy:

Go to the Google Cloud Console and navigate to "Cloud Storage".

Select the bucket you identified.

Go to the "Retention" tab and set a retention policy to retain objects for seven years.

Enable Bucket Lock:

Once the retention policy is set, you need to lock the bucket to make the retention policy permanent.

This is done by enabling the bucket lock. Go to the "Retention" tab and click "Lock".

Confirm and Monitor:

Confirm that the bucket lock is applied.

Monitor the bucket using log-based alerts to ensure compliance.

Reference::

Cloud Storage Retention Policy

Cloud Storage Bucket Lock

Question: 209

Your organization wants to protect all workloads that run on Compute Engine VM to ensure that the instances weren't compromised by boot-level or kernel-level malware. Also, you need to ensure that data in use on the VM cannot be read by the underlying host system by using a hardware-based solution.

What should you do?

- A. • 1 Use Google Shielded VM including secure boot Virtual Trusted Platform Module (vTPM) and integrity monitoring
- 2 Create a Cloud Run function to check for the VM settings generate metrics and run the function regularly
- B. • 1 Activate Virtual Machine Threat Detection in Security Command Center (SCC Premium)
- 2 Monitor the findings in SCC
- C. * 1 Use Google Shielded VM including secure boot Virtual Trusted Platform Module (vTPM) and integrity monitoring
- 2 Activate Confidential Computing
- 3 Enforce these actions by using organization policies
- D. • 1 Use secure hardened images from the Google Cloud Marketplace
- 2 When deploying the images activate the Confidential Computing option
- 3 Enforce the use of the correct images and Confidential Computing by using organization policies

Answer: C

Explanation:

Use Google Shielded VM including secure boot Virtual Trusted Platform Module (vTPM) and integrity monitoring: Shielded VMs provide verifiable integrity of the VM by ensuring that it was not tampered with or compromised at the boot level. They use features like Secure Boot, vTPM, and integrity monitoring to detect and prevent malicious changes to the VM's operating system and firmware. Activate Confidential Computing: Confidential Computing provides a secure environment for processing sensitive data. It uses hardware-based enclaves to protect data in use by ensuring it cannot be accessed by the underlying host or any other unauthorized entity. By leveraging Intel SGX or AMD SEV, it ensures that data remains encrypted even when it is being processed.

Enforce these actions by using organization policies: Organization policies can enforce the use of Shielded VMs and Confidential Computing across your organization. This ensures that all VMs comply with these security measures without requiring manual configuration for each VM.

Reference:

Shielded VMs documentation

Confidential Computing documentation

Organization Policies documentation

Question: 210

Your company uses Google Cloud and has publicly exposed network assets. You want to discover the assets and perform a security audit on these assets by using a software tool in the least amount of time.

What should you do?

- A. Run a platform security scanner on all instances in the organization.
- B. Notify Google about the pending audit and wait for confirmation before performing the scan.
- C. Contact a Google approved security vendor to perform the audit.
- D. Identify all external assets by using Cloud Asset Inventory and then run a network security scanner against them.

Answer: D

Explanation:

Cloud Asset Inventory: Using Cloud Asset Inventory allows you to quickly identify all the external assets and resources in your Google Cloud environment. This includes information about your projects, instances, storage buckets, and more.

This step is crucial for understanding the scope of your audit. Network Security Scanner: Once you have identified the external assets, you can run a network security scanner to assess the security of these assets. Network security scanners can help identify vulnerabilities and potential security risks quickly.

Question: 211

Your organization is using Active Directory and wants to configure Security Assertion Markup Language (SAML). You must set up and enforce single sign-on (SSO) for all users.

What should you do?

- A.
 - 1. Manage SAML profile assignments.

- 2. Enable OpenID Connect (OIDC) in your Active Directory (AD) tenant.
- 3. Verify the domain.

B.

- 1. Create a new SAML profile.
- 2. Upload the X.509 certificate.
- 3. Enable the change password URL.
- 4. Configure Entity ID and ACS URL in your IdP.

C.

- 1- Create a new SAML profile.
- 2. Populate the sign-in and sign-out page URLs.
- 3. Upload the X.509 certificate.
- 4. Configure Entity ID and ACS URL in your IdP

D.

- 1. Configure prerequisites for OpenID Connect (OIDC) in your Active Directory (AD) tenant
- 2. Verify the AD domain.
- 3. Decide which users should use SAML.
- 4. Assign the pre-configured profile to the select organizational units (OUs) and groups.

Answer: C

Explanation:

When configuring SAML-based Single Sign-On (SSO) in an organization that's using Active Directory, the general steps would involve setting up a SAML profile, specifying the necessary URLs for sign-in and sign-out processes, uploading an X.509 certificate for secure communication, and setting up the Entity ID and Assertion Consumer Service (ACS) URL in the Identity Provider (which in this case would be Active Directory).

Question: 212

Your customer has an on-premises Public Key Infrastructure (PKI) with a certificate authority (CA).

You need to issue certificates for many HTTP load balancer frontends. The on-premises PKI should be minimally affected due to many manual processes, and the solution needs to scale.

What should you do?

- A. Use Certificate Manager to issue Google managed public certificates and configure it at HTTP the load balancers in your infrastructure as code (IaC).
- B. Use Certificate Manager to import certificates issued from on-premises PKI and for the frontends. Leverage the gcloud tool for importing
- C. Use a subordinate CA in the Google Certificate Authority Service from the on-premises PKI system to issue certificates for the load balancers.
- D. Use the web applications with PKCS12 certificates issued from subordinate CA based on OpenSSL on-premises Use the gcloud tool for importing. Use the External TCP/UDP Network load balancer instead of an external HTTP Load Balancer.

Answer: C

Explanation:

This approach allows you to leverage your existing on-premises PKI infrastructure while minimizing its impact and manual processes. By creating a subordinate CA in Google's Certificate Authority Service, you can automate the process of issuing certificates for your HTTP load balancer frontends. This solution scales well as the number of load balancers increases.

Question: 213

Your organization wants full control of the keys used to encrypt data at rest in their Google Cloud environments. Keys must be generated and stored outside of Google and integrate with many Google Services including BigQuery. What should you do?

- A. Create a Cloud Key Management Service (KMS) key with imported key material Wrap the key for protection during import. Import the key generated on a trusted system in Cloud KMS.
- B. Create a KMS key that is stored on a Google managed FIPS 140-2 level 3 Hardware Security Module (HSM) Manage the Identity and Access Management (IAM) permissions settings, and set up the key rotation period.
- C. Use Cloud External Key Management (EKM) that integrates with an external Hardware Security Module (HSM) system from supported vendors.
- D. Use customer-supplied encryption keys (CSEK) with keys generated on trusted external systems Provide the raw CSEK as part of the API call.

Answer: C

Explanation:

Use Cloud External Key Management (EKM) that integrates with an external Hardware Security Module (HSM) system from supported vendors: Cloud EKM allows you to use encryption keys that are managed externally to Google Cloud. This means you can generate and store your keys in an onpremises HSM or another supported external HSM service, and integrate these keys with various Google Cloud services.

Integration with Google Services: Cloud EKM integrates seamlessly with many Google Cloud services, including BigQuery, Cloud Storage, Compute Engine, and more. This provides you with full control over your encryption keys while still taking advantage of Google Cloud's powerful services.

Reference:

Cloud External Key Management (EKM) documentation

External Key Management overview

Question: 214

Your organization wants to be General Data Protection Regulation (GDPR) compliant You want to ensure that your DevOps teams can only create Google Cloud resources in the Europe regions. What should you do?

- A. Use the org policy constraint "Restrict Resource Service Usage"* on your Google Cloud organization node.
- B. Use Identity and Access Management (IAM) custom roles to ensure that your DevOps team can only create resources in the Europe regions
- C. Use the org policy constraint "Google Cloud Platform - Resource Location Restriction" on your Google Cloud organization node.
- D. Use Identity-Aware Proxy (IAP) with Access Context Manager to restrict the location of Google Cloud resources.

Answer: C

Explanation:

Use the org policy constraint "Google Cloud Platform - Resource Location Restriction" on your Google Cloud organization

node: This organizational policy constraint allows you to restrict the locations where your resources can be created. By setting this constraint to allow only Europe regions, you can ensure compliance with GDPR and other regional regulations.

Implementation: To implement this, you need to configure the organization policy with the constraint constraints/gcp.resourceLocations. You can specify allowed regions such as europe-west1 and europe-west4 to ensure resources are only created in these locations.

Reference:

Resource Location Restriction documentation
GDPR compliance on Google Cloud

Question: 215

Your organization uses BigQuery to process highly sensitive, structured datasets. Following the "need to know" principle, you need to create the Identity and Access Management (IAM) design to meet the needs of these users:

- Business user must access curated reports.
- Data engineer: must administrate the data lifecycle in the platform.
- Security operator: must review user activity on the data platform.

What should you do?

- A. Configure data access log for BigQuery services, and grant Project Viewer role to security operators.
- B. Generate a CSV data file based on the business user's needs, and send the data to their email addresses.
- C. Create curated tables in a separate dataset and assign the role roles/bigquery.dataViewer.
- D. Set row-based access control based on the "region" column, and filter the record from the United States for data engineers.

Answer: C

Explanation:

This option directly addresses the needs of the business user who must access curated reports. By creating curated tables in a separate dataset, you can control access to specific data. Assigning the roles/bigquery.dataViewer role allows the business user to view the data in BigQuery.

Question: 216

Your organization operates Virtual Machines (VMs) with only private IPs in the Virtual Private Cloud (VPC) with internet access through Cloud NAT. Everyday, you must patch all VMs with critical OS updates and provide summary reports

What should you do?

- A. Validate that the egress firewall rules allow any outgoing traffic Log in to each VM and execute OS specific update commands Configure the Cloud Scheduler job to update with critical patches daily for daily updates.
- B. Ensure that VM Manager is installed and running on the VMs. In the OS patch management service. configure the patch jobs to update with critical patches daily.
- C. Assign public IPs to VMs. Validate that the egress firewall rules allow any outgoing traffic Log in to each VM. and configure a daily cron job to enable for OS updates at night during low activity periods.
- D. Copy the latest patches to the

Cloud Storage bucket. Log in to each VM. download the patches from the bucket, and install them.

Answer: B

Explanation:

VM Manager is a suite of tools that can be used to manage operating systems for large virtual machine (VM) fleets running Windows and Linux on Compute Engine. It helps drive efficiency through automation and reduces the operational burden of maintaining these VM fleets. VM Manager includes several services such as OS patch management, OS inventory management, and OS configuration management. By using VM Manager, you can apply patches, collect operating system information, and install, remove, or auto-update software packages. The suite provides a high level of control and automation for managing large VM fleets on Google Cloud.

<https://cloud.google.com/compute/docs/vm-manager>

Question: 217

A company is using Google Kubernetes Engine (GKE) with container images of a mission-critical application. The company wants to scan the images for known security issues and securely share the report with the security team without exposing them outside Google Cloud.

What should you do?

A. 1. Enable Container Threat Detection in the Security Command Center Premium tier.

• 2. Upgrade all clusters that are not on a supported version of GKE to the latest possible GKE version.

• 3. View and share the results from the Security Command Center

B.

• 1. Use an open source tool in Cloud Build to scan the images.

• 2. Upload reports to publicly accessible buckets in Cloud Storage by using gsutil

• 3. Share the scan report link with your security department.

C.

• 1. Enable vulnerability scanning in the Artifact Registry settings.

• 2. Use Cloud Build to build the images

• 3. Push the images to the Artifact Registry for automatic scanning.

• 4. View the reports in the Artifact Registry.

D.

• 1. Get a GitHub subscription.

• 2. Build the images in Cloud Build and store them in GitHub for automatic scanning

• 3. Download the report from GitHub and share with the Security Team

Answer: C

Explanation:

"The service evaluates all changes and remote access attempts to detect runtime attacks in near-real time."

<https://cloud.google.com/security-command-center/docs/concepts-container-threat-detection-overview> This has nothing to do with KNOWN security Vulns in images

Question: 218

An administrative application is running on a virtual machine (VM) in a managed group at port 5601 inside a Virtual

Private Cloud (VPC) instance without access to the internet currently. You want to expose the web interface at port 5601 to users and enforce authentication and authorization Google credentials. What should you do?

- A. Modify the VPC routing with the default route point to the default internet gateway Modify the VPC Firewall rule to allow access from the internet 0.0.0.0/0 to port 5601 on the application instance.
- B. Configure the bastion host with OS Login enabled and allow connection to port 5601 at VPC firewall Log in to the bastion host from the Google Cloud console by using SSH-in-browser and then to the web application
- C. Configure an HTTP Load Balancing instance that points to the managed group with Identity-Aware Proxy (IAP) protection with Google credentials Modify the VPC firewall to allow access from IAP network range
- D. Configure Secure Shell Access (SSH) bastion host in a public network, and allow only the bastion host to connect to the application on port 5601. Use a bastion host as a jump host to connect to the application

Answer: C

Explanation:

This approach allows you to expose the web interface securely by using Identity-Aware Proxy (IAP), which provides authentication and authorization with Google credentials. The HTTP Load Balancer can distribute traffic to the VMs in the managed group, and the VPC firewall rule ensures that access is allowed from the IAP network range.

Question: 219

Your company is concerned about unauthorized parties gaming access to the Google Cloud environment by using a fake login page. You must implement a solution to protect against person-in- the-middle attacks. Which security measure should you use?

- A. Text message or phone call code
- B. Security key
- C. Google Authenticator application
- D. Google prompt

Answer: B

Explanation:

A security key is a physical device that you can use for two-step verification, providing an additional layer of security for your Google Account. Security keys can defend against phishing and man-in-the- middle attacks, making your login process more secure.

Question: 220

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. Change the encryption type on the bucket to CMEK, and rewrite the objects

Answer: D

Explanation:

Rewriting the objects in-place within the same bucket, specifying the new CMEK for encryption, allows you to re-encrypt the data without downloading and re-uploading it, thus minimizing costs and time.

<https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys>

Question: 221

Your company conducts clinical trials and needs to analyze the results of a recent study that are stored in BigQuery. The interval when the medicine was taken contains start and stop dates. The interval data is critical to the analysis, but specific dates may identify a particular batch and introduce bias. You need to obfuscate the start and end dates for each row and preserve the interval data. What should you do?

- A. Use bucketing to shift values to a predetermined date based on the initial value.
- B. Extract the date using TimePartConfig from each date field and append a random month and year.
- C. Use date shifting with the context set to the unique ID of the test subject.
- D. Use the FFX mode of format preserving encryption (FPE) and maintain data consistency.

Answer: A

Explanation:

"Date shifting techniques randomly shift a set of dates but preserve the sequence and duration of a period of time. Shifting dates is usually done in context to an individual or an entity. That is, each individual's dates are shifted by an amount of time that is unique to that individual."

