Free Questions for Professional-Cloud-Developer by actualtestdumps

Shared by Alexander on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Your team is building an application for a financial institution. The application's frontend runs on Compute Engine, and the data resides in Cloud SQL and one Cloud Storage bucket. The application will collect data containing PII, which will be stored in the Cloud SQL database and the Cloud Storage bucket. You need to secure the PII dat

a. What should you do?

Options:

A- 1) Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database
2) Using IAM, allow only the frontend service account to access the Cloud Storage bucket

B- 1) Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database
2) Enable private access to allow the frontend to access the Cloud Storage bucket privately

C- 1) Configure a private IP address for Cloud SQL
2) Use VPC-SC to create a service perimeter
3) Add the Cloud SQL database and the Cloud Storage bucket to the same service perimeter

D- 1) Configure a private IP address for Cloud SQL
2) Use VPC-SC to create a service perimeter
3) Add the Cloud SQL database and the Cloud Storage bucket to different service perimeters

Answer:

C

Question 2

Question Type: MultipleChoice

You are designing an application that consists of several microservices. Each microservice has its own RESTful API and will be deployed as a separate Kubernetes Service. You want to ensure that the consumers of these APIs aren't impacted when there is a change to your API, and also ensure that third-party systems aren't interrupted when new versions of the API are released. How should you configure the connection to the application following Google-recommended best practices?

Options:

A- Use an Ingress that uses the API's URL to route requests to the appropriate backend.

B- Leverage a Service Discovery system, and connect to the backend specified by the request.

C- Use multiple clusters, and use DNS entries to route requests to separate versioned backends.

D- Combine multiple versions in the same service, and then specify the API version in the POST request.

Answer:

D

Question 3

Question Type: MultipleChoice

Your company has deployed a new API to a Compute Engine instance. During testing, the API is not behaving as expected. You want to monitor the application over 12 hours to diagnose the problem within the application code without redeploying the application. Which tool should you use?

Options:

A- Cloud Trace

B- Cloud Monitoring

C- Cloud Debugger logpoints

D- Cloud Debugger snapshots

Answer:

C

Explanation:

https://cloud.google.com/debugger/docs/using/logpoints

Logpoints allow you to inject logging into running services without restarting or interfering with the normal function of the service

Question 4

Question Type: MultipleChoice

You recently migrated an on-premises monolithic application to a microservices application on Google Kubernetes Engine (GKE). The application has dependencies on backend services on-premises, including a CRM system and a MySQL database that contains personally identifiable information (PII). The backend services must remain on-premises to meet regulatory requirements.

You established a Cloud VPN connection between your on-premises data center and Google Cloud. You notice that some requests from your microservices application on GKE to the backend services are failing due to latency issues caused by fluctuating bandwidth, which is causing the application to crash. How should you address the latency issues?

Options:

A- Use Memorystore to cache frequently accessed PII data from the on-premises MySQL database

B- Use Istio to create a service mesh that includes the microservices on GKE and the on-premises services

C- Increase the number of Cloud VPN tunnels for the connection between Google Cloud and the on-premises services

D- Decrease the network layer packet size by decreasing the Maximum Transmission Unit (MTU) value from its default value on Cloud VPN

Answer:

C

Explanation:

https://cloud.google.com/network-connectivity/docs/vpn/concepts/choosing-networks-routing#route-alignment

Question 5

Question Type: MultipleChoice

You manage a microservices application on Google Kubernetes Engine (GKE) using Istio. You secure the communication channels between your microservices by implementing an Istio AuthorizationPolicy, a Kubernetes NetworkPolicy, and mTLS on your GKE cluster. You discover that HTTP requests between two Pods to specific URLs fail, while other requests to other URLs succeed. What is the

cause of the connection issue?

Options:

A- A Kubernetes NetworkPolicy resource is blocking HTTP traffic between the Pods.

B- The Pod initiating the HTTP requests is attempting to connect to the target Pod via an incorrect TCP port.

C- The Authorization Policy of your cluster is blocking HTTP requests for specific paths within your application.

D- The cluster has mTLS configured in permissive mode, but the Pod's sidecar proxy is sending unencrypted traffic in plain text.

Answer:

C

Question 6

Question Type: **MultipleChoice**

Your development team has built several Cloud Functions using Java along with corresponding integration and service tests. You are building and deploying the functions and launching the tests using Cloud Build. Your Cloud Build job is reporting deployment failures immediately after successfully validating the code. What should you do?

Options:

A- Check the maximum number of Cloud Function instances.

B- Verify that your Cloud Build trigger has the correct build parameters.

C- Retry the tests using the truncated exponential backoff polling strategy.

D- Verify that the Cloud Build service account is assigned the Cloud Functions Developer role.

Answer:

D

Explanation:

https://cloud.google.com/build/docs/securing-builds/configure-access-for-cloud-build-service-account

Question 7

Question Type: MultipleChoice

Your organization has recently begun an initiative to replatform their legacy applications onto Google Kubernetes Engine. You need to decompose a monolithic application into microservices. Multiple instances have read and write access to a configuration file, which is

stored on a shared file system. You want to minimize the effort required to manage this transition, and you want to avoid rewriting the application code. What should you do?

Options:

A- Create a new Cloud Storage bucket, and mount it via FUSE in the container.

