PROFESSIONAL-CLOUD-DEVELOPER$^{Q\&As}$

Professional Cloud Developer

Pass Google PROFESSIONAL-CLOUD-DEVELOPER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.certbus.com/professional-cloud-developer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

QUESTION 1

Your company has a BigQuery dataset named "Master" that keeps information about employee travel and expenses. This information is organized by employee department. That means employees should only be able to view information for their department. You want to apply a security framework to enforce this requirement with the minimum number of steps.

What should you do?

A. Create a separate dataset for each department. Create a view with an appropriate WHERE clause to select records from a particular dataset for the specific department. Authorize this view to access records from your Master dataset. Give employees the permission to this department-specific dataset.

B. Create a separate dataset for each department. Create a data pipeline for each department to copy appropriate information from the Master dataset to the specific dataset for the department. Give employees the permission to this department-specific dataset.

C. Create a dataset named Master dataset. Create a separate view for each department in the Master dataset. Give employees access to the specific view for their department.

D. Create a dataset named Master dataset. Create a separate table for each department in the Master dataset. Give employees access to the specific table for their department.

Correct Answer: B

QUESTION 2

You are developing a flower ordering application Currently you have three microservices.

?Order Service (receives the orders).

?Order Fulfillment Service (processes the orders).

?Notification Service (notifies the customer when the order is filled).

You need to determine how the services will communicate with each other. You want incoming orders to be processed quickly and you need to collect order information for fulfillment. You also want to make sure orders are not lost between your services and are able to communicate asynchronously. How should the requests be processed?

A.

B.

C.

D.

Correct Answer: A

QUESTION 3

You have an application that uses an HTTP Cloud Function to process user activity from both desktop browser and mobile application clients. This function will serve as the endpoint for all metric submissions using HTTP POST.

Due to legacy restrictions, the function must be mapped to a domain that is separate from the domain requested by users on web or mobile sessions. The domain for the Cloud Function is https://fn.example.com. Desktop and mobile clients

use the domain https://www.example.com. You need to add a header to the function\\'s HTTP response so that only those browser and mobile sessions can submit metrics to the Cloud Function.

Which response header should you add?

A. Access-Control-Allow-Origin: *

B. Access-Control-Allow-Origin: https://*.example.com

C. Access-Control-Allow-Origin: https://fn.example.com

D. Access-Control-Allow-origin: https://www.example.com

Correct Answer: D

QUESTION 4

You have deployed a Java application to Cloud Run. Your application requires access to a database hosted on Cloud SQL Due to regulatory requirements: your connection to the Cloud SQL instance must use its internal IP address. How should you configure the connectivity while following Google-recommended best practices\\'?

A. Configure your Cloud Run service with a Cloud SQL connection.

B. Configure your Cloud Run service to use a Serverless VPC Access connector

C. Configure your application to use the Cloud SQL Java connector

D. Configure your application to connect to an instance of the Cloud SQL Auth proxy

Correct Answer: A

QUESTION 5

You are developing an application that will be launched on Compute Engine instances into multiple distinct projects, each corresponding to the environments in your software development process (development, QA, staging, and production).

The instances in each project have the same application code but a different configuration. During deployment, each instance should receive the application\\'s configuration based on the environment it serves. You want to minimize the

number of steps to configure this flow.

What should you do?

A. When creating your instances, configure a startup script using the gcloud command to determine the project name that indicates the correct environment.

B. In each project, configure a metadata key "environment" whose value is the environment it serves. Use your deployment tool to query the instance metadata and configure the application based on the "environment" value.

C. Deploy your chosen deployment tool on an instance in each project. Use a deployment job to retrieve the appropriate configuration file from your version control system, and apply the configuration when deploying the application on each instance.

D. During each instance launch, configure an instance custom-metadata key named "environment" whose value is the environment the instance serves. Use your deployment tool to query the instance metadata, and configure the application based on the "environment" value.

Correct Answer: B

Reference: https://cloud.google.com/compute/docs/metadata/overview

---

QUESTION 6

You want to re-architect a monolithic application so that it follows a microservices model.

You want to

accomplish this efficiently while minimizing the impact of this change to the business.

Which approach should you take?

A. Deploy the application to Compute Engine and turn on autoscaling.

B. Replace the application\\'s features with appropriate microservices in phases.

C. Refactor the monolithic application with appropriate microservices in a single effort and deploy it.

D. Build a new application with the appropriate microservices separate from the monolith and replace it when it is complete.

Correct Answer: C

Reference: https://cloud.google.com/solutions/migrating-a-monolithic-app-to-microservices- gke

---

QUESTION 7

You need to load-test a set of REST API endpoints that are deployed to Cloud Run. The API responds to HTTP POST requests Your load tests must meet the following requirements:

?Load is initiated from multiple parallel threads

?User traffic to the API originates from multiple source IP addresses.

?Load can be scaled up using additional test instances

You want to follow Google-recommended best practices How should you configure the load testing\\'?

A. Create an image that has cURL installed and configure cURLto run a test plan Deploy the image in a managed instance group, and run one instance of the image for each VM.

B. Create an image that has cURL installed and configure cURL to run a test plan Deploy the image in an unmanaged instance group, and run one instance of the image for each VM.

C. Deploy a distributed load testing framework on a private Google Kubernetes Engine Cluster Deploy additional Pods as needed to initiate more traffic and support the number of concurrent users.

D. Download the container image of a distributed load testing framework on Cloud Shell Sequentially start several instances of the container on Cloud Shell to increase the load on the API.

Correct Answer: C

QUESTION 8

Your team is developing a new application using a PostgreSQL database and Cloud Run. You are responsible for ensuring that all traffic is kept private on Google Cloud. You want to use managed services and follow Google-recommended best practices. What should you do?

A. 1. Enable Cloud SQL and Cloud Run in the same project.

2.

Configure a private IP address for Cloud SQL. Enable private services access.

3.

Create a Serverless VPC Access connector.

4.

Configure Cloud Run to use the connector to connect to Cloud SQL.

B. 1. Install PostgreSQL on a Compute Engine virtual machine (VM), and enable Cloud Run in the same project.

2.

Configure a private IP address for the VM. Enable private services access.

3.

Create a Serverless VPC Access connector.

4.

Configure Cloud Run to use the connector to connect to the VM hosting PostgreSQL.

C. 1. Use Cloud SQL and Cloud Run in different projects.

2.

Configure a private IP address for Cloud SQL. Enable private services access.

3.

Create a Serverless VPC Access connector.

4.

Set up a VPN connection between the two projects. Configure Cloud Run to use the connector to connect to Cloud SQL.

D. 1. Install PostgreSQL on a Compute Engine VM, and enable Cloud Run in different projects.

2.

Configure a private IP address for the VM. Enable private services access.

3.

Create a Serverless VPC Access connector.

4.

Set up a VPN connection between the two projects. Configure Cloud Run to use the connector to access the VM hosting PostgreSQL

Correct Answer: A

Explanation: https://cloud.google.com/sql/docs/postgres/connect-run#private-ip

---

QUESTION 9

You are running a web application on Google Kubernetes Engine that you inherited. You want to determine whether the application is using libraries with known vulnerabilities or is vulnerable to XSS attacks. Which service should you use?

A. Google Cloud Armor

B. Debugger

C. Web Security Scanner

D. Error Reporting

Correct Answer: C

https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner- overview Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to exercise as many user inputs and event handlers as possible.

QUESTION 10

You are designing a schema for a table that will be moved from MySQL to Cloud Bigtable. The MySQL table is as

follows:

```
AccountActivity
(
Account_id int,
Event_timestamp datetime,
Transaction_type string,
Amount numeric(18, 4)
) primary key (Account_id, Event_timestamp)
```

How should you design a row key for Cloud Bigtable for this table?

A. Set Account_id as a key.

B. Set Account_id_Event_timestamp as a key.

C. Set Event_timestamp_Account_id as a key.

D. Set Event_timestamp as a key.

Correct Answer: C

QUESTION 11

You are developing an ecommerce application that stores customer, order, and inventory data as relational tables inside Cloud Spanner. During a recent load test, you discover that Spanner performance is not scaling linearly as expected. Which of the following is the cause?

A. The use of 64-bit numeric types for 32-bit numbers.

B. The use of the STRING data type for arbitrary-precision values.

C. The use of Version 1 UUIDs as primary keys that increase monotonically.

D. The use of LIKE instead of STARTS_WITH keyword for parameterized SQL queries.

Correct Answer: C

QUESTION 12

You are building an API that will be used by Android and iOS apps The API must: ?Support HTTPs ?Minimize bandwidth cost ?Integrate easily with mobile apps Which API architecture should you use?

A. RESTful APIs

B. MQTT for APIs

C. gRPC-based APIs

D. SOAP-based APIs

Correct Answer: A

Reference: https://www.devteam.space/blog/how-to-build-restful-api-for-your-mobile-app/

QUESTION 13

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?

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.

Correct Answer: D

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 14

You are developing a new public-facing application that needs to retrieve specific properties in the metadata of users\\' objects in their respective Cloud Storage buckets. Due to privacy and data residency requirements, you must retrieve only the metadata and not the object data. You want to maximize the performance of the retrieval process. How should you retrieve the metadata?

A. Use the patch method.

B. Use the compose method.

C. Use the copy method.

D. Use the fields request parameter.

Correct Answer: D

Explanation: https://cloud.google.com/storage/docs/json_api/v1/objects/get

QUESTION 15

You need to deploy an internet-facing microservices application to Google Kubernetes Engine (GKE). You want to validate new features using the A/B testing method. You have the following requirements for deploying new container image releases

?There is no downtime when new container images are deployed.

?New production releases are tested and verified using a subset of production users.

What should you do?

A. 1 Configure your CI/CD pipeline to update the Deployment manifest file by replacing the container version with the latest version. 2 Recreate the Pods in your cluster by applying the Deployment manifest file. 3 Validate the application\\'s performance by comparing its functionality with the previous release version and roll back if an issue arises.

B. 1 install the Anthos Service Mesh on your GKE cluster. 2 Create two Deployments on the GKE cluster and label them with different version names. 3 Create a VirtualService with a routing rule to send a small percentage of traffic to the Deployment that references the new version of the application.

C. 1 Create a second namespace on GKE for the new release version. 2 Create a Deployment configuration for the second namespace with the desired number of Pods. 3 Deploy new container versions in the second namespace. 4 Update the ingress configuration to route traffic to the namespace with the new container versions.

D. 1. Implement a rolling update pattern by replacing the Pods gradually with the new release versify. 2 Validate the application\\'s performance for the new subset of users during the rollout and roll back if an issue arises.

Correct Answer: C