Question: 222

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 "prj-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:

When you use a customer-managed encryption key (CMEK) to secure a Cloud Storage bucket, the key and the bucket

must be located in the same region. In this case, the key is in europe-west3 and the bucket is in europe-west1, which is why you're unable to access the key.

Question: 223

You are deploying regulated workloads on Google Cloud. The regulation has data residency and data access requirements. It also requires that support is provided from the same geographical location as where the data resides.

What should you do?

- A. Enable Access Transparency Logging.
- B. Deploy resources only to regions permitted by data residency requirements
- C. Use Data Access logging and Access Transparency logging to confirm that no users are accessing data from another region.
- D. Deploy Assured Workloads.

Answer: D

Explanation:

Assured Workloads for Google Cloud allows you to deploy regulated workloads with data residency, access, and support requirements. It helps you configure your environment in a manner that aligns with specific compliance frameworks and standards.

Question: 224

You are migrating an application into the cloud The application will need to read data from a Cloud Storage bucket. Due to local regulatory requirements, you need to hold the key material used for encryption fully under your control and you require a valid rationale for accessing the key material. What should you do?

- A. Encrypt the data in the Cloud Storage bucket by using Customer Managed Encryption Keys. Configure an IAM deny policy for unauthorized groups
- B. Encrypt the data in the Cloud Storage bucket by using Customer Managed Encryption Keys backed by a Cloud Hardware Security Module (HSM). Enable data access logs.
- C. Generate a key in your on-premises environment and store it in a Hardware Security Module (HSM) that is managed on-premises Use this key as an external key in the Cloud Key Management Service (KMS). Activate Key Access Justifications (KAJ) and set the external key system to reject unauthorized accesses.
- D. Generate a key in your on-premises environment to encrypt the data before you upload the data to the Cloud Storage bucket Upload the key to the Cloud Key Management Service (KMS). Activate Key Access Justifications (KAJ) and have the external key system reject unauthorized accesses.

Answer: C

Explanation:

By generating a key in your on-premises environment and storing it in an HSM that you manage, you're ensuring that the key material is fully under your control. Using the key as an external key in Cloud KMS allows you to use the key with Google Cloud services without having the key stored on Google Cloud. Activating Key Access Justifications (KAJ) provides a reason every time the key is accessed, and you can configure the external key system to reject unauthorized access

attempts.

Question: 225

Your organization develops software involved in many open source projects and is concerned about software supply chain threats. You need to deliver provenance for the build to demonstrate the software is untampered.

What should you do?

A.

- 1- Generate Supply Chain Levels for Software Artifacts (SLSA) level 3 assurance by using Cloud

Build.

- 2. View the build provenance in the Security insights side panel within the Google Cloud console.

B.

- 1. Review the software process.
- 2. Generate private and public key pairs and use Pretty Good Privacy (PGP) protocols to sign the output software artifacts together with a file containing the address of your enterprise and point of contact.
- 3. Publish the PGP signed attestation to your public web page.

C.

- 1. Publish the software code on GitHub as open source.
- 2. Establish a bug bounty program, and encourage the open source community to review, report, and fix the vulnerabilities.

D.

- 1. Hire an external auditor to review and provide provenance.
- 2. Define the scope and conditions.
- 3. Get support from the Security department or representative.
- 4. Publish the attestation to your public web page.

Answer: A

Explanation:

Generate Supply Chain Levels for Software Artifacts (SLSA) level 3 assurance by using Cloud Build: SLSA is a framework for ensuring the integrity of software artifacts. By using Cloud Build, you can automate the build process and generate SLSA level 3 compliance, which includes verifiable build steps and provenance.

View the build provenance in the Security insights side panel within the Google Cloud console: The build provenance provides a detailed history of how the software was built, including the source

code, build process, and any dependencies. This information is accessible through the Security insights side panel in the Google Cloud console, allowing you to verify the integrity and authenticity of your software artifacts.

Reference:

[Supply Chain Levels for Software Artifacts \(SLSA\) documentation](#)

Cloud Build documentation

Security insights in Google Cloud console

Question: 226

You control network traffic for a folder in your Google Cloud environment. Your folder includes multiple projects and Virtual Private Cloud (VPC) networks. You want to enforce on the folder level that egress connections are limited only to IP range 10.58.5.0/24 and only from the VPC network dev-vpc. You want to minimize implementation and maintenance effort. What should you do?

A.

- 1. Attach external IP addresses to the VMs in scope.
- 2. Configure a VPC Firewall rule in "dev-vpc" that allows egress connectivity to IP range

10.58.5. 0/24 for all source addresses in this network.

B.

- 1. Attach external IP addresses to the VMs in scope.
- 2. Define and apply a hierarchical firewall policy on folder level to deny all egress connections and to allow egress to IP range 10.58.5.0/24 from network dev-vpc.

C.

- 1. Leave the network configuration of the VMs in scope unchanged.
- 2. Create a new project including a new VPC network "new-vpc."
- 3 Deploy a network appliance in "new-vpc" to filter access requests and only allow egress connections from -dev-vpc" to 10.58.5.0/24.

D.

- 1 Leave the network configuration of the VMs in scope unchanged
- 2 Enable Cloud NAT for dev-vpc" and restrict the target range in Cloud NAT to 10.58.5.0/24.

Answer: B

Explanation:

This approach allows you to control network traffic at the folder level. By attaching external IP addresses to the VMs in scope, you can ensure that the VMs have a unique, routable IP address for outbound connections. Then, by defining and applying a hierarchical firewall policy at the folder level, you can enforce that egress connections are limited to the specified IP range and only from the specified VPC network.

Question: 227

Your Google Cloud environment has one organization node, one folder named Apps." and several projects within that folder The organizational node enforces the constraints/iam.allowedPolicyMemberDomains organization policy, which allows members from the terramearth.com organization The "Apps" folder enforces the constraints/iam.allowedPolicyMemberDomains organization policy, which allows members from the flowlogistic.com organization. It also has the inheritFromParent: false property.

You attempt to grant access to a project in the Apps folder to the user testuser@terramearth.com. What is the result of your action and why?

- A. The action fails because a constraints/iam.allowedPolicyMemberDomains organization policy must be defined on the current project to deactivate the constraint temporarily.
- B. The action fails because a constraints/iam.allowedPolicyMemberDomains organization policy is in place and only members from the flowlogistic.com organization are allowed.
- C. The action succeeds because members from both organizations, terramearth. com or flowlogistic.com, are allowed on projects in the "Apps" folder
- D. The action succeeds and the new member is successfully added to the project's Identity and Access Management (IAM) policy because all policies are inherited by underlying folders and projects.

Answer: B

Explanation:

The action fails because a constraints/iam.allowedPolicyMemberDomains organization policy is in place and only members from the flowlogistic.com organization are allowed. The inheritFromParent: false property on the "Apps" folder means that it does not inherit the organization policy from the organization node. Therefore, only the policy set at the folder level applies, which allows only members from the flowlogistic.com organization. As a result, the attempt to grant access to the user testuser@terramearth.com fails because this user is not a member of the flowlogistic.com organization.

Question: 228

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:

Question: 229

Employees at your company use their personal computers to access your organization's Google Cloud console. You need to ensure that users can only access the Google Cloud console from their corporate-issued devices and verify that they have a valid enterprise certificate.

What should you do?

- A. Implement an Identity and Access Management (IAM) conditional policy to verify the device certificate
- B. Implement a VPC firewall policy. Activate packet inspection and create an allow rule to validate and verify the device certificate.
- C. Implement an organization policy to verify the certificate from the access context.
- D. Implement an Access Policy in BeyondCorp Enterprise to verify the device certificate. Create an access binding with the access policy just created.

Answer: D

Explanation:

Question: 230

You are developing a new application that uses exclusively Compute Engine VMs. Once a day, this application will execute five different batch jobs. Each of the batch jobs requires a dedicated set of permissions on Google Cloud resources outside of your application. You need to design a secure access concept for the batch jobs that adheres to the least-privilege principle. What should you do?

- A.
 - 1. Create a general service account "g-sa" to execute the batch jobs.
 - 2. Grant the permissions required to execute the batch jobs to g-sa.
 - 3. Execute the batch jobs with the permissions granted to g-sa.
- B.
 - 1. Create a general service account "g-sa" to orchestrate the batch jobs.
 - 2. Create one service account per batch job "Mb-sa-[1-5]," and grant only the permissions required

to run the individual batch jobs to the service accounts.

- 3. Grant the Service Account Token Creator role to g-sa Use g-sa to obtain short-lived access tokens for b-sa-[1-5] and to execute the batch jobs with the permissions of b-sa-[1-5].

C. 1. Create a workload identity pool and configure workload identity pool providers for each batch job

- 2 Assign the workload identity user role to each of the identities configured in the providers.
- 3. Create one service account per batch job Mb-sa-[1-5]" and grant only the permissions required

to run the individual batch jobs to the service accounts

- 4 Generate credential configuration files for each of the providers Use these files to execute the batch jobs with the permissions of b-sa-[1-5].

D.

- 1. Create a general service account "g-sa" to orchestrate the batch jobs.
- 2 Create one service account per batch job 'b-sa-[1-5]\ Grant only the permissions required to run the individual batch jobs to the service accounts and generate service account keys for each of these service accounts
- 3. Store the service account keys in Secret Manager. Grant g-sa access to Secret Manager and run the batch jobs with the permissions of b-sa-[1-5].

Answer: B

Explanation:

Question: 231

Your company's users access data in a BigQuery table. You want to ensure they can only access the data during working hours.

What should you do?

- A. Assign a BigQuery Data Viewer role along with an IAM condition that limits the access to specified working hours.
- B. Configure Cloud Scheduler so that it triggers a Cloud Functions instance that modifies the organizational policy constraints for BigQuery during the specified working hours.
- C. Assign a BigQuery Data Viewer role to a service account that adds and removes the users daily during the specified working hours
- D. Run a gsutil script that assigns a BigQuery Data Viewer role, and remove it only during the specified working hours.

Answer: A

Explanation:

To ensure that users can only access the data in a BigQuery table during working hours, you can assign the BigQuery Data Viewer role with an IAM condition that specifies the allowed access times. This method leverages IAM Conditions, which allow you to define and enforce time-based access policies. Here's how to do it:

Identify the BigQuery Table: Determine which BigQuery table(s) require restricted access.

Create an IAM Policy with Conditions: Define an IAM policy that includes a condition for time-based access. You can do this using the Google Cloud Console, gcloud command-line tool, or directly editing the IAM policy JSON.

Specify Working Hours: In the IAM condition, specify the time frame during which access is allowed. For example, you can set access to be allowed from 9 AM to 5 PM on weekdays.

Assign the Role with Conditions: Apply the policy to the users or groups who need access. Ensure that the condition is correctly attached to the BigQuery Data Viewer role.

Example using gcloud:

gcloud projects add-iam-policy-binding [PROJECT_ID] \

--member=user:[USER_EMAIL] \

--role=roles/bigquery.dataViewer \

--condition=expression="(request.time.getFullYear() == 2024) && (request.time.getDayOfWeek() in [1, 2, 3, 4, 5]) && (request.time.getHours() >= 9) && (request.time.getHours() < 17)",title="Working hours condition",description="Access limited to working hours" Reference:

Google Cloud IAM Conditions

Google Cloud BigQuery IAM Roles

Question: 232

Your organization wants to be compliant with the General Data Protection Regulation (GDPR) on Google Cloud You must implement data residency and operational sovereignty in the EU.

What should you do?

Choose 2 answers

- A. Limit the physical location of a new resource with the Organization Policy Service resource locations constraint."
- B. Use Cloud IDS to get east-west and north-south traffic visibility in the EU to monitor intra-VPC and mter-VPC communication.
- C. Limit Google personnel access based on predefined attributes such as their citizenship or geographic location by using Key Access Justifications
- D. Use identity federation to limit access to Google Cloud resources from non-EU entities.
- E. Use VPC Flow Logs to monitor intra-VPC and inter-VPC traffic in the EU.

Answer: A, C

Explanation:

https://cloud.google.com/architecture/framework/security/data-residency-sovereignty#manage_your_operational_sovereignty

To ensure compliance with GDPR and implement data residency and operational sovereignty in the EU, the following steps can be taken:

Limit Physical Location of Resources: Use the Organization Policy Service to enforce the resource locations constraint.

This ensures that all new resources are created within the specified regions (EU in this case).

Configure Organization Policy: Set up an organization policy that restricts the locations where new resources can be created. This is done through the Google Cloud Console or via the gcloud command-line tool.

Example:

```
gcloud resource-manager org-policies allow constraints/gcp.resourceLocations [europe- west1,europe-west2] --organization=YOUR_ORG_ID
```

Key Access Justifications (KAJ): Use Key Access Justifications to limit Google personnel's access to encryption keys based on attributes like their geographic location or citizenship.

Set Up KAJ: Implement KAJ policies to ensure that only authorized personnel within the EU can access encryption keys.

Reference:

Organization Policy Service

Key Access Justifications

Question: 233

You manage a mission-critical workload for your organization, which is in a highly regulated industry. The workload uses Compute Engine VMs to analyze and process the sensitive data after it is uploaded to Cloud Storage from the endpoint computers. Your compliance team has detected that this workload does not meet the data protection requirements for sensitive data. You need to meet these requirements;

- Manage the data encryption key (DEK) outside the Google Cloud boundary.
- Maintain full control of encryption keys through a third-party provider.
- Encrypt the sensitive data before uploading it to Cloud Storage
- Decrypt the sensitive data during processing in the Compute Engine VMs
- Encrypt the sensitive data in memory while in use in the Compute Engine VMs

What should you do?

Choose 2 answers

- A. Create a VPC Service Controls service perimeter across your existing Compute Engine VMs and Cloud Storage buckets
- B. Migrate the Compute Engine VMs to Confidential VMs to access the sensitive data.
- C. Configure Cloud External Key Manager to encrypt the sensitive data before it is uploaded to Cloud Storage and decrypt the sensitive data after it is downloaded into your VMs
- D. Create Confidential VMs to access the sensitive data.
- E. Configure Customer Managed Encryption Keys to encrypt the sensitive data before it is uploaded to Cloud Storage, and decrypt the sensitive data after it is downloaded into your VMs.

Answer: C, D

Explanation:

<https://cloud.google.com/confidential-computing/confidential-vm/docs/creating-cvm-instance#considerations>

Confidential VM does not support live migration. You can only enable Confidential Computing on a

VM when you first create the instance. [https://cloud.google.com/confidential-](https://cloud.google.com/confidential-computing/confidential-vm/docs/creating-cvm-instance)

[computing/confidential-vm/docs/creating-cvm-instance](https://cloud.google.com/confidential-computing/confidential-vm/docs/creating-cvm-instance)

Question: 234

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

Answer: B

Explanation:

To provide temporary write access to a Cloud Storage bucket with the minimum permissions necessary, you should: Identify the Compute Engine instance's default service account: Each Compute Engine instance has a default service account that is used to interact with other Google Cloud services.

