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.pass4itsure.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

PROFESSIONAL-CLOUD-DEVELOPER PDF Dumps | PROFESSIONAL-CLOUD-DEVELOPER Study Guide |
PROFESSIONAL-CLOUD-DEVELOPER Exam Questions

1 / 9

QUESTION 1

You need to copy directory local-scripts and all of its contents from your local workstation to a Compute Engine virtual machine instance. Which command should you use?

A. gsutil cp --project "my-gcp-project" -r ~/local-scripts/ gcp-instance-name:~/ server-scripts/ --zone "us-east1-b"

B. gsutil cp --project "my-gcp-project" -R ~/local-scripts/ gcp-instance-name:~/ server-scripts/ --zone "us-east1-b"

C. gcloud compute scp --project "my-gcp-project" --recurse ~/local-scripts/ gcpinstance-name:~/server-scripts/ --zone "us-east1-b"

D. gcloud compute mv --project "my-gcp-project" --recurse ~/local-scripts/ gcpinstance-name:~/server-scripts/ --zone "us-east1-b"

Correct Answer: C

Reference: https://cloud.google.com/sdk/gcloud/reference/compute/copy-files

QUESTION 2

Before promoting your new application code to production, you want to conduct testing across a variety of different users. Although this plan is risky, you want to test the new version of the application with production users and you want to control which users are forwarded to the new version of the application based on their operating system. If bugs are discovered in the new version, you want to roll back the newly deployed version of the application as quickly as possible.

What should you do?

A. Deploy your application on Cloud Run. Use traffic splitting to direct a subset of user traffic to the new version based on the revision tag.

B. Deploy your application on Google Kubernetes Engine with Anthos Service Mesh. Use traffic splitting to direct a subset of user traffic to the new version based on the user-agent header.

C. Deploy your application on App Engine. Use traffic splitting to direct a subset of user traffic to the new version based on the IP address.

D. Deploy your application on Compute Engine. Use Traffic Director to direct a subset of user traffic to the new version based on predefined weights.

Correct Answer: B

QUESTION 3

Your company has a BigQuery data mart that provides analytics information to hundreds of employees. One user of wants to run jobs without interrupting important workloads. This user isn\'t concerned about the time it takes to run these jobs. You want to fulfill this request while minimizing cost to the company and the effort required on your part.

What should you do?

PROFESSIONAL-CLOUD-DEVELOPER PDF Dumps | PROFESSIONAL-CLOUD-DEVELOPER Study Guide | PROFESSIONAL-CLOUD-DEVELOPER Exam Questions

3 / 9

A. Ask the user to run the jobs as batch jobs.

B. Create a separate project for the user to run jobs.

C. Add the user as a job.user role in the existing project.

D. Allow the user to run jobs when important workloads are not running.

Correct Answer: B

QUESTION 4

HipLocal\\'s data science team wants to analyze user reviews. How should they prepare the data?

A. Use the Cloud Data Loss Prevention API for redaction of the review dataset.

B. Use the Cloud Data Loss Prevention API for de-identification of the review dataset.

C. Use the Cloud Natural Language Processing API for redaction of the review dataset.

D. Use the Cloud Natural Language Processing API for de-identification of the review dataset.

Correct Answer: B

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

QUESTION 5

You have an HTTP Cloud Function that is called via POST. Each submission\\'s request body has a flat, unnested JSON structure containing numeric and text data. After the Cloud Function completes, the collected data should be immediately available for ongoing and complex analytics by many users in parallel. How should you persist the submissions?

A. Directly persist each POST request\\'s JSON data into Datastore.

B. Transform the POST request\\'s JSON data, and stream it into BigQuery.

C. Transform the POST request\\'s JSON data, and store it in a regional Cloud SQL cluster.

D. Persist each POST request\\'s JSON data as an individual file within Cloud Storage, with the file name containing the request identifier.

Correct Answer: D

QUESTION 6

You manage an application that runs in a Compute Engine instance. You also have multiple backend services executing in stand-alone Docker containers running in Compute Engine instances. The Compute Engine instances supporting the backend services are scaled by managed instance groups in multiple regions. You want your calling application to be

PROFESSIONAL-CLOUD-DEVELOPER PDF Dumps | PROFESSIONAL-CLOUD-DEVELOPER Study Guide | PROFESSIONAL-CLOUD-DEVELOPER Exam Questions

4 / 9

loosely coupled. You need to be able to invoke distinct service implementations that are chosen based on the value of an HTTP header found in the request. Which Google Cloud feature should you use to invoke the backend services?

A. Traffic Director

B. Service Directory

C. Anthos Service Mesh

D. Internal HTTP(S) Load Balancing

Correct Answer: D

QUESTION 7

You are using Cloud Build build to promote a Docker image to Development, Test, and Production environments. You need to ensure that the same Docker image is deployed to each of these environments. How should you identify the Docker image in your build?

A. Use the latest Docker image tag.

B. Use a unique Docker image name.

C. Use the digest of the Docker image.

D. Use a semantic version Docker image tag.

Correct Answer: D

QUESTION 8

You are using the Cloud Client Library to upload an image in your application to Cloud Storage. Users of the application report that occasionally the upload does not complete and the client library reports an HTTP 504 Gateway Timeout error. You want to make the application more resilient to errors. What changes to the application should you make?

A. Write an exponential backoff process around the client library call.

B. Write a one-second wait time backoff process around the client library call.

C. Design a retry button in the application and ask users to click if the error occurs.

D. Create a queue for the object and inform the users that the application will try again in 10 minutes.

Correct Answer: A

QUESTION 9

You have an application running in App Engine. Your application is instrumented with Stackdriver Trace. The /product-details request reports details about four known unique products at /sku-details as shown below. You want to reduce the time it takes for the request to complete. What should you do?
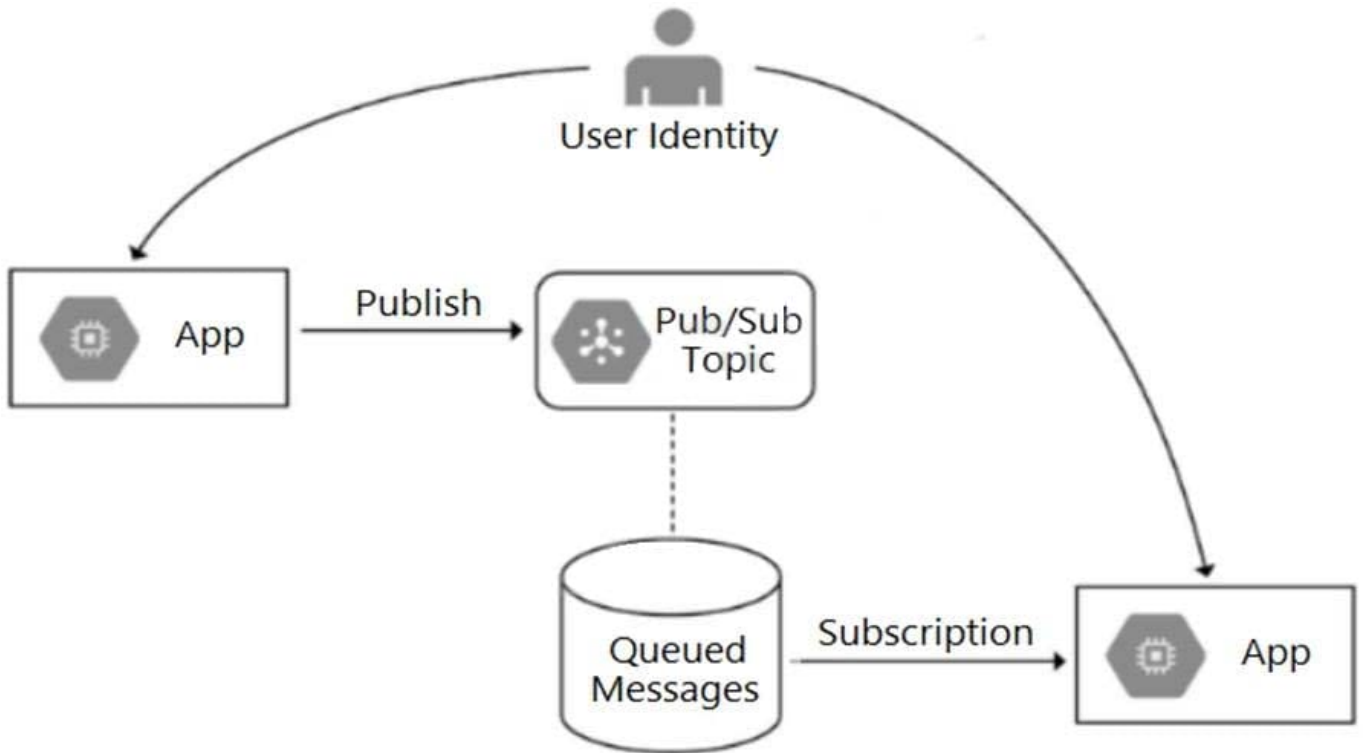
A. Increase the size of the instance class.