B- Create a new persistent disk, and mount the volume as a shared PersistentVolume.

C- Create a new Filestore instance, and mount the volume as an NFS PersistentVolume.

D- Create a new ConfigMap and volumeMount to store the contents of the configuration file.

Answer:

D

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/configmap

ConfigMaps bind non-sensitive configuration artifacts such as configuration files, command-line arguments, and environment variables to your Pod containers and system components at runtime.

A ConfigMap separates your configurations from your Pod and components, which helps keep your workloads portable. This makes their configurations easier to change and manage, and prevents hardcoding configuration data to Pod specifications.

Question 8

Question Type: MultipleChoice

You need to migrate a standalone Java application running in an on-premises Linux virtual machine (VM) to Google Cloud in a cost-effective manner. You decide not to take the lift-and-shift approach, and instead you plan to modernize the application by converting it to a container. How should you accomplish this task?

Options:

A- Use Migrate for Anthos to migrate the VM to your Google Kubernetes Engine (GKE) cluster as a container.

B- Export the VM as a raw disk and import it as an image. Create a Compute Engine instance from the Imported image.

C- Use Migrate for Compute Engine to migrate the VM to a Compute Engine instance, and use Cloud Build to convert it to a container.

D- Use Jib to build a Docker image from your source code, and upload it to Artifact Registry. Deploy the application in a GKE cluster, and test the application.

Answer:

D

Explanation:

https://cloud.google.com/blog/products/application-development/introducing-jib-build-java-docker-images-better

Question 9

Question Type: MultipleChoice

Your company has a new security initiative that requires all data stored in Google Cloud to be encrypted by customer-managed encryption keys. You plan to use Cloud Key Management Service (KMS) to configure access to the keys. You need to follow the "separation of duties" principle and Google-recommended best practices. What should you do? (Choose two.)

Options:

A- Provision Cloud KMS in its own project.

B- Do not assign an owner to the Cloud KMS project.

C- Provision Cloud KMS in the project where the keys are being used.

D- Grant the roles/cloudkms.admin role to the owner of the project where the keys from Cloud KMS are being used.

E- Grant an owner role for the Cloud KMS project to a different user than the owner of the project where the keys from Cloud KMS are

being used.

Answer:

A, B

Explanation:

https://cloud.google.com/kms/docs/separation-of-duties#using_separate_project

Question 10

Question Type: MultipleChoice

You manage your company's ecommerce platform's payment system, which runs on Google Cloud. Your company must retain user logs for 1 year for internal auditing purposes and for 3 years to meet compliance requirements. You need to store new user logs on Google Cloud to minimize on-premises storage usage and ensure that they are easily searchable. You want to minimize effort while ensuring that the logs are stored correctly. What should you do?

Options:

A- Store the logs in a Cloud Storage bucket with bucket lock turned on.

B- Store the logs in a Cloud Storage bucket with a 3-year retention period.

C- Store the logs in Cloud Logging as custom logs with a custom retention period.

D- Store the logs in a Cloud Storage bucket with a 1-year retention period. After 1 year, move the logs to another bucket with a 2-year retention period.

Answer:

C

Explanation:

https://cloud.google.com/logging/docs/buckets#custom-retention

Question 11

Question Type: MultipleChoice

You are developing an application that consists of several microservices running in a Google Kubernetes Engine cluster. One microservice needs to connect to a third-party database running on-premises. You need to store credentials to the database and ensure that these credentials can be rotated while following security best practices. What should you do?

Options:

A- Store the credentials in a sidecar container proxy, and use it to connect to the third-party database.

B- Configure a service mesh to allow or restrict traffic from the Pods in your microservice to the database.

C- Store the credentials in an encrypted volume mount, and associate a Persistent Volume Claim with the client Pod.

D- Store the credentials as a Kubernetes Secret, and use the Cloud Key Management Service plugin to handle encryption and decryption.

Answer:

D

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/encrypting-secrets

By default, Google Kubernetes Engine (GKE) encrypts customer content stored at rest, including Secrets. GKE handles and manages this default encryption for you without any additional action on your part.

Application-layer secrets encryption provides an additional layer of security for sensitive data, such as Secrets, stored in etcd. Using this functionality, you can use a key managed with Cloud KMS to encrypt data at the application layer. This encryption protects against attackers who gain access to an offline copy of etcd.

Question 12

Question Type: MultipleChoice

You are in the final stage of migrating an on-premises data center to Google Cloud. You are quickly approaching your deadline, and discover that a web API is running on a server slated for decommissioning. You need to recommend a solution to modernize this API while migrating to Google Cloud. The modernized web API must meet the following requirements:

* Autoscales during high traffic periods at the end of each month

* Written in Python 3.x

* Developers must be able to rapidly deploy new versions in response to frequent code changes

You want to minimize cost, effort, and operational overhead of this migration. What should you do?

Options:

A- Modernize and deploy the code on App Engine flexible environment.

B- Modernize and deploy the code on App Engine standard environment.

C- Deploy the modernized application to an n1-standard-1 Compute Engine instance.

D- Ask the development team to re-write the application to run as a Docker container on Google Kubernetes Engine.

Answer:

B

Explanation:

https://cloud.google.com/appengine/docs/standard