

DUMPSQUEEN

Google Cloud Certified - Professional Cloud Security Engineer

Google Professional-Cloud-Security-Engineer

Version Demo

Total Demo Questions: 10

Total Premium Questions: 134

Buy Premium PDF

<https://dumpsqueen.com>

support@dumpsqueen.com

dumpsqueen.com

QUESTION NO: 1

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: A D

QUESTION NO: 2

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: C D

QUESTION NO: 3

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: A C

Explanation:

Reference: <https://cloud.google.com/kms/docs/secret-management>

QUESTION NO: 4

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: B E

QUESTION NO: 5

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: B D

QUESTION NO: 6

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:

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

QUESTION NO: 7

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

QUESTION NO: 8

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>

QUESTION NO: 9

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

QUESTION NO: 10

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: D

Explanation:

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

- **Encryption type:** The kind of encryption used in the de-identification transformation.
- **Supported input values:** Minimum requirements for input values.
- **Surrogate annotation:** A user-specified annotation that is prepended to encrypted values to provide context to users and to provide information for Cloud DLP to use in the re-identification of a de-identified value. A surrogate annotation is required for re-identification of unstructured data. It is optional when transforming a column of structured, or tabular, data with a [RecordTransformation](#).