Assign the storage.objectCreator role: This predefined IAM role grants permissions to create objects in a Cloud Storage bucket, which is sufficient for temporary write access. It does not grant permissions to read or delete objects, thus adhering to the principle of least privilege. Avoid using full permissions or long-lived keys: Options A and C suggest using broader permissions than necessary or embedding long-lived keys, which could pose a security risk if compromised. Service account impersonation (Option D) is not necessary for this task and would be more appropriate for scenarios where you need to assume a different identity with different permissions. Reference::

[Google Cloud documentation on IAM roles for Cloud Storage, which lists the storage.objectCreator role as providing permissions to create objects without granting full administrative access to the bucket1.](#)

[Best practices for access control in Cloud Storage recommend using the least privilege necessary and avoiding the use of long-lived service account keys2.](#)

Question: 235

Your organization is using Vertex AI Workbench Instances. You must ensure that newly deployed instances are automatically kept up-to-date and that users cannot accidentally alter settings in the operating system. What should you do?

- A. Enable the VM Manager and ensure the corresponding Google Compute Engine instances are added.
- B. Enforce the disableRootAccess and requireAutoUpgradeSchedule organization policies for newly deployed instances.
- C. Assign the AI Notebooks Runner and AI Notebooks Viewer roles to the users of the AI Workbench Instances.
- D. Implement a firewall rule that prevents Secure Shell access to the corresponding Google Compute Engine instances by using tags.

Answer: B

Explanation:

To ensure that Vertex AI Workbench Instances (formerly AI Platform Notebooks) are automatically updated and that users cannot modify operating system settings, it's crucial to implement organizational policies that enforce these requirements.

disableRootAccess Organization Policy:

This policy prevents users from obtaining root access on virtual machines. By enforcing this policy, you ensure that users cannot make unauthorized changes to the operating system settings, maintaining the integrity and security of the instances.

requireAutoUpgradeSchedule Organization Policy:

This policy mandates that virtual machines have an auto-upgrade schedule for their operating systems. By enforcing this policy, you ensure that instances are automatically kept up-to-date with the latest security patches and updates, reducing the risk of vulnerabilities.

Given the options:

Option A: Enabling VM Manager helps in managing updates and configurations but does not inherently prevent users from altering OS settings.

Option B: Enforcing the `disableRootAccess` and `requireAutoUpgradeSchedule` organization policies directly addresses both requirements: preventing unauthorized OS modifications and ensuring automatic updates.

Option C: Assigning specific roles controls user permissions but does not enforce OS-level restrictions or automatic updates.

Option D: Implementing firewall rules to prevent SSH access adds a layer of security but does not ensure automatic updates or prevent OS modifications through other means.

Therefore, Option B is the most effective approach, as it directly enforces the necessary policies to meet both requirements.

Reference::

Organization Policy Service

VM Manager Overview

Question: 236

You have just created a new log bucket to replace the `_Default` log bucket. You want to route all log entries that are currently routed to the `_Default` log bucket to this new log bucket in the most efficient manner. What should you do?

- A. Create a user-defined sink with inclusion filters copied from the `_Default` sink. Select the new log bucket as the sink destination.
- B. Create exclusion filters for the `_Default` sink to prevent it from receiving new logs. Create a user-defined sink, and select the new log bucket as the sink destination.
- C. Disable the `_Default` sink. Create a user-defined sink and select the new log bucket as the sink destination.
- D. Edit the `_Default` sink, and select the new log bucket as the sink destination.

Answer: D

Explanation:

In Google Cloud's Logging service, log entries are automatically routed to the `_Default` log bucket unless configured otherwise. When you create a new log bucket and intend to redirect all log entries from the `_Default` bucket to this new bucket, the most efficient approach is to modify the existing `_Default` sink to point to the new log bucket.

Option A: Creating a new user-defined sink with filters replicated from the `_Default` sink is redundant and may lead to configuration complexities.

Option B: Implementing exclusion filters on the `_Default` sink and then creating a new sink introduces unnecessary steps and potential for misconfiguration.

Option C: Disabling the `_Default` sink would stop all log routing to it, but creating a new sink to replicate its functionality is inefficient.

Option D: Editing the `_Default` sink to change its destination to the new log bucket ensures a seamless transition of log routing without additional configurations.

Therefore, Option D is the most efficient and straightforward method to achieve the desired log routing.

Reference::

Routing and Storage Overview

[Configure Default Log Router Settings](#)

Question: 237

Your organization wants to publish yearly reports of your website usage analytics. You must ensure that no data with personally identifiable information (PII) is published by using the Cloud Data Loss Prevention (Cloud DLP) API. Data integrity must be preserved. What should you do?

- A. Encrypt the PII from the report by using the Cloud DLP API.
- B. Discover and transform PII data in your reports by using the Cloud DLP API.
- C. Detect all PII in storage by using the Cloud DLP API. Create a cloud function to delete the PII.
- D. Discover and quarantine your PII data in your storage by using the Cloud DLP API.

Answer: B

Explanation:

To ensure that no personally identifiable information (PII) is published in your yearly website usage analytics reports while preserving data integrity, the Cloud Data Loss Prevention (Cloud DLP) API can be utilized to identify and transform PII within your datasets.

Option A: Encrypting PII does not remove it from the reports; it merely obscures it, which may not be sufficient for compliance or privacy requirements.

Option B: Discovering and transforming PII ensures that sensitive information is either masked, tokenized, or otherwise obfuscated, effectively removing PII from the reports while maintaining the overall structure and utility of the data.

Option C: Detecting and deleting PII could lead to loss of valuable data and may disrupt the integrity of the reports.

Option D: Quarantining PII data implies isolating it, which doesn't address the need to publish reports without PII. Therefore, Option B is the most appropriate approach, as it leverages the Cloud DLP API to identify and transform PII, ensuring that the published reports are free from sensitive information while preserving data integrity.

Reference::

Cloud DLP Overview

De-identifying Sensitive Data

Question: 238

Your organization is building a real-time recommendation engine using ML models that process live user activity data stored in BigQuery and Cloud Storage. Each new model developed is saved to Artifact Registry. This new system deploys models to Google Kubernetes Engine and uses Pub/Sub for message queues. Recent industry news has been reporting attacks exploiting ML model supply chains. You need to enhance the security in this serverless architecture, specifically against risks to the development and deployment pipeline. What should you do?

- A. Limit external libraries and dependencies that are used for the ML models as much as possible. Continuously rotate encryption keys that are used to access the user data from BigQuery and Cloud Storage.
- B. Enable container image vulnerability scanning during development and pre-deployment. Enforce Binary Authorization on images deployed from Artifact Registry to your continuous integration and continuous deployment (CI/CD) pipeline.
- C. Thoroughly sanitize all training data prior to model development to reduce risk of poisoning attacks. Use IAM for authorization, and apply role-based restrictions to code repositories and cloud services.
- D. Develop strict firewall rules to limit external traffic to Cloud Run instances. Integrate intrusion detection systems (IDS) for real-time anomaly detection on Pub/Sub message flows.

Answer: B

Explanation:

To enhance the security of your machine learning (ML) model supply chain within a serverless architecture, it's crucial to implement measures that protect both the development and deployment pipelines.

Option A: While limiting external dependencies and rotating encryption keys are good security practices, they do not directly address the risks associated with the ML model supply chain. Option B: Implementing container image vulnerability scanning during development and predeployment helps identify and mitigate known vulnerabilities in your container images. Enforcing Binary Authorization ensures that only trusted and verified images are deployed in your environment. This combination directly strengthens the security of the ML model supply chain by validating the integrity of container images before deployment.

Option C: Sanitizing training data and applying role-based access controls are important security practices but do not specifically safeguard the deployment pipeline against compromised container images.

Option D: While strict firewall rules and intrusion detection systems enhance network security, they do not specifically address vulnerabilities within the container images or the deployment process. Therefore, Option B is the most effective approach, as it directly addresses the security of the development and deployment pipeline by ensuring that only vetted and secure container images are used in your environment.

Reference::

[Container Scanning Overview](#) [Binary Authorization Overview](#)

Question: 239

Your organization strives to be a market leader in software innovation. You provided a large number of Google Cloud environments so developers can test the integration of Gemini in Vertex AI into their existing applications or create new projects. Your organization has 200 developers and a five-person security team. You must prevent and detect proper security policies across the Google Cloud environments. What should you do? (Choose 2 answers)

- A. Apply a predefined AI-recommended security posture template for Gemini in Vertex AI in Security Command Center Enterprise or Premium tiers.
- B. Publish internal policies and clear guidelines to securely develop applications.
- C. Implement the least privileged access Identity and Access Management roles to prevent misconfigurations.
- D. Apply organization policy constraints. Detect and monitor drifts by using Security Health Analytics. E. Use Cloud Logging to create log filters to detect misconfigurations. Trigger Cloud Run functions to remediate misconfigurations.

Answer: C, D

Explanation:

To maintain proper security policies across numerous Google Cloud environments, especially with a large developer base and a small security team, it's crucial to implement automated and scalable security measures.

Option A: While applying AI-recommended security posture templates can be beneficial, as of now, there isn't a specific predefined template for Gemini in Vertex AI within the Security Command Center.

Option B: Publishing internal policies and guidelines is essential for promoting secure development practices but may not be sufficient alone to enforce or detect security policies.

Option C: Implementing the principle of least privilege through Identity and Access Management (IAM) roles minimizes the risk of misconfigurations and unauthorized access by ensuring users have only the permissions necessary for

their tasks.

Option D: Applying organization policy constraints enforces specific configurations and restrictions across projects.

Utilizing Security Health Analytics helps in detecting and monitoring deviations from these policies, providing automated insights into potential security issues.

Option E: Using Cloud Logging to detect misconfigurations and triggering Cloud Run functions for remediation introduces complexity and may require significant maintenance, making it less practical for a small security team.

Therefore, Options C and D are the most effective strategies. They provide automated enforcement and monitoring of security policies, aligning with the need for scalable solutions given the organization's size and resources.

Reference::

Identity and Access Management (IAM) Overview

Organization Policy Service Overview

Security Health Analytics Overview

Question: 240

Your organization must follow the Payment Card Industry Data Security Standard (PCI DSS). To prepare for an audit, you must detect deviations at an infrastructure-as-a-service level in your Google Cloud landing zone. What should you do?

- A. Create a data profile covering all payment-relevant data types. Configure Data Discovery and a risk analysis job in Google Cloud Sensitive Data Protection to analyze findings.
- B. Use the Google Cloud Compliance Reports Manager to download the latest version of the PCI DSS report. Analyze the report to detect deviations.
- C. Create an Assured Workloads folder in your Google Cloud organization. Migrate existing projects into the folder and monitor for deviations in the PCI DSS.
- D. Activate Security Command Center Premium. Use the Compliance Monitoring product to filter findings that may not be PCI DSS compliant.

Answer: D

Explanation:

To ensure compliance with the Payment Card Industry Data Security Standard (PCI DSS) at the infrastructure-as-a-service (IaaS) level within Google Cloud, it's essential to have continuous monitoring and assessment tools that can detect deviations from compliance requirements. Option A: Creating data profiles and configuring Data Discovery jobs in Google Cloud Sensitive Data Protection focuses on identifying and analyzing sensitive data but does not directly address infrastructure compliance monitoring.

Option B: Downloading the latest PCI DSS report from the Compliance Reports Manager provides a static compliance report but does not offer real-time detection of deviations within your specific environment.

Option C: Utilizing Assured Workloads helps in creating environments that meet specific compliance requirements, but migrating existing projects into such folders does not actively detect deviations; it primarily ensures that new workloads comply with predefined policies.

Option D: Activating Security Command Center (SCC) Premium and leveraging its Compliance Monitoring capabilities allows for continuous assessment of your Google Cloud environment against PCI DSS requirements. SCC can identify misconfigurations, vulnerabilities, and compliance violations in real-time, providing actionable insights to address any issues promptly.

Therefore, Option D is the most effective approach to detect deviations at the IaaS level in preparation for a PCI DSS audit.

Reference::

Security Command Center Overview

Security Command Center Compliance Monitoring

Question: 241

A security audit uncovered several inconsistencies in your project's Identity and Access Management (IAM) configuration. Some service accounts have overly permissive roles, and a few external collaborators have more access than necessary. You need to gain detailed visibility into changes to IAM policies, user activity, service account behavior, and access to sensitive projects. What should you do?

- A. Enable the metrics explorer in Cloud Monitoring to follow the service account authentication events and build alerts linked on it.
- B. Use Cloud Audit Logs. Create log export sinks to send these logs to a security information and event management (SIEM) solution for correlation with other event sources.
- C. Configure Google Cloud Functions to be triggered by changes to IAM policies. Analyze changes by using the policy simulator, send alerts upon risky modifications, and store event details.
- D. Deploy the OS Config Management agent to your VMs. Use OS Config Management to create patch management jobs and monitor system modifications.

Answer: B

Explanation:

To address inconsistencies in your project's Identity and Access Management (IAM) configuration and gain comprehensive visibility into IAM policy changes, user activity, service account behavior, and access to sensitive projects, leveraging Google Cloud's auditing capabilities is essential.

Option A: While Cloud Monitoring's metrics explorer can track certain metrics, it is not designed to provide detailed logs of IAM policy changes or user activities.

Option B: Cloud Audit Logs offer detailed records of administrative activities, including IAM policy changes and authentications. By creating log export sinks, you can forward these logs to a Security Information and Event Management (SIEM) solution, enabling correlation with other event sources and comprehensive analysis. This approach provides the necessary visibility into IAM configurations and user activities.

Option C: Triggering Cloud Functions based on IAM policy changes and analyzing them with a policy simulator is a proactive approach. However, it may not provide the depth of historical data and comprehensive analysis capabilities that a SIEM solution offers.

Option D: Deploying the OS Config Management agent focuses on VM configuration and patch management, which does not directly address IAM policy monitoring or user activity tracking. Therefore, Option B is the most effective solution to gain detailed visibility into IAM-related activities and address the identified inconsistencies.

Reference::

Cloud Audit Logs Overview

Exporting Logs to a SIEM

Question: 242

Your organization uses a microservices architecture based on Google Kubernetes Engine (GKE). Security reviews recommend tighter controls around deployed container images to reduce potential vulnerabilities and maintain compliance. You need to implement an automated system by using managed services to ensure that only

approved container images are deployed to the GKE clusters. What should you do?

- A. Enforce Binary Authorization in your GKE clusters. Integrate container image vulnerability scanning into the CI/CD pipeline and require vulnerability scan results to be used for Binary Authorization policy decisions.
- B. Develop custom organization policies that restrict GKE cluster deployments to container images hosted within a specific Artifact Registry project where your approved images reside.
- C. Build a system using third-party vulnerability databases and custom scripts to identify potential Common Vulnerabilities and Exposures (CVEs) in your container images. Prevent image deployment if the CVE impact score is beyond a specified threshold.
- D. Automatically deploy new container images upon successful CI/CD builds by using Cloud Build triggers. Set up firewall rules to limit and control access to instances to mitigate malware injection.

Answer: A

Explanation:

To enhance the security of your microservices architecture on Google Kubernetes Engine (GKE) and ensure that only approved container images are deployed, implementing Binary Authorization is a **robust solution**.

Option A: Enforcing Binary Authorization in your GKE clusters ensures that only container images that meet your organization's security policies are deployed. By integrating container image vulnerability scanning into your Continuous Integration/Continuous Deployment (CI/CD) pipeline, you can assess images for known vulnerabilities before they are deployed. Binary Authorization can be configured to use these vulnerability scan results to make policy decisions, effectively preventing the deployment of insecure images. This approach leverages managed services provided by Google Cloud, ensuring

scalability and compliance with security standards.

Option B: Developing custom organization policies to restrict deployments to images within a specific Artifact Registry project helps in controlling the source of images but does not inherently assess the security posture of those images. Without integrated vulnerability scanning and enforcement mechanisms, this approach may not fully mitigate the risk of deploying vulnerable images.

Option C: Building a system using third-party vulnerability databases and custom scripts requires significant maintenance and may not integrate seamlessly with GKE. This approach can be error-prone and lacks the efficiency of managed services designed for this purpose.

Option D: Automatically deploying new images upon successful CI/CD builds ensures rapid deployment but does not address the need for security assessments of the images. While setting up firewall rules is good practice, it does not prevent the deployment of potentially vulnerable images. Therefore, Option A is the most effective approach, as it utilizes Google Cloud's managed services to enforce security policies and integrate vulnerability assessments directly into the deployment process, ensuring that only approved and secure container images are deployed to your GKE clusters.

Reference::

[Binary Authorization Documentation](#)

[Container Analysis Documentation](#)

Question: 243

Your organization's application is being integrated with a partner application that requires read access to customer data to process customer orders. The customer data is stored in one of your Cloud Storage buckets. You have evaluated different options and determined that this activity requires the use of service account keys. You must advise the partner on how to minimize the risk of a compromised service account key causing a loss of data. What should you advise the partner to do?

- A. Define a VPC Service Controls perimeter, and restrict the Cloud Storage API. Add an ingress rule to the perimeter to

allow access to the Cloud Storage API for the service account from outside of the perimeter.

B. Scan the Cloud Storage bucket with Sensitive Data Protection when new data is added, and automatically mask all customer data.

C. Ensure that all data for the application that is accessed through the relevant service accounts is encrypted at rest by using customer-managed encryption keys (CMEK).

D. Implement a secret management service. Configure the service to frequently rotate the service account key.

Configure proper access control to the key, and restrict who can create service account keys.

Answer: D

Explanation:

When integrating applications that require access to sensitive data stored in Cloud Storage, managing service account keys securely is crucial to prevent unauthorized access or data loss. Option A: Defining a VPC Service Controls perimeter enhances security by restricting access to Google Cloud services. However, configuring ingress rules to allow external access for the service account may introduce complexities and potential security gaps, especially if the partner's infrastructure is outside the defined perimeter.

Option B: Scanning and masking customer data addresses data sensitivity but does not mitigate risks associated with compromised service account keys. This approach focuses on data content rather than access control mechanisms.

Option C: Encrypting data at rest using customer-managed encryption keys (CMEK) ensures data confidentiality but does not directly address the security of service account keys or access controls. Option D: Implementing a secret management service to handle service account keys is a best practice. By configuring the service to frequently rotate keys, you reduce the window of opportunity for malicious actors to exploit compromised keys. Additionally, enforcing strict access controls ensures that only authorized personnel can create or manage service account keys, minimizing the risk of unauthorized access. This approach directly addresses the security concerns related to service account key management.

Therefore, Option D is the most appropriate recommendation, as it focuses on securely managing service account keys through rotation and access controls, thereby minimizing the risk of data loss due to compromised keys.

Reference::

[Best Practices for Managing Service Account Keys](#)

Secret Manager Documentation

Question: 244

Your organization is using Vertex AI Workbench Instances. You must ensure that newly deployed instances are automatically kept up-to-date and that users cannot accidentally alter settings in the operating system. What should you do?

A. Enable the VM Manager and ensure the corresponding Google Compute Engine instances are added.

B. Enforce the disableRootAccess and requireAutoUpgradeSchedule organization policies for newly deployed instances.

C. Assign the AI Notebooks Runner and AI Notebooks Viewer roles to the users of the AI Workbench Instances.

D. Implement a firewall rule that prevents Secure Shell access to the corresponding Google Compute Engine instances by using tags.

Answer: B

Explanation:

To ensure that Vertex AI Workbench Instances are automatically kept up-to-date and that users cannot alter operating system settings, implementing specific organization policies is essential. Option A: Enabling VM Manager and adding Compute Engine instances assists in managing and monitoring VM instances but does not enforce automatic updates or restrict user modifications to the operating system.

Option B: Enforcing the disableRootAccess organization policy prevents users from gaining root access, thereby restricting unauthorized changes to the operating system. Additionally, the requireAutoUpgradeSchedule policy ensures that instances are automatically updated according to a defined schedule. Together, these policies maintain system integrity and compliance with update requirements.

Option C: Assigning AI Notebooks Runner and AI Notebooks Viewer roles controls user permissions related to running and viewing notebooks but does not directly influence operating system settings or update mechanisms.

Option D: Implementing firewall rules to prevent SSH access limits direct access to instances but does not ensure automatic updates or prevent alterations through other means.

Therefore, Option B is the most appropriate action, as it directly addresses both the enforcement of automatic updates and the prevention of unauthorized operating system modifications.

Reference::

Organization Policy Constraints

VM Manager Overview

Question: 245

You are implementing a new web application on Google Cloud that will be accessed from your onpremises network. To provide protection from threats like malware, you must implement transport layer security (TLS) interception for incoming traffic to your application. What should you do?

- A. Configure Secure Web Proxy. Offload the TLS traffic in the load balancer, inspect the traffic, and forward the traffic to the web application.
- B. Configure an internal proxy load balancer. Offload the TLS traffic in the load balancer, inspect the traffic, and forward the traffic to the web application.
- C. Configure a hierarchical firewall policy. Enable TLS interception by using Cloud Next Generation Firewall (NGFW) Enterprise.
- D. Configure a VPC firewall rule. Enable TLS interception by using Cloud Next Generation Firewall (NGFW) Enterprise.

Answer: A

Explanation:

To protect your web application from threats like malware by implementing TLS interception for incoming traffic, configuring a Secure Web Proxy with TLS offloading at the load balancer is an effective approach.

Option A: By configuring a Secure Web Proxy, you can offload TLS traffic at the load balancer, inspect the decrypted traffic for threats such as malware, and then forward the inspected traffic to your web application. This approach ensures that encrypted traffic is securely analyzed without compromising the security of the data in transit.

Option B: An internal proxy load balancer is designed for distributing traffic within a private network and may not support

TLS interception capabilities required for inspecting incoming traffic from external sources.

Option C: Hierarchical firewall policies in Google Cloud are used to enforce security rules across your organization but do not provide TLS interception capabilities.

Option D: VPC firewall rules control traffic to and from VM instances based on specified rules but do not have the capability to perform TLS interception or traffic inspection.

Therefore, Option A is the most suitable solution, as it allows for TLS interception through a Secure Web Proxy, enabling the inspection of incoming encrypted traffic to detect and mitigate threats like malware before the traffic reaches your web application.

Reference::

Secure Web Proxy Overview

Cloud Load Balancing Overview

Question: 246

You work for a healthcare provider that is expanding into the cloud to store and process sensitive patient data.

a. You must ensure the chosen Google Cloud configuration meets these strict regulatory requirements:

Data must reside within specific geographic regions.

Certain administrative actions on patient data require explicit approval from designated compliance officers.

Access to patient data must be auditable.

What should you do?

A. Select multiple standard Google Cloud regions for high availability. Implement Access Control Lists (ACLs) on individual storage objects containing patient data. Enable Cloud Audit Logs.

B. Deploy an Assured Workloads environment in multiple regions for redundancy. Utilize custom IAM roles with granular permissions. Isolate network-level data by using VPC Service Controls.

C. Deploy an Assured Workloads environment in an approved region. Configure Access Approval for sensitive operations on patient data. Enable both Cloud Audit Logs and Access Transparency.

D. Select a standard Google Cloud region. Restrict access to patient data based on user location and job function by using Access Context Manager. Enable both Cloud Audit Logging and Access Transparency.

Answer: C

Explanation:

To ensure compliance with strict regulatory requirements for storing and processing sensitive patient data in the cloud, the following measures should be implemented:

Assured Workloads: Deploying an Assured Workloads environment in an approved region ensures that data residency requirements are met by restricting data storage and processing to specific geographic locations. Assured Workloads provide predefined controls and configurations tailored to meet regulatory compliance needs.

Access Approval: Configuring Access Approval ensures that certain administrative actions on patient data require explicit approval from designated compliance officers. This adds a layer of control over sensitive operations, aligning with the need for explicit approvals.

Cloud Audit Logs and Access Transparency: Enabling Cloud Audit Logs provides a detailed record of actions taken on your data, supporting the requirement for auditability. Access Transparency logs offer visibility into Google's administrative access to your content, enhancing transparency and compliance.

Therefore, Option C is the most appropriate choice, as it comprehensively addresses data residency, administrative control, and auditability requirements.

Reference::

Assured Workloads Overview

Access Approval Documentation

Cloud Audit Logs Overview

Access Transparency Overview

Question: 247

Your organization leverages folders to represent different teams within your Google Cloud environment. To support Infrastructure as Code (IaC) practices, each team receives a dedicated service account upon onboarding. You want to ensure that teams have comprehensive permissions to manage resources within their assigned folders while adhering to the principle of least privilege. You must design the permissions for these team-based service accounts in the most effective way possible. What should you do?

- A. Grant each service account the folder administrator role on its respective folder.
- B. Grant each service account the project creator role at the organization level and use folder-level IAM conditions to restrict project creation to specific folders. [Reddit](#)
- C. Assign each service account the project editor role at the organization level and instruct teams to use IAM bindings at the folder level for fine-grained permissions.
- D. Assign each service account the folder IAM administrator role on its respective folder to allow teams to create and manage additional custom roles if needed.

Answer: A

Explanation:

To ensure that each team's service account has the necessary permissions to manage resources within their assigned folders while adhering to the principle of least privilege, the following considerations apply:

Folder Administrator Role: Granting each service account the Folder Administrator role on its respective folder provides comprehensive permissions to manage resources within that folder, including creating, updating, and deleting projects and resources. This approach ensures that teams have the necessary control over their environments without extending permissions beyond their assigned scope.

Principle of Least Privilege: By assigning permissions at the folder level, you limit the service account's access to only the resources within its designated folder, aligning with the principle of least privilege and reducing the risk of unauthorized access to other parts of the organization.

Therefore, Option A is the most effective approach, as it provides the necessary permissions for teams to manage their resources within their assigned folders while adhering to security best practices.

Reference::

Understanding Roles

[Best Practices for Enterprise Organizations](#)

Question: 248

You must ensure that the keys used for at-rest encryption of your data are compliant with your organization's security controls. One security control mandates that keys get rotated every 90 days. You must implement an effective detection strategy to validate if keys are rotated as required. What should you do?

- A. Analyze the crypto key versions of the keys by using data from Cloud Asset Inventory. If an active key is older than 90 days, send an alert message through your incident notification channel.
- B. Identify keys that have not been rotated by using Security Health Analytics. If a key is not rotated after 90 days, a finding in Security Command Center is raised.

- C. Assess the keys in the Cloud Key Management Service by implementing code in Cloud Run. If a key is not rotated after 90 days, raise a finding in Security Command Center.
- D. Define a metric that checks for timely key updates by using Cloud Logging. If a key is not rotated after 90 days, send an alert message through your incident notification channel.

Answer: A

Explanation:

Question: 249

Your organization deploys a large number of containerized applications on Google Kubernetes Engine (GKE). Node updates are currently applied manually. Audit findings show that a critical patch has not been installed due to a missed notification. You need to design a more reliable, cloud-first, and scalable process for node updates. What should you do?

- A. Migrate the cluster infrastructure to a self-managed Kubernetes environment for greater control over the patching process.
- B. Develop a custom script to continuously check for patch availability, download patches, and apply the patches across all components of the cluster.
- C. Schedule a daily reboot for all nodes to automatically upgrade.
- D. Configure node auto-upgrades for node pools in the maintenance windows.

Answer: D

Explanation:

To establish a reliable, cloud-native, and scalable process for updating nodes in your GKE clusters, configuring node auto-upgrades within designated maintenance windows is the most effective approach.

Option A: Migrating to a self-managed Kubernetes environment would increase operational overhead and complexity, as your team would be responsible for managing the entire infrastructure, including patching and updates. This contradicts the goal of adopting a cloud-first strategy and does not inherently provide a more reliable update process.

Option B: Developing custom scripts for patch management introduces potential risks and maintenance burdens. Ensuring the reliability, security, and scalability of such scripts can be challenging, and this approach may not align with best practices for managing GKE environments. Option C: Scheduling daily reboots does not guarantee that nodes will apply the latest patches or updates. Without a mechanism to manage and apply updates, reboots alone are insufficient to maintain node security and compliance.

Option D: Configuring node auto-upgrades ensures that GKE automatically keeps your nodes up-to-date with the latest stable versions, reducing the risk of missed critical patches. By setting maintenance windows, you can control when these upgrades occur, minimizing disruptions to your workloads. This approach leverages GKE's managed services to maintain security and compliance efficiently.

Therefore, Option D is the optimal solution, as it aligns with a cloud-first strategy and leverages GKE's native capabilities to automate and schedule node updates effectively.

Reference::

[Auto-upgrading nodes | Google Kubernetes Engine \(GKE\) Maintenance windows and exclusions | Google Kubernetes Engine](#)

Question: 250

Your company's detection and response team requires break-glass access to the Google Cloud organization in the event of a security investigation. At the end of each day, all security group membership is removed. You need to automate user provisioning to a Cloud Identity security group. You have created a service account to provision group memberships. Your solution must follow Google-recommended practices and comply with the principle of least privilege. What should you do?

- A. In Google Workspace, grant the service account client ID access to the scope, <https://www.googleapis.com/auth/admindirectorygroup>, by using domain-wide delegation, and use a service account key.
- B. In Google Workspace, grant the service account client ID access to the scope, <https://www.googleapis.com/auth/admindirectorygroup>, by using domain-wide delegation. Use Application Default Credentials with the resource-attached service account.
- C. In Google Workspace, grant the Groups Editor role to the service account. Enable the Cloud Identity API. Use a service account key.
- D. In Google Workspace, grant the Groups Editor role to the service account, enable the Cloud Identity API, and use Application Default Credentials with the resource-attached service account.

Answer: B

Explanation:

The problem requires automating user provisioning to a Cloud Identity security group using a service account, adhering to Google-recommended practices and the principle of least privilege.

Cloud Identity Groups and Google Workspace: Cloud Identity groups are managed as part of Google Workspace. To programmatically manage Google Workspace resources (like groups), you typically use the Admin SDK APIs.

Domain-Wide Delegation: Service accounts cannot directly authenticate to Google Workspace APIs using IAM roles.

Instead, they require "domain-wide delegation" to impersonate a user with the necessary administrative privileges within Google Workspace. This allows a service account to access user data or perform administrative tasks across the domain.

The correct scope for managing groups is <https://www.googleapis.com/auth/admindirectorygroup>. Extract Reference: "To allow a service account to access user data on behalf of users in a Google Workspace domain, you must delegate domain-wide authority to your service account" (Google Cloud documentation: <https://developers.google.com/identity/protocols/oauth2/service-account#delegating>).

Extract Reference (Admin SDK Scopes): The <https://www.googleapis.com/auth/admindirectorygroup> scope is explicitly listed for "View and manage all groups on the domain" (Google Workspace Admin SDK documentation: <https://developers.google.com/admin-sdk/directory/v1/scopes>).

Application Default Credentials (ADC) with Resource-Attached Service Account: Google-

recommended practices strongly advise against using service account keys directly for authentication when running on Google Cloud infrastructure. Instead, it's recommended to use Application Default Credentials (ADC) with a service account attached to the resource (eg, a Compute Engine VM, Cloud Run service, or Cloud Functions). This method manages credentials automatically and securely, reducing the risk associated with managing and rotating keys. Extract Reference: "For most Google Cloud services, Application Default Credentials (ADC) is the recommended way to authenticate" and "When running code in a Google Cloud environment, such as Compute Engine, Cloud Run, or Cloud Functions, use the built-in service account to authenticate automatically with ADC. This is the most secure approach, as

you don't need to manually create or manage service account keys" (Google Cloud documentation:

<https://cloud.google.com/docs/authentication/production>)

Options C and D are incorrect because granting an IAM role like "Groups Editor" in Google Cloud does not enable a service account to manage Google Workspace (Cloud Identity) group memberships; domain-wide delegation is required for that Option A uses a service account key, which is less secure than ADC with a resource-attached service account according to Google's recommendations

Therefore, option B is the most aligned with Google's recommended practices for securely automating group provisioning using a service account and domain-wide delegation

Question: 251

Your financial services company has an audit requirement under a strict regulatory framework that requires comprehensive, immutable audit trails for all administrative and data access activity that ensures that data is kept for seven years Your current logging is fragmented across individual projects You need to establish a centralized, tamper-proof, long-term logging solution accessible for audits What should you do?

- A. Implement Pub/Sub to stream all audit logs from each project in real-time to an external Security Information and Event Management (SIEM) for long-term analysis
- B. Establish organization-level Cloud Logging sinks to export Cloud Audit Logs to a dedicated Cloud Storage bucket with object retention lock
- C. Enable Security Command Center across the organization to gain centralized visibility into threats and manage compliance posture for all Google Cloud projects
- D. Individually configure Cloud Audit Logs for all Google Cloud services in each project Store the logs in regional Cloud Logging buckets with 30-day retention policies

Answer: B

Explanation:

The core requirements are: centralized, tamper-proof, long-term (seven years) immutable audit trails for administrative and data access activity

Centralization: The current logging is fragmented To centralize, you need to collect logs from across the organization Cloud Logging sinks configured at the organization level are designed for this purpose They allow you to route logs from all projects within an organization to a single destination Extract Reference: "Aggregated exports allow you to export logs from multiple Google Cloud projects, folders, or your entire organization An aggregated export can include all logs from all

included resources, or you can use queries to include only specific logs" (Google Cloud documentation:

https://cloud.google.com/logging/docs/export/aggregated_exports) Long-Term Storage (Seven Years): Cloud Logging buckets have default retention periods (eg, 30 days for Data Access logs, 400 days for Admin Activity logs) which are not sufficient for a seven-year requirement Cloud Storage is ideal for long-term archival Extract Reference: "Cloud Storage is a highly scalable and durable object storage service suitable for archiving large volumes of data for extended periods" (Google Cloud documentation, general overview of Cloud Storage features)

Tamper-Proof / Immutability: This is a critical requirement for audit trails in financial services under strict regulatory frameworks Cloud Storage's "object retention lock" feature provides immutability Once an object retention lock is set on a bucket, objects within that bucket cannot be deleted or overwritten for a specified duration, ensuring data integrity for

compliance purposes Extract Reference: "Object Retention Lock helps you meet compliance requirements by preventing data from being deleted or modified for a fixed amount of time or indefinitely This feature satisfies SEC Rule 17a-4(f), FINRA Rule 4511(c), and CFTC Regulation 131(c)-(d) requirements" (Google Cloud documentation: <https://cloud.google.com/storage/docs/bucket-lock>)

Let's evaluate the other options:

A Implement Pub/Sub to stream all audit logs from each project in real-time to an external SIEM: While Pub/Sub can centralize real-time streaming to a SIEM, the solution described does not inherently guarantee tamper-proof storage or 7-year immutability within Google Cloud The SIEM itself would need to provide those capabilities, which is outside the scope of Google Cloud's direct offering for this specific requirement

C Enable Security Command Center across the organization: Security Command Center (SCC) provides centralized visibility into security posture, threats, and compliance However, SCC is a security management and monitoring platform; it does not serve as the primary long-term, immutable storage for raw audit logs It consumes information from logs but doesn't store them in a way that meets the 7-year immutable archival requirement

D Individually configure Cloud Audit Logs for all Google Cloud services in each project Store the logs in regional Cloud Logging buckets with 30-day retention policies: This fails on multiple counts: it's not centralized (requires individual configuration), and the 30-day retention in Cloud Logging buckets is far short of the seven-year requirement It also doesn't explicitly guarantee tamper-proof storage beyond the default logging immutability

Therefore, option B directly addresses all aspects of the requirement: centralization via organization-level sinks, long-term storage with Cloud Storage, and immutability/tamper-proofing with object retention lock

Question: 252

Your organization has Google Cloud applications that require access to external web services You must monitor, control, and log access to these services What should you do?

- A. Configure VPC firewall rules to allow the services to access the IP addresses of required external web services
- B. Set up a Secure Web Proxy that allows access to the specific external web services Configure applications to use the proxy for the web service requests
- C. Configure Google Cloud Armor to monitor and protect your applications by checking incoming traffic patterns for attack patterns
- D. Set up a Cloud NAT instance to allow egress traffic from your VPC

Answer: B

Explanation:

The problem states that Google Cloud applications need to access external web services and requires the ability to monitor, control, and log this access

Monitoring, Controlling, and Logging external web access: This specifically points to a proxy solution, which can intercept, inspect, and log HTTP/S traffic

Secure Web Proxy (SWP): Google Cloud's Secure Web Proxy is designed for exactly this use case It acts as an explicit forward proxy for HTTP(S) traffic, allowing organizations to implement granular access controls, inspect traffic for security threats, and log all outbound web requests from their Google Cloud environment Extract Reference: "Secure Web Proxy is a managed service that lets you deploy and manage an explicit forward proxy to protect your organization's internal resources from web-based threats and to control access to external web applications" and "With Secure Web Proxy, you can: Enforce granular access policies based on different attributes, Log all HTTP(S) requests that are handled by the proxy, and Monitor web traffic for threats" (Google Cloud documentation: <https://cloud.google.com/secure-web-proxy>)

Let's evaluate the other options:

A Configure VPC firewall rules to allow the services to access the IP addresses of required external web services: VPC firewall rules operate at Layer 4 (TCP/UDP) and Layer 3 (IP) While they can allow or deny traffic to specific IP addresses and ports, they cannot monitor, control, or log HTTP/S requests at the application layer They don't provide granular control over which web services are accessed or inspect the content of the requests

C Configure Google Cloud Armor to monitor and protect your applications by checking incoming traffic patterns for attack patterns: Google Cloud Armor is primarily a Distributed Denial of Service (DDoS) protection and Web Application Firewall (WAF) service It focuses on protecting applications from incoming threats (ingress traffic), not controlling and logging outgoing access to external web services

D Set up a Cloud NAT instance to allow egress traffic from your VPC: Cloud NAT allows instances without external IP addresses to connect to the internet While it enables egress, it does not provide monitoring, control, or logging capabilities for specific web services at the application layer It's a network address translation service, not an application-layer proxy

Therefore, setting up a Secure Web Proxy is the most appropriate solution to meet the requirements of monitoring, controlling, and logging access to external web services from Google Cloud applications

Question: 253

Your company is developing a new application for your organization The application consists of two Cloud Run services, service A and service B Service A provides a web-based user front-end Service B provides back-end services that are called by service A You need to set up identity and access management for the application Your solution should follow the principle of least privilege What should you do?

A. Create a new service account with the permissions to run service A and service B Require

authentication for service B Permit only the new service account to call the backend

B. Create two separate service accounts Grant one service account the permissions to execute service A, and grant the other service account the permissions to execute service B Require authentication for service B Permit only the service account for service A to call the back-end C. Use the Compute Engine default service account to run service A and service B Require authentication for service B Permit only the default service account to call the backend

D. Create three separate service accounts Grant one service account the permissions to execute service A Grant the second service account the permissions to run service B Grant the third service account the permissions to communicate between both services A and B Require authentication for service B Call the back-end by authenticating with a service account key for the third service account

Answer: B

Explanation:

The problem describes an application with two Cloud Run services (Service A - frontend, Service B - backend) and requires setting up IAM with the principle of least privilege Service A calls Service B

Principle of Least Privilege: This principle dictates that each entity (in this case, a Cloud Run service) should only have the minimum permissions necessary to perform its function

Separate Service Accounts for Separate Services: To adhere to the principle of least privilege, it's best practice to assign a unique service account to each distinct service or component This ensures that a compromise of one service account

does not grant excessive permissions across other services Service A needs permissions to run itself and to invoke Service B Service B only needs permissions to run itself Extract Reference: "Assign a service account to a Cloud Run service The service account acts as the identity for your service and determines what permissions your revisions have when executing requests It is a best practice to grant each service account only the permissions that are required to run the specific service (principle of least privilege)" (Google Cloud documentation: <https://cloud.google.com/run/docs/configuring/service-accounts>)

Authentication for Cloud Run Services: When one Cloud Run service (caller) needs to invoke another Cloud Run service (callee), the caller must be authorized to do so This is typically achieved by assigning the roles/runinvoker role on the callee service to the caller's service account Extract Reference: "To allow a service to invoke another service, grant the roles/runinvoker role on the called service to the caller's service account" (Google Cloud documentation: <https://cloud.google.com/run/docs/securing/service-to-service>)

Let's evaluate the options:

A Create a new service account with the permissions to run service A and service B Require authentication for service B Permit only the new service account to call the backend This violates the principle of least privilege by giving a single service account permissions for both services If that service account were compromised, both services would be affected

B Create two separate service accounts Grant one service account the permissions to execute service A, and grant the other service account the permissions to execute service B Require authentication for service B Permit only the service account for service A to call the back-end This aligns perfectly with least privilege Service A gets its own identity, Service B gets its own identity Service A's service account is then granted runinvoker permissions on Service B, allowing it to call the backend while Service B requires authentication This is the recommended approach

C Use the Compute Engine default service account to run service A and service B Require authentication for service B Permit only the default service account to call the backend The Compute Engine default service account often has broad permissions (eg, editor role in its project) Using it violates the principle of least privilege and is generally discouraged for production applications due to the potential for excessive permissions

D Create three separate service accounts Grant one service account the permissions to execute service A Grant the second service account the permissions to run service B Grant the third service account the permissions to communicate between both services A and B Require authentication for service B Call the back-end by authenticating with a service account key for the third service account This introduces unnecessary complexity with a third service account just for communication More critically, using a service account key for authentication is generally discouraged in Cloud Run environments where ADC (Application Default Credentials) can be used, as managing keys securely becomes an operational overhead and security risk Cloud Run services automatically use their attached service accounts for authentication when making calls to other Google Cloud services, including other Cloud Run services Therefore, option B is the best solution, adhering to the principle of least privilege and Google Cloud best practices for Cloud Run service-to-service authentication

Question: 254

location and to deploy different types of models in a consistent way You must ensure that your users can only access the approved models What should you do?

- A. Configure IAM permissions on individual Model Garden to restrict access to specific models
- B. Regularly audit user activity logs in Vertex AI to identify and revoke access to unapproved models
- C. Train custom models within your Vertex AI project and restrict user access to these models
- D. Implement an organization policy that restricts the vertexai.allowedModels constraint

Answer: D

Explanation:

The problem states that the organization is using Model Garden and needs to ensure users can only access approved models. This implies a need for a central, enforceable control mechanism. Organization Policies and Constraints: Google Cloud Organization Policy Service allows administrators to centrally control resources across an organization. Constraints are specific types of restrictions that can be applied. For AI Platform (which includes Vertex AI and Model Garden), there are specific constraints designed to control model usage. `vertexai.allowedModels` Constraint: This specific organization policy constraint is designed precisely to restrict which models can be used within a given organization, folder, or project. It provides a centralized way to define a list of approved models that users are allowed to access. Extract Reference: "The `vertexai.allowedModels` constraint allows you to specify a list of model URIs that are allowed to be used within the resource hierarchy" and "This constraint helps organizations enforce compliance and control which models are consumed by their users" (Google Cloud documentation, typically found under Organization Policy Service constraints for Vertex AI or AI Platform). Let's evaluate the other options:

A. Configure IAM permissions on individual Model Garden to restrict access to specific models: IAM

(Identity and Access Management) typically grants permissions at a broader resource level (eg, project, dataset, model resource). While you can control who can manage models, directly restricting access to specific models within Model Garden for consumption via IAM roles on individual models is not the primary mechanism for enforcing a list of approved models across an organization in a preventative way. Organization policies are designed for this kind of broad, preventative control.

B. Regularly audit user activity logs in Vertex AI to identify and revoke access to unapproved models: Auditing logs is a reactive measure. While important for monitoring and detecting violations, it does not prevent users from accessing unapproved models in the first place. The requirement is to ensure they can only access approved models, implying a proactive control.

C. Train custom models within your Vertex AI project and restrict user access to these models: This is about managing access to custom-trained models, not about controlling access to the collection of models in Model Garden, which often includes pre-trained or publicly available models that need to be whitelisted. It doesn't address the requirement of ensuring users only access approved models from the broader Model Garden collection.

Therefore, implementing an organization policy with the `vertexai.allowedModels` constraint is the most effective and Google-recommended way to centrally ensure that users can only access approved models within an organization using Model Garden.

Question: 255

A centralized security service has been implemented by your company. All applications running in Google Cloud are required to send data to this service. You need to ensure that developers have high autonomy to configure firewall rules within their projects, while preventing accidental blockage of access to the central security service. What should you do?

A. Deploy a central Secure Web Proxy and connect it to all VPC networks. Create a Secure Web Proxy policy to allow traffic to the central security service.

B. Implement a hierarchical firewall policy that prioritizes the central security service by allowing its connections and directing all other traffic to the subsequent firewall level.

C. Create a central project to manage Shared VPC networks which will be accessible to all other projects. Administer all firewall rules centrally within this project.

D. Use Terraform to automate the creation of the required firewall rule in all projects Restrict rule change permissions solely to the Terraform service account

Answer: B

Explanation:

The problem has two key requirements:

All applications must send data to a centralized security service

Developers need high autonomy over firewall rules within their projects Prevent accidental blockage of access to the central security service This scenario requires a mechanism to enforce critical network policies at a higher level of the resource hierarchy while still allowing project-level flexibility

Hierarchical Firewall Policies: Google Cloud's Hierarchical Firewall Policies (HFP) are designed precisely for this purpose They allow administrators to define firewall rules at the organization or folder level, and these rules are inherited by all projects and VPC networks within that hierarchy Crucially, HFP rules can be prioritized Rules with higher priority (lower numerical value) are evaluated first This means you can create high-priority "allow" rules for critical services that cannot be overridden or blocked by project-level firewall rules Extract Reference: "Hierarchical firewall policies allow you to define and enforce consistent network security policies across your organization Policies can be applied at the organization or folder level, and they are inherited by all projects and VPC networks within that hierarchy" and "Rules in a hierarchical firewall policy can take precedence over VPC network firewall rules based on priority A rule with a lower priority value takes precedence over a rule with a higher priority value" (Google Cloud documentation: <https://cloud.google.com/vpc/docs/firewall-policies-overview>)

Preventing Accidental Blockage while Allowing Autonomy: By setting a high-priority "allow" rule for the central security service in a hierarchical firewall policy, you guarantee that this traffic will always be permitted, regardless of what project-level firewall rules developers might configure This ensures the critical connectivity while still allowing developers to manage other, less critical firewall rules within their projects with high autonomy

Let's evaluate the other options:

A Deploy a central Secure Web Proxy and connect it to all VPC networks Create a Secure Web Proxy policy to allow traffic to the central security service A Secure Web Proxy is for HTTP/S outbound traffic to external web services The central security service might not be an external web service, and this solution is focused on application-layer proxies, not general network connectivity like sending data to an internal service Also, it doesn't directly address the challenge of developers blocking access with project-level firewall rules

C Create a central project to manage Shared VPC networks which will be accessible to all other projects Administer all firewall rules centrally within this project While Shared VPC centralizes network management, it means all firewall rules are administered centrally This directly contradicts the requirement for developers to have "high autonomy to configure firewall rules within their projects" Shared VPC would centralize too much control for this specific scenario

D Use Terraform to automate the creation of the required firewall rule in all projects Restrict rule change permissions solely to the Terraform service account This approach automates the creation but doesn't prevent developers from creating conflicting or overriding rules in their projects (unless Terraform is used to manage all rules, again removing autonomy) It also relies on restricting IAM permissions for all firewall rules, which is against the "high autonomy" requirement for developers Hierarchical firewall policies offer a more robust and native solution for overriding and enforcing specific rules

Therefore, implementing a hierarchical firewall policy is the most effective solution, as it allows for the enforcement of critical security service connectivity at a higher level, while still granting developers the desired autonomy over their project-specific firewall rules

Question: 256

You are setting up Cloud Identity for your company's Google Cloud organization. User accounts will be provisioned from Microsoft Entra ID through Directory Sync, and there will be single sign-on through Entra ID. You need to secure the super administrator accounts for the organization. Your solution must follow the principle of least privilege and implement strong authentication. What should you do?

- A. Create dedicated accounts for super administrators. Ensure that 2-step verification is enforced for the super administrator accounts in Entra ID.
- B. Create dedicated accounts for super administrators. Enforce Google 2-step verification for the super administrator accounts.
- C. Create accounts that combine the organization administrator and the super administrator privileges. Ensure that 2-step verification is enforced for the super administrator accounts in Entra ID.
- D. Create accounts that combine the organization administrators and the super administrator privileges. Enforce Google 2-step verification for the super administrator accounts.

Answer: B

Explanation:

The problem focuses on securing "super administrator accounts for the organization" when Cloud Identity is synced with Microsoft Entra ID and uses Entra ID for SSO. The key requirements are the principle of least privilege and strong authentication.

Principle of Least Privilege & Dedicated Accounts: Google's best practices strongly recommend creating dedicated, non-federated accounts for super administrators that are distinct from regular user accounts. These accounts should only be used for super administrator tasks and not for daily activities. This segregation ensures that the highest privilege accounts are isolated and adhere to the principle of least privilege by not having combined responsibilities.

Extract Reference: "Designate Organization Administrators We recommend keeping your super admin account separate from your Organization Administrator group" and "Give super admins a separate account that requires a separate login. For example, user `alice@example.com` could have a super admin account `alice-admin@example.com`" and "Use the super admin account only when needed. Delegate administrator tasks to user accounts with limited admin roles. Use the least privilege approach" (Google Cloud Documentation: "Super administrator account best practices | Resource Manager Documentation" - <https://cloud.google.com/resource-manager/docs/super-admin-best-practices>)

Strong Authentication (Google 2-Step Verification): Even when using a third-party identity provider like Microsoft Entra ID for most users, Google recommends enforcing Google's own 2-Step Verification for the critical super administrator accounts. This provides a "break-glass" mechanism that is independent of the external IdP. If the Entra ID integration were to fail or become compromised, the Google-managed super administrator accounts, protected by Google's own 2SV, would still be accessible for emergency recovery.

Extract Reference: "Even when using the legacy SSO profile, super admins can't sign in with SSO in these cases: Admin console. When super administrators try to sign in to an SSO-enabled domain via `admin.google.com`, they must enter their full Google administrator account email address and associated Google password (not their SSO username and password), and click Sign in to directly access the Admin console. Google doesn't redirect them to the SSO sign-in page" (Google Cloud Identity Help: "Super administrator SSO" -

<https://support.google.com/cloudidentity/answer/6341409>) - This highlights that super admin accounts can bypass SSO for direct Admin console access, making Google's 2SV crucial.

Extract Reference: "It's especially important for super admins to use 2SV because their accounts control access to all business and employee data in the organization. Protect your business with 2-Step Verification. Use security keys for 2-Step Verification" (Cloud Identity Help: "Security best practices for administrator accounts" -

<https://support.google.com/cloudidentity/answer/9011373>) Options C and D are incorrect because combining "organization administrator" (IAM role for GCP resources) and "super administrator" (Google Workspace/Cloud Identity domain-level control) privileges violates the principle of least privilege. Option A is less secure than B because relying solely on Entra ID's 2SV for super administrators means a compromise of Entra ID or an outage would leave the Google Cloud organization vulnerable without an independent break-glass mechanism.

Question: 257

Your company has deployed an artificial intelligence model in a central project. As this model has a lot of sensitive intellectual property and must be kept strictly isolated from the internet, you must expose the model endpoint only to a defined list of projects in your organization. What should you do?

- A. Create a central project to host Shared VPC networks that are provided to all other projects. Centrally administer all firewall rules in this project to grant access to the model.
- B. Activate Private Google Access in both the model project as well as in each project that needs to connect to the model. Create a firewall policy to allow connectivity to Private Google Access addresses.
- C. Within the model project, create an internal Application Load Balancer that points to the model endpoint. Expose this

load balancer with Private Service Connect to a configured list of projects D. Within the model project, create an external Application Load Balancer that points to the model endpoint Create a Cloud Armor policy to restrict IP addresses to Google Cloud

Answer: C

Explanation:

The problem requires exposing a sensitive AI model endpoint internally (strictly isolated from the internet) to a defined list of projects within the organization

Internal Exposure and Isolation: An "internal Application Load Balancer" is suitable for exposing services within your VPC network, ensuring they are not accessible from the internet

Private Service Connect (PSC): This is the key technology for securely and privately exposing services from one VPC network (the service producer, where the model is) to other VPC networks (the service consumers, the defined list of projects) within the same or different organizations PSC allows consumers to access services using internal IP addresses, with traffic remaining on Google's private network You can configure a service attachment that points to the internal load balancer, and then permit specific consumer projects to connect to this service attachment

Extract Reference: "Private Service Connect is a capability of Google Cloud networking that allows consumers to access managed services privately from inside their VPC network Similarly, it allows managed service producers to host these services in their own separate VPC networks and offer a private connection to their consumers" (Google Cloud Documentation: "Private Service Connect | VPC" - <https://cloud.google.com/vpc/docs/private-service-connect>)

Extract Reference: "Private Service Connect endpoints are internal IP addresses in a consumer VPC network that can be directly accessed by clients in that network Endpoints are created by deploying a forwarding rule that references a service attachment or a bundle of Google APIs" (Google Cloud Documentation: "About Private Service Connect | VPC" - <https://cloud.google.com/vpc/docs/private-service-connect>)

Extract Reference: "Private Service Connect can be used to access managed services that are owned by Google, third-party software as a service (SaaS) companies, or other teams within the consumer's own company Both published services and Google APIs can be targets of Private Service Connect" (Google Cloud Documentation: "About Private Service Connect | VPC" - <https://cloud.google.com/vpc/docs/private-service-connect>)

Let's evaluate the other options:

A Shared VPC and central firewall rules: While Shared VPC centralizes network management, it does not provide a direct managed service exposure mechanism like PSC for a model endpoint to specific projects It's more about sharing subnets and network resources Administering all firewall rules

centrally would also not meet the need for exposing only this specific model to a defined list of projects in a managed, private service pattern

B Activate Private Google Access (PGA): Private Google Access allows VMs without external IP addresses to access Google APIs and services (like Cloud Storage, BigQuery, etc) privately from within their VPC network It's for consuming Google services, not for exposing custom services hosted in a Google Cloud project to other projects

D External Application Load Balancer + Cloud Armor: An "external Application Load Balancer" exposes the service to the internet While Cloud Armor can restrict access based on IP addresses, it still involves internet exposure, which contradicts the "strictly isolated from the internet" requirement Restricting to "Google Cloud IP addresses" doesn't guarantee access only to a defined list of projects and still exposes the service externally

Therefore, creating an internal Application Load Balancer and exposing it via Private Service Connect is the most suitable and secure solution for this scenario

Question: 258

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:

The goal is to detect cryptocurrency mining software using Security Command Center (SCC). Security Command Center Threat Detection Services: SCC Premium and Enterprise tiers offer various specialized threat detection services. Virtual Machine Threat Detection (VMTD): This service is explicitly designed to scan virtual machines (Compute Engine instances and GKE nodes) for specific threats, including cryptocurrency mining software. It operates at the hypervisor level, performing deep scans of VM memory and disks. Extract Reference: "Virtual Machine Threat Detection (VMTD) helps you detect potential threats, such as cryptocurrency mining and malware, within your Compute Engine instances and GKE nodes" (Google Cloud Documentation: "Virtual Machine Threat Detection overview | Security Command Center" - <https://cloud.google.com/security-command-center/docs/concepts-vm-threat-detection-overview>) Extract Reference: "This service scans virtual machines to detect potentially malicious applications, such as cryptocurrency mining software, kernel-mode rootkits, and malware running in compromised cloud environments" (Google Cloud Documentation: "Virtual Machine Threat Detection overview | Security Command Center" - <https://cloud.google.com/security-command-center/docs/concepts-vm-threat-detection-overview>)

Let's evaluate the other options:

A Web Security Scanner: This service scans for common web application vulnerabilities like XSS, Flash injection, and mixed content. It is not designed to detect runtime threats like cryptocurrency mining

software

B Container Threat Detection: While Container Threat Detection (CTD) also detects cryptocurrency mining, it specifically focuses on runtime threats within GKE containers. The question asks for detection of "cryptocurrency mining software" generally, and VMs are a common target for such activity (and GKE nodes are VMs). VMTD provides a more general detection across Compute Engine VMs and GKE nodes for this specific type of threat. If the context explicitly mentioned containers or Cloud Run, CTD would be the more specific answer. However, for a general detection of "software" on "workloads", and given that VMTD explicitly lists "cryptocurrency mining software" for VMs, it is the most direct and broadly applicable answer among the choices.

C Rapid Vulnerability Detection: This service actively scans internet-exposed assets for network vulnerabilities and misconfigurations. It focuses on finding known vulnerabilities, not detecting active malicious processes like cryptocurrency mining.

Given the direct and explicit mention of cryptocurrency mining detection for VMs in its documentation, Virtual Machine Threat Detection is the correct SCC service to use.

Question: 259

Your organization is using Google Cloud to develop and host its applications. Following Google-recommended practices, the team has created dedicated projects for development and production. Your development team is located in Canada and Germany. The operations team works exclusively from Germany to adhere to local laws. You need to ensure that admin access to Google Cloud APIs is restricted to these countries and environments. What should you do?

- A. Create dedicated firewall policies for each environment at the organization level, and then apply these policies to the projects. Create a rule to restrict access based on geolocations.
- B. Group all development and production projects in separate folders. Activate the organization policy on the folders to restrict resource location according to the requirements.
- C. Create dedicated VPC Service Controls perimeters for development and production projects. Configure distinct ingress policies to allow access from the respective countries.
- D. Create dedicated IAM Groups for the Canadian and German developers. Grant access to the development and production projects according to the requirements.

Answer: C

Explanation:

The problem requires restricting admin access to Google Cloud APIs based on geographic location (Canada and Germany) and environment (development and production projects).

VPC Service Controls (VPC SC): VPC Service Controls is designed to create security perimeters around Google Cloud resources and services. Its primary purpose is to prevent data exfiltration and control access to Google Cloud APIs based on the context of the request, which includes the source IP address.

Extract Reference: "VPC Service Controls provides an extra layer of security defense for Google Cloud services that is independent of Identity and Access Management (IAM). While IAM enables granular identity-based access control, VPC Service Controls enables broader context-based perimeter security, including controlling data egress across the perimeter" (Google Cloud Documentation:

"Overview of VPC Service Controls" - <https://cloud.google.com/vpc-service-controls/docs/overview>) Service Perimeters for Environments: Creating dedicated perimeters for development and production projects allows for logical separation of environments, which aligns with the "dedicated projects for development and production" structure.

Ingress Policies with Geographic Restrictions: VPC Service Controls uses "ingress rules" to define who and from where requests can enter a service perimeter. These ingress rules can be configured to allow access based on various attributes, including the source IP address of the request. By allowing access from specific IP ranges corresponding to Canada and Germany, you effectively restrict administrative access to APIs from those countries. You can define "access levels" (which can include IP subnets or geographical origins) and attach them to ingress policies.

Extract Reference: "To allow ingress to resources, VPC Service Controls evaluates sources and identityType attributes as an AND condition. You must specify an accessLevel or a resource (Google Cloud project or VPC network), or set accessLevel attribute to '*' (Google Cloud Documentation: "Ingress and egress rules | VPC Service Controls" -

<https://cloud.google.com/vpc-service-controls/docs/ingress-egress-rules>)

Extract Reference (for Context-Aware Access which underpins access levels): "You can create different types of Context-Aware Access policies for accessing apps: IP, device, geographic origin, and custom access-level attributes" (Google Workspace Admin Help: "Protect your business with Context-Aware Access" -

<https://support.google.com/a/answer/9275380>) - While this references Workspace apps, the underlying mechanism of

Access Context Manager (used by VPC SC) supports geographic restrictions

Let's evaluate the other options:

A Create dedicated firewall policies restrict access based on geolocations: VPC firewall rules operate at the network level (Layers 3/4) within a VPC They control traffic between VM instances or to/from the internet for network services They do not directly control admin access to Google Cloud APIs (eg, via the console or gcloud CLI calls) originating from outside the VPC

B Activate the organization policy on the folders to restrict resource location: The Resource Location Restriction organization policy constraint restricts where new resources can be created or stored (eg, data residency requirements)

It does not restrict where administrators can connect from to manage these resources or access APIs

D Create dedicated IAM Groups Grant access: IAM (Identity and Access Management) controls who can access what resources and what actions they can perform It does not natively provide control over where the access originates from (eg, country-specific IP addresses)

Therefore, VPC Service Controls with properly configured ingress policies based on source IP/Access Levels is the recommended and most effective method for restricting admin access to Google Cloud APIs by geographic location and environment

Question: 260

You are responsible for the operation of your company's application that runs on Google Cloud The database for the application will be maintained by an external partner You need to give the partner team access to the database This access must be restricted solely to the database and cannot extend to any other resources within your company's network Your solution should follow Google- recommended practices What should you do?

- A. Add a public IP address to the application's database Create database users for each of the partner's employees Securely distribute the credentials for these users to the partner team
- B. Create accounts for the partner team in your corporate identity provider Synchronize these accounts with Google Cloud Identity Grant the accounts access to the database
- C. Ask the partner team to set up Cloud Identity accounts within their own corporate environment and identity provider Grant the partner's Cloud Identity accounts access to the database
- D. Configure Workforce Identity Federation for the partner Connect the identity pool provider to the partner's identity provider Grant the workforce pool resources access to the database

Answer: D

Explanation:

The problem requires granting an external partner team access solely to a database, without extending to other network resources, and following Google-recommended practices

Workforce Identity Federation: This Google Cloud IAM feature is specifically designed for scenarios where an organization needs to grant Google Cloud access to external identities (like partners, contractors, or customers) who are managed by their own identity provider (IdP) It allows these external users to authenticate using their existing credentials and then gain access to specified Google Cloud resources

Extract Reference: "Workforce Identity Federation lets you use an external identity provider (IdP) to authenticate and authorize a workforce—a group of users, such as employees, partners, and contractors—using IAM, so that the users can access Google Cloud services" (Google Cloud Documentation: "Workforce Identity Federation | IAM Documentation" -

<https://cloud.google.com/iam/docs/workforce-identity-federation>)

Extract Reference: "Secure access for partners and vendors Workforce Identity Federation can enable enterprises to selectively federate users from partner or vendor IdPs without requiring IT teams to sync or create a separate identity store to use Google Cloud resources" (Google Cloud Documentation: "Introducing Workforce Identity Federation" - <https://www.wazalio.io/introducing-workforce-identity-federation-to-easily-manage-workforce-access-to-google-cloud/>)

Least Privilege and Isolation: With Workforce Identity Federation, you create an identity pool and a provider that trusts the partner's IdP. You then grant IAM roles only to the workforce pool (or specific identities within it) on the specific database resource. This ensures fine-grained access control and prevents access to other resources in your network, directly addressing the least privilege and isolation requirements. The partner's identities are never synced into your internal Cloud Identity directory.

Let's evaluate the other options:

A Add a public IP address Securely distribute credentials: Adding a public IP address exposes the database to the internet, which is a major security risk and contradicts "restricted solely to the database and can not extend to any other resources within your company's network" as it allows any external network to potentially reach it. Distributing credentials manually is also not a Google-recommended secure practice.

B Create accounts for the partner team in your corporate identity provider Synchronize these accounts with Google Cloud Identity: This means you become responsible for managing the partner's identities within your own corporate IdP and syncing them. This is an unnecessary operational burden and blurs the lines of identity management. It also may inadvertently grant them broader network access if your corporate IdP is connected to your internal network resources.

C Ask the partner team to set up Cloud Identity accounts within their own corporate environment

and identity provider Grant the partner's Cloud Identity accounts access: While better than B, this implies the partner managing Cloud Identity accounts themselves and you directly granting IAM roles to their Cloud Identity users. Workforce Identity Federation is a more robust and scalable solution for federating any external IdP with Google Cloud IAM, rather than requiring partners to adopt Cloud Identity directly. Workforce Identity Federation is the explicit pattern for cross-organization access using existing external IdPs.

Therefore, Workforce Identity Federation is the most secure, scalable, and Google-recommended solution for granting restricted access to external partner teams.

Question: 261

Your organization needs to restrict the types of Google Cloud services that can be deployed within specific folders to enforce compliance requirements. You must apply these restrictions only to the designated folders without affecting other parts of the resource hierarchy. You want to use the most efficient and simple method. What should you do?

- A. Create an organization policy at the folder level using the "Restrict Resource Service Usage" constraint and define the allowed services per folder.
- B. Implement IAM conditions on service account creation within each folder.
- C. Create a global organization policy at the organization level with the "Restrict Resource Service Usage" constraint and apply exceptions for other folders.
- D. Configure VPC Service Controls perimeters around each folder and define the allowed services within the perimeter.

Answer: A

Explanation:

The problem requires restricting the types of Google Cloud services that can be deployed within specific folders to enforce compliance, without affecting other parts of the resource hierarchy, using the most efficient and simple method

Organization Policies: Organization policies allow you to define centralized, programmatic controls over your Google Cloud resources. They apply hierarchically, meaning a policy set on a folder applies to all projects and resources within that folder and its descendants.

Restrict Resource Service Usage Constraint: This specific organization policy constraint is designed precisely for controlling which Google Cloud services can be used (and thus deployed/created resources for) within a given part of the resource hierarchy. It supports both allowlists and denylists of service API identifiers.

Extract Reference: "The Restrict Resource Service Usage constraint controls the runtime access to all in-scope resources" and "This constraint can be used in two mutually exclusive ways: Denylist - resources of any service that isn't denied are allowed. Allowlist - resources of any service that isn't allowed are denied" (Google Cloud Documentation: "Restricting resource usage | Resource Manager Documentation" - <https://cloud.google.com/resource-manager/docs/organization-policy/restricting-resources>)

Folder-Level Application: Applying this organization policy at the folder level directly meets the requirement of applying restrictions "only to the designated folders without affecting other parts of

the resource hierarchy." This is more efficient and simpler than applying a global policy with numerous exceptions. Let's evaluate the other options:

B Implement IAM conditions on service account creation within each folder: IAM conditions control permissions for who can do what. While they can be used for very fine-grained access control, they are not designed to restrict the types of services that can be deployed directly. Controlling service account creation doesn't prevent a user with appropriate permissions from deploying other resources.

C Create a global organization policy and apply exceptions: While technically possible, this is less efficient and simple if the goal is to only restrict specific folders. Managing exceptions for the entire rest of the organization would be more complex than simply applying the policy directly where it's needed.

D Configure VPC Service Controls perimeters around each folder: VPC Service Controls primarily prevent data exfiltration and restrict API access at a network perimeter level. They are not designed to restrict which types of Google Cloud services can be deployed within a project or folder; rather, they control how allowed services interact with each other and with external endpoints. Therefore, creating an organization policy with the "Restrict Resource Service Usage" constraint at the folder level is the most efficient, simple, and direct method to achieve the stated goal.

Question: 262

Your organization has 3 TB of information in BigQuery and Cloud SQL. You need to develop a cost-effective, scalable, and secure strategy to anonymize the personally identifiable information (PII) that exists today. What should you do?

A. Scan your BigQuery and Cloud SQL data using the Cloud DLP data profiling feature. Use the data profiling results to create a de-identification strategy with either Cloud Sensitive Data Protection's de-identification templates or custom configurations.

B. Create a new BigQuery dataset and Cloud SQL instance. Copy a small subset of the data to these new locations. Use

- Cloud Data Loss Prevention API to scan this subset for PII Based on the results, create a custom anonymization script and apply the script to the entire 3 TB dataset in the original locations
- C. Export all 3TB of data from BigQuery and Cloud SQL to Cloud Storage Use Cloud Sensitive Data Protection to anonymize the exported data Re-import the anonymized data back into BigQuery and Cloud SQL
- D. Inspect a representative sample of the data in BigQuery and Cloud SQL to identify PII Based on this analysis, develop a custom script to anonymize the identified PII

Answer: A

Explanation:

The problem requires a cost-effective, scalable, and secure strategy to anonymize 3 TB of PII in BigQuery and Cloud SQL

Cloud Sensitive Data Protection (Cloud DLP): Cloud DLP is Google Cloud's fully managed service designed for discovering, classifying, and de-identifying sensitive data, including PII, across various

Google Cloud services and on-premises environments It is specifically built for scale and security Extract Reference:

"Sensitive Data Protection helps you discover, classify, and de-identify sensitive data" and "Sensitive Data Protection includes advanced de-identification techniques, such as tokenization, masking, and format-preserving encryption, to help you protect your sensitive data while preserving its utility for analysis" (Google Cloud Documentation: "Sensitive Data Protection overview" - <https://cloud.google.com/sensitive-data-protection/docs/overview>)

Data Profiling: Cloud DLP's data profiling feature automatically scans your data (eg, BigQuery tables, Cloud SQL databases) to identify where sensitive data resides and to understand its characteristics This is crucial for developing an effective de-identification strategy

Extract Reference: "Data profiles provide metrics and insights about your sensitive and high-risk data, helping you make informed decisions about data protection, access, and storage" and "When you configure Sensitive Data Protection to profile data in BigQuery, Cloud Storage, and Datastore, it automatically scans your data for sensitive information" (Google Cloud Documentation: "About data profiles" - <https://cloud.google.com/sensitive-data-protection/docs/concepts-data-profiles>) De-identification Templates and Custom Configurations: After profiling identifies the PII, Cloud DLP offers various de-identification methods, which can be applied using pre-built templates or custom configurations This allows for flexible and targeted anonymization For BigQuery specifically, DLP can integrate with remote functions for de-identification at query time, minimizing data movement Extract Reference: "De-identification techniques, like encryption, obfuscate raw sensitive identifiers in your data These techniques let you preserve the utility of your data for joining or analytics, while reducing the risk of handling the data" and "You can use this tutorial to replace that pipeline with a SQL query for only re-identification or both de-identification and re-identification" (Google Cloud Documentation: "De-identify BigQuery data at query time | Sensitive Data Protection Documentation" - <https://cloud.google.com/sensitive-data-protection/docs/deidentify-bq-tutorial>) Let's evaluate the other options:

B Copy small subset create a custom anonymization script apply the script to the entire 3 TB dataset: Relying on a "custom anonymization script" for 3 TB of data is generally not scalable, cost-effective, or secure compared to a managed service like Cloud DLP Custom scripts require significant development, testing, maintenance, and robust error handling for large datasets, and might introduce security vulnerabilities

C Export all 3TB of data to Cloud Storage anonymize Re-import: While feasible, exporting and reimporting 3 TB of data is a very time-consuming and potentially costly process due to data transfer and storage operations Cloud DLP can often process data in-place or integrate more efficiently, especially with BigQuery This option might not be the most cost-effective or efficient

D Inspect a representative sample develop a custom script: Similar to option B, this relies on a custom script, which is less

ideal for scalability, cost-effectiveness, and security than a managed service like Cloud DLP for 3 TB of sensitive data. Sample inspection is a good initial step, but the subsequent custom script for anonymization is the weak point.

Therefore, using Cloud DLP's data profiling to inform a de-identification strategy with its managed templates and configurations is the most appropriate, scalable, secure, and cost-effective approach for anonymizing PII in BigQuery and Cloud SQL.

Question: 263

Your organization is deploying a serverless web application on Cloud Run that must be publicly accessible over HTTPS. To meet security requirements, you need to terminate TLS at the edge, apply

threat mitigation, and prepare for geo-based access restrictions. What should you do?

- A. Make the Cloud Run service public by enabling allUsers access. Configure Identity-Aware Proxy (IAP) for authentication and IP-based access control. Use custom SSL certificates for HTTPS.
- B. Assign a custom domain to the Cloud Run service. Enable HTTPS. Configure IAM to allow allUsers to invoke the service. Use firewall rules and VPC Service Controls for geo-based restriction and traffic filtering.
- C. Deploy an external HTTP(S) load balancer with a serverless NEG that points to the Cloud Run service. Use a Google-managed certificate for TLS termination. Configure a Cloud Armor policy with geo-based access control.
- D. Create a Cloud DNS public zone for the Cloud Run URL. Bind a static IP to the service. Use VPC firewall rules to restrict incoming traffic based on IP ranges and threat signatures.

Answer: C

Explanation:

The problem requires a publicly accessible Cloud Run service with HTTPS, TLS termination at the edge, threat mitigation, and geo-based access restrictions.

External HTTP(S) Load Balancer: This is the standard Google Cloud component for exposing public-facing web applications, providing a single global IP address, global load balancing, and crucially, TLS termination at the edge of Google's network.

Serverless Network Endpoint Group (NEG): A serverless NEG connects an HTTP(S) Load Balancer to a Cloud Run service (or other serverless backends like Cloud Functions or App Engine), allowing the load balancer to route traffic to the serverless application. **Extract Reference:** "You can use an external HTTP(S) Load Balancer with a serverless network endpoint group (NEG) to integrate Cloud Run services with advanced load balancing features, such as Cloud CDN, Google Cloud Armor, and Cloud Identity-Aware Proxy" (Google Cloud Documentation: "Connect a Cloud Run service to an HTTP(S) Load Balancer | Cloud Run Documentation" - <https://cloud.google.com/run/docs/integrating/load-balancers>).

Google-Managed Certificate for TLS Termination: Using Google-managed SSL certificates simplifies the management of HTTPS, as Google handles the provisioning, renewal, and deployment of certificates. The external HTTP(S) Load Balancer terminates TLS using these certificates. **Extract Reference:** "Google-managed SSL certificates let you use Google Cloud's globally distributed infrastructure to automatically provision and renew your certificates" (Google Cloud Documentation: "Google-managed SSL certificates overview | Load Balancing" - <https://cloud.google.com/load-balancing/docs/ssl-certificates/google-managed-certs>).

Cloud Armor for Threat Mitigation and Geo-Based Access Control: Cloud Armor is a Web Application Firewall (WAF) service that integrates with HTTP(S) Load Balancers. It provides DDoS protection, allows custom rules for threat mitigation (eg, SQL injection, XSS), and supports geo-based access control rules to allow or deny traffic based on the source's

geographic region Extract Reference: "Google Cloud Armor helps protect your Google Cloud deployments from various threats, including Distributed Denial of Service (DDoS) attacks and application attacks such as cross-site scripting (XSS) and SQL injection (SQLi)" and "Google Cloud Armor supports geo-based access control, allowing you to filter requests based on geographic region" (Google Cloud Documentation: "Google Cloud Armor overview" -

<https://cloud.google.com/armor/docs/overview>)

Let's evaluate the other options:

A IAP for authentication and IP-based access control + custom SSL certs: IAP is primarily an authentication layer for users/identities, not a WAF for threat mitigation or geo-blocking for a publicly accessible service before any authentication takes place While it can apply IP-based access, it's typically for post-authentication controls

B Assign custom domain, enable HTTPS, allUsers IAM, firewall rules & VPC SC: While Cloud Run supports custom domains and HTTPS, and allUsers makes it public, direct Cloud Run services do not leverage traditional VPC firewall rules for public ingress VPC Service Controls are for API access and data exfiltration, not for WAF or geo-blocking of public web traffic

D Cloud DNS public zone, static IP, VPC firewall rules, threat signatures: Cloud Run services are managed and do not typically expose a static IP address that can be directly associated with VPC firewall rules for public traffic VPC firewall rules apply to VMs within a VPC, not to the managed global infrastructure of Cloud Run's public endpoints

Therefore, deploying an external HTTP(S) Load Balancer with a serverless NEG and integrating it with Google-managed certificates and Cloud Armor is the most comprehensive and Google- recommended solution for meeting all the specified security requirements for a publicly accessible Cloud Run application

Question: 264

A security audit uncovered several inconsistencies in your project's Identity and Access Management (IAM) configuration Some service accounts have overly permissive roles, and a few external collaborators have more access than necessary You need to gain detailed visibility into changes to IAM policies, user activity, service account behavior, and access to sensitive projects What should you do?

- A. Deploy the OS Config Management agent to your VMs Use OS Config Management to create patch management jobs and monitor system modifications
- B. Enable the metrics explorer in Cloud Monitoring to follow the service account authentication events and build alerts linked on it
- C. Use Cloud Audit Logs Create log export sinks to send these logs to a security information and event management (SIEM) solution for correlation with other event sources
- D. Configure Google Cloud Functions to be triggered by changes to IAM policies Analyze changes by using the policy simulator, send alerts upon risky modifications, and store event details

Answer: C

Explanation:

The problem requires gaining "detailed visibility into changes to IAM policies, user activity, service account behavior, and access to sensitive projects" due to security inconsistencies

Cloud Audit Logs: Cloud Audit Logs records administrative activities, data access, and system events across Google Cloud These logs are the primary source of truth for tracking "who did what, where, and when" in your Google Cloud environment

Extract Reference: "Cloud Audit Logs maintains the following audit logs for each project, folder, and organization: Admin Activity audit logs, Data Access audit logs, System Event audit logs, Policy Denied audit logs"

Extract Reference: "Admin Activity audit logs contain log entries for API calls or other actions that

modify the configuration or metadata of resources Data Access audit logs record API calls that read the configuration or metadata of resources, as well as user-provided data" (Google Cloud Documentation: "Cloud Audit Logs overview" - <https://cloud.google.com/logging/docs/audit>) These logs directly capture: Changes to IAM policies: Recorded in Admin Activity logs User activity: Recorded in Admin Activity and Data Access logs

Service account behavior: Actions performed by service accounts are logged in the same way as user actions

Access to sensitive projects: Data Access logs, especially for sensitive data services, record access events

Log Export Sinks: To gain "detailed visibility" and enable "correlation with other event sources," these audit logs should be exported to a centralized Security Information and Event Management (SIEM) solution Log sinks allow you to route logs from Cloud Logging to various destinations, including BigQuery, Cloud Storage, or Pub/Sub (which can then feed into a SIEM)

Extract Reference: "You can use sinks to route some or all of your logs to supported destinations" and "Many security information and event management (SIEM) systems can ingest logs through Cloud Pub/Sub" (Google Cloud Documentation: "Routing and storage overview | Cloud Logging" - <https://cloud.google.com/logging/docs/routing-overview>)

Let's evaluate the other options:

A OS Config Management agent: This service manages operating system configurations, patching, and inventory on VMs It is not designed to monitor or log IAM policy changes, user activity, or service account behavior within Google Cloud's IAM system

B Metrics Explorer in Cloud Monitoring: While Cloud Monitoring can provide some metrics related to service account authentication, it focuses on time-series data and operational health metrics It does not provide the detailed, event-level audit records necessary for forensic analysis of IAM policy changes, specific user actions, or granular access events to sensitive data that Cloud Audit Logs offer D Cloud Functions triggered by IAM policy changes + Policy Simulator: This describes a reactive automation pattern for some IAM changes While useful for immediate alerting on risky modifications, it's a custom solution for a subset of the requirements It doesn't inherently provide "detailed visibility" into all user activity or comprehensive service account behavior across all projects, nor does it replace the robust logging and correlation capabilities of a SIEM solution ingesting raw audit logs Cloud Audit Logs are the fundamental data source this approach would rely on

Therefore, leveraging Cloud Audit Logs and exporting them to a SIEM is the most comprehensive and recommended approach for gaining detailed visibility into IAM-related changes and activities across your Google Cloud organization

Question: 265

Your financial services company is migrating its operations to Google Cloud You are implementing a centralized logging strategy to meet strict regulatory compliance requirements Your company's Google Cloud organization has a dedicated folder for all production projects All audit logs, including Data Access logs from all current and future projects within this production folder, must be securely collected and stored in a central BigQuery dataset for long-term retention and analysis To prevent duplicate log storage and to enforce centralized control, you need to implement a logging solution that intercepts and overrides any project-level log sinks for these audit logs, to ensure that logs are not inadvertently routed elsewhere What should you do?

- A. Create an aggregated log sink at the production folder level with a destination of the central BigQuery dataset Configure an inclusion filter for all audit and Data Access logs Grant the Logs Bucket Writer role to the sink's service account on the production folder
- B. Create a log sink in each production project to route audit logs to the central BigQuery dataset Set the writer_identity

- field of each sink to a service account with BigQuery Data Editor permissions on the central dataset
- C. Create an aggregated log sink at the organization level with a destination of the central BigQuery dataset and a filter for all audit logs Use the `--include-children` flag and configure a log view for the production folder
- D. Create an intercepting aggregated log sink at the production folder level with the central BigQuery dataset as the destination Configure an inclusion filter for the necessary audit logs Grant the appropriate IAM permissions to the sink's `writer_identity` on the BigQuery dataset

Answer: D

Explanation:

The problem's critical requirements are: centralized collection of all audit logs (including Data Access logs) from a production folder to a central BigQuery dataset, with long-term retention, and most importantly, the ability to intercept and override any project-level log sinks to prevent duplicate storage or misrouting

Aggregated Log Sink at Folder Level: To centralize logs from all current and future projects within the "production folder," an aggregated sink configured at the folder level is the correct approach Logs generated in child projects will flow up to the folder level and be matched by this sink

Extract Reference: "Aggregated exports allow you to export logs from multiple Google Cloud projects, folders, or your entire organization An aggregated export can include all logs from all included resources, or you can use queries to include only specific logs" (Google Cloud Documentation: "Route logs to supported destinations | Cloud Logging" - https://cloud.google.com/logging/docs/export/aggregated_exports)

Intercepting Sink (`--intercept-logs` / `overrideDestinations`): This is the crucial feature to meet the "intercept and override" requirement When an aggregated sink is configured as an "intercepting" sink, any log entries that match its filter are immediately routed to its destination and are not processed by any lower-level sinks (eg, project-level sinks) This ensures that logs are not inadvertently routed elsewhere and prevents duplicate storage

Extract Reference: "An intercepting sink is an aggregated sink that, if it includes the `overrideDestinations` field set to true, stops matched log entries from propagating to lower-level sinks in the Cloud Logging resource hierarchy" (Google Cloud Documentation: "Route logs to supported destinations | Cloud Logging" -

https://cloud.google.com/logging/docs/export/aggregated_exports)

BigQuery Destination and IAM Permissions: BigQuery is specified as the long-term retention destination The sink's `writer_identity` (a service account automatically created for the sink) needs appropriate IAM permissions (eg, BigQuery Data Editor or BigQuery User) on the target BigQuery dataset to write logs

Inclusion Filter for Audit Logs: An inclusion filter is necessary to ensure only the required audit logs, including Data Access logs (`logName:"cloudaudit.googleapis.com"` OR `logName:"data_access"`), are routed

Let's evaluate the other options:

A Standard aggregated log sink Logs Bucket Writer: A standard aggregated sink does not intercept or override lower-level sinks Also, Logs Bucket Writer role is for Cloud Logging buckets, not BigQuery The correct role would be for BigQuery

B Log sink in each production project: This is not a centralized solution and would require manual configuration for every project, which is inefficient and error-prone for "all current and future projects" It also doesn't provide any override mechanism

C Aggregated log sink at the organization level configure a log view: While an organization-level sink offers broad centralization, if the requirement is specifically for a production folder and not the entire organization, a folder-level intercepting sink is more targeted A log view is for displaying logs, not for routing or overriding The `--include-children` flag is implied for aggregated sinks but doesn't provide the intercepting behavior

Therefore, creating an intercepting aggregated log sink at the production folder level, configured to send audit logs to BigQuery, is the precise solution to meet all the stated requirements, especially the crucial "intercept and override" condition

Question: 266

Your company is deploying a large number of containerized applications to GKE. The existing CI/CD pipeline uses Cloud Build to construct container images, transfers the images to Artifact Registry, and then deploys the images to GKE. You need to ensure that only images that have passed vulnerability scanning and meet specific corporate policies are allowed to be deployed. The process needs to be automated and integrated into the existing CI/CD pipeline. What should you do?

- A. Implement a custom script in the Cloud Build pipeline that uses a third-party vulnerability scanning tool. Fail the build if vulnerabilities are found.
- B. Configure GKE to use only images from a specific, trusted Artifact Registry repository. Manually inspect all images before pushing them to this repository.
- C. Configure a policy in Binary Authorization to use Artifact Analysis vulnerability scanning to only allow images that pass the scan to deploy to your GKE clusters.
- D. Enable Artifact Analysis vulnerability scanning and regularly scan images in Artifact Registry. Remove any images that do not meet the vulnerability requirements before deployment.

Answer: C

Explanation:

The problem requires ensuring that only images that have passed vulnerability scanning and meet corporate policies are allowed to be deployed to GKE, with the process being automated and integrated into the existing CI/CD pipeline.

Binary Authorization: This Google Cloud service is purpose-built to enforce deployment policies on images before they are run on Google Kubernetes Engine (GKE), Cloud Run, and other deployable platforms. It acts as a policy gate that prevents the deployment of non-compliant images.

Extract Reference: "Binary Authorization is a deploy-time security control that ensures only trusted container images are deployed on Google Kubernetes Engine (GKE), Cloud Run, and Anthos clusters" and "With Binary Authorization, you can require images to be signed by trusted authorities and enforce validation policies during deployment" (Google Cloud Documentation: "Binary Authorization overview" - <https://cloud.google.com/binary-authorization/docs/overview>)

Artifact Analysis (part of Container Analysis): Artifact Analysis (which includes Container Analysis) provides vulnerability scanning capabilities for container images stored in Artifact Registry. It generates findings and metadata about vulnerabilities.

Extract Reference: "Container Analysis is a service that scans your images for known vulnerabilities and provides metadata about them" (Google Cloud Documentation: "Overview | Container Analysis" - <https://cloud.google.com/container-analysis/docs/overview>)

Binary Authorization can be configured to integrate with Artifact Analysis (or other attestors) to check for vulnerability scan results as part of its deployment policy.

Integration and Automation: Binary Authorization policies can require attestations before deployment. An attestation confirms that an image meets specific criteria (eg, it has passed a vulnerability scan, it was signed by an approved CI/CD process, it adheres to corporate policies). Cloud Build can be configured to generate these attestations after a successful vulnerability scan (using Artifact Analysis). This fully automates the process and integrates directly into the CI/CD pipeline.

Extract Reference: "With Binary Authorization, you create a policy that enforces your requirements. The policy defines rules that govern deployment. For example, a policy can require all images to be signed by a trusted authority before deployment" (Google Cloud Documentation: "Binary Authorization overview" - <https://cloud.google.com/binary-authorization/docs/overview>) Let's evaluate the other options:

A Custom script in Cloud Build Fail the build: While scanning during the build is good practice (shift-left security), failing the build only prevents the image from being pushed. It doesn't prevent a developer or an automated process from manually

deploying an old or non-compliant image that might already exist in Artifact Registry, or from bypassing the build system

The enforcement needs to **happen at deployment time**

B Configure GKE to use only images from a specific trusted Artifact Registry repository Manually inspect all images: Manually inspecting images is not automated and does not scale for a "large number of containerized applications" It also doesn't programmatically enforce vulnerability scan results or corporate policies

D Enable Artifact Analysis vulnerability scanning and regularly scan images Remove any images that do not meet before deployment: This describes scanning and remediation, which are important However, it's a reactive approach ("remove any images") rather than a proactive enforcement ("only images that are allowed to be deployed") There's still a window where a non-compliant image could be deployed before it's removed Binary Authorization is the enforcement gate

Therefore, configuring Binary Authorization with a policy that integrates with Artifact Analysis (or requires attestations based on its findings) is the most robust, automated, and Google-recommended solution for enforcing deployment policies based on vulnerability scanning and corporate compliance