B. Change the Persistent Disk type to SSD.

C. Change /product-details to perform the requests in parallel.

D. Store the /sku-details information in a database, and replace the webservice call with a database query.

Correct Answer: C

QUESTION 10

You recently migrated a monolithic application to Google Cloud by breaking it down into microservices. One of the microservices is deployed using Cloud Functions. As you modernize the application, you make a change to the API of the service that is backward-incompatible. You need to support both existing callers who use the original API and new callers who use the new API. What should you do?

A. Leave the original Cloud Function as-is and deploy a second Cloud Function with the new API. Use a load balancer to distribute calls between the versions.

B. Leave the original Cloud Function as-is and deploy a second Cloud Function that includes only the changed API. Calls are automatically routed to the correct function.

C. Leave the original Cloud Function as-is and deploy a second Cloud Function with the new API. Use Cloud Endpoints to provide an API gateway that exposes a versioned API.

D. Re-deploy the Cloud Function after making code changes to support the new API. Requests for both versions of the API are fulfilled based on a version identifier included in the call.

Correct Answer: D

Reference: https://cloud.google.com/endpoints/docs/openapi/versioning-an-api

QUESTION 11

You are deploying a microservices application to Google Kubernetes Engine (GKE). The application will receive daily
updates. You expect to deploy a large number of distinct containers that will run on the Linux operating system (OS).
You want to be alerted to any known OS vulnerabilities in the new containers. You want to follow Google-recommended
best practices. What should you do?

A. Use the gcloud CLI to call Container Analysis to scan new container images. Review the vulnerability results before
each deployment.

B. Enable Container Analysis, and upload new container images to Artifact Registry. Review the vulnerability results
before each deployment.

C. Enable Container Analysis, and upload new container images to Artifact Registry. Review the critical vulnerability
results before each deployment.

D. Use the Container Analysis REST API to call Container Analysis to scan new container images. Review the
vulnerability results before each deployment.

Correct Answer: B

https://cloud.google.com/container-analysis/docs/automated-scanning-howto https://cloud.google.com/container-
analysis/docs/os-overview says: The Container Scanning API allows you to automate OS vulnerability detection,
scanning each time you push an image to Container Registry or Artifact Registry. Enabling this API also triggers
language package scans for Go and Java vulnerabilities (Preview).

QUESTION 12

Your team develops services that run on Google Kubernetes Engine. You need to standardize their log data using
Google-recommended practices and make the data more useful in the fewest number of steps. What should you do?
(Choose two.)

A. Create aggregated exports on application logs to BigQuery to facilitate log analytics.

B. Create aggregated exports on application logs to Cloud Storage to facilitate log analytics.

C. Write log output to standard output (stdout) as single-line JSON to be ingested into Cloud Logging as structured
logs.

D. Mandate the use of the Logging API in the application code to write structured logs to Cloud Logging.

E. Mandate the use of the Pub/Sub API to write structured data to Pub/Sub and create a Dataflow streaming pipeline to
normalize logs and write them to BigQuery for analytics.

Correct Answer: AC

https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs#best_practices

QUESTION 13

PROFESSIONAL-CLOUD-DEVELOPER PDF Dumps | PROFESSIONAL-CLOUD-DEVELOPER Study Guide |
PROFESSIONAL-CLOUD-DEVELOPER Exam Questions

7 / 9

You have an application deployed in Google Kubernetes Engine (GKE). You need to update the application to make authorized requests to Google Cloud managed services. You want this to be a one-time setup, and you need to follow security best practices of auto-rotating your security keys and storing them in an encrypted store. You already created a service account with appropriate access to the Google Cloud service. What should you do next?

A. Assign the Google Cloud service account to your GKE Pod using Workload Identity.

B. Export the Google Cloud service account, and share it with the Pod as a Kubernetes Secret.

C. Export the Google Cloud service account, and embed it in the source code of the application.

D. Export the Google Cloud service account, and upload it to HashiCorp Vault to generate a dynamic service account for your application.

Correct Answer: A

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

Applications running on GKE might need access to Google Cloud APIs such as Compute Engine API, BigQuery Storage API, or Machine Learning APIs.

Workload Identity allows a Kubernetes service account in your GKE cluster to act as an IAM service account. Pods that use the configured Kubernetes service account automatically authenticate as the IAM service account when accessing

Google Cloud APIs. Using Workload Identity allows you to assign distinct, fine-grained identities and authorization for each application in your cluster.

QUESTION 14

Your website is deployed on Compute Engine. Your marketing team wants to test conversion rates between 3 different website designs. Which approach should you use?

A. Deploy the website on App Engine and use traffic splitting.

B. Deploy the website on App Engine as three separate services.

C. Deploy the website on Cloud Functions and use traffic splitting.

D. Deploy the website on Cloud Functions as three separate functions.

Correct Answer: A

Reference: https://cloud.google.com/appengine/docs/standard/python/splitting-traffic

---

QUESTION 15

You are developing a web application that contains private images and videos stored in a Cloud Storage bucket. Your users are anonymous and do not have Google Accounts. You want to use your application-specific logic to control access to the images and videos. How should you configure access?

A. Cache each web application user\'s IP address to create a named IP table using Google Cloud Armor. Create a Google Cloud Armor security policy that allows users to access the backend bucket.

PROFESSIONAL-CLOUD-DEVELOPER PDF Dumps | PROFESSIONAL-CLOUD-DEVELOPER Study Guide | PROFESSIONAL-CLOUD-DEVELOPER Exam Questions

8 / 9

B. Grant the Storage Object Viewer IAM role to allUsers. Allow users to access the bucket after authenticating through your web application.

C. Configure Identity-Aware Proxy (IAP) to authenticate users into the web application. Allow users to access the bucket after authenticating through IAP.

D. Generate a signed URL that grants read access to the bucket. Allow users to access the URL after authenticating through your web application.

Correct Answer: D

https://cloud.google.com/storage/docs/access-control/signed-urls#should-you-use In some scenarios, you might not want to require your users to have a Google account in order to access Cloud Storage, but you still want to control access using your application-specific logic. The typical way to address this use case is to provide a signed URL to a user, which gives the user read, write, or delete access to that resource for a limited time. You specify an expiration time when you create the signed URL. Anyone who knows the URL can access the resource until the expiration time for the URL is reached or the key used to sign the URL is rotated.

[PROFESSIONAL-CLOUD-DEVELOPER PDF Dumps](#) [PROFESSIONAL-CLOUD-DEVELOPER Study Guide](#) [PROFESSIONAL-CLOUD-DEVELOPER Exam Questions](#)