



Vendor: Google

Exam Code: Professional-Cloud-Developer

Exam Name: Professional Cloud Developer

Version: DEMO

QUESTION 1

You are a developer working with the CI/CD team to troubleshoot a new feature that your team introduced. The CI/CD team used HashiCorp Packer to create a new Compute Engine image from your development branch. The image was successfully built, but is not booting up. You need to investigate the issue with the CI/CD team. What should you do?

- A. Create a new feature branch, and ask the build team to rebuild the image.
- B. Shut down the deployed virtual machine, export the disk, and then mount the disk locally to access the boot logs.
- C. Install Packer locally, build the Compute Engine image locally, and then run it in your personal Google Cloud project.
- D. Check Compute Engine OS logs using the serial port, and check the Cloud Logging logs to confirm access to the serial port.

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-using-serial-console>

QUESTION 2

You are porting an existing Apache/MySQL/PHP application stack from a single machine to Google Kubernetes Engine. You need to determine how to containerize the application. Your approach should follow Google-recommended best practices for availability. What should you do?

- A. Package each component in a separate container. Implement readiness and liveness probes.
- B. Package the application in a single container. Use a process management tool to manage each component.
- C. Package each component in a separate container. Use a script to orchestrate the launch of the components.
- D. Package the application in a single container. Use a bash script as an entrypoint to the container, and then spawn each component as a background job.

Answer: A

Explanation:

https://cloud.google.com/architecture/best-practices-for-building-containers#package_a_single_app_per_container

When you start working with containers, it's a common mistake to treat them as virtual machines that can run many different things simultaneously. A container can work this way, but doing so reduces most of the advantages of the container model. For example, take a classic Apache/MySQL/PHP stack: you might be tempted to run all the components in a single container. However, the best practice is to use two or three different containers: one for Apache, one for MySQL, and potentially one for PHP if you are running PHP-FPM.

QUESTION 3

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.

Answer: B

Explanation:

<https://cloud.google.com/compute/docs/metadata/setting-custom-metadata#set-custom-project-wide-metadata>

QUESTION 4

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.

Answer: C

Explanation:

https://cloud.google.com/spanner/docs/schema-and-data-model#choosing_a_primary_key

QUESTION 5

You are developing an application that reads credit card data from a Pub/Sub subscription. You have written code and completed unit testing. You need to test the Pub/Sub integration before deploying to Google Cloud. What should you do?

- A. Create a service to publish messages, and deploy the Pub/Sub emulator. Generate random content in the publishing service, and publish to the emulator.
- B. Create a service to publish messages to your application. Collect the messages from Pub/Sub in production, and replay them through the publishing service.
- C. Create a service to publish messages, and deploy the Pub/Sub emulator. Collect the messages from Pub/Sub in production, and publish them to the emulator.
- D. Create a service to publish messages, and deploy the Pub/Sub emulator. Publish a standard set of testing messages from the publishing service to the emulator.

Answer: D

Explanation:

<https://cloud.google.com/pubsub/docs/emulator>

QUESTION 6

You are designing an application that will subscribe to and receive messages from a single Pub/Sub topic and insert corresponding rows into a database. Your application runs on Linux and leverages preemptible virtual machines to reduce costs. You need to create a shutdown script that will initiate a graceful shutdown.

What should you do?

- A. Write a shutdown script that uses inter-process signals to notify the application process to disconnect from the database.
- B. Write a shutdown script that broadcasts a message to all signed-in users that the Compute Engine instance is going down and instructs them to save current work and sign out.
- C. Write a shutdown script that writes a file in a location that is being polled by the application once every five minutes. After the file is read, the application disconnects from the database.
- D. Write a shutdown script that publishes a message to the Pub/Sub topic announcing that a shutdown is in progress. After the application reads the message, it disconnects from the database.

Answer: A

Explanation:

Compute Engine sends a preemption notice to the instance in the form of an ACPI G2 Soft Off signal. You can use a shutdown script to handle the preemption notice and complete cleanup actions before the instance stops.

<https://cloud.google.com/compute/docs/instances/preemptible#preemption>

QUESTION 7

You work for a web development team at a small startup. Your team is developing a Node.js application using Google Cloud services, including Cloud Storage and Cloud Build. The team uses a Git repository for version control. Your manager calls you over the weekend and instructs you to make an emergency update to one of the company's websites, and you're the only developer available. You need to access Google Cloud to make the update, but you don't have your work laptop. You are not allowed to store source code locally on a non-corporate computer. How should you set up your developer environment?

- A. Use a text editor and the Git command line to send your source code updates as pull requests from a public computer.
- B. Use a text editor and the Git command line to send your source code updates as pull requests from a virtual machine running on a public computer.
- C. Use Cloud Shell and the built-in code editor for development. Send your source code updates as pull requests.
- D. Use a Cloud Storage bucket to store the source code that you need to edit. Mount the bucket to a public computer as a drive, and use a code editor to update the code. Turn on versioning for the bucket, and point it to the team's Git repository.

Answer: C

Explanation:

<https://cloud.google.com/shell/docs>

Cloud Shell is an interactive shell environment for Google Cloud that lets you learn and experiment with Google Cloud and manage your projects and resources from your web browser.

With Cloud Shell, the Google Cloud CLI and other utilities you need are pre-installed, fully authenticated, up-to-date, and always available when you need them. Cloud Shell comes with a built-in code editor with an integrated Cloud Code experience, allowing you to develop, build, debug, and deploy your cloud-based apps entirely in the cloud.

QUESTION 8

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.

Answer: CD

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/about-logs#best_practices

QUESTION 9

You are designing a deployment technique for your new applications on Google Cloud. As part of your deployment planning, you want to use live traffic to gather performance metrics for both new and existing applications. You need to test against the full production load prior to launch. What should you do?

- A. Use canary deployment
- B. Use blue/green deployment
- C. Use rolling updates deployment
- D. Use A/B testing with traffic mirroring during deployment

Answer: D

Explanation:

You need to test against the full production load prior to launch" It's impossible with canary. A/B testing with traffic mirroring during deployment" is the only one possibility we have to test the entire traffic before the roll out.

QUESTION 10

You support an application that uses the Cloud Storage API. You review the logs and discover multiple HTTP 503 Service Unavailable error responses from the API. Your application logs the error and does not take any further action. You want to implement Google-recommended retry logic to improve success rates. Which approach should you take?

- A. Retry the failures in batch after a set number of failures is logged.
- B. Retry each failure at a set time interval up to a maximum number of times.
- C. Retry each failure at increasing time intervals up to a maximum number of tries.
- D. Retry each failure at decreasing time intervals up to a maximum number of tries.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/retry-strategy>

QUESTION 11

You need to redesign the ingestion of audit events from your authentication service to allow it to handle a large increase in traffic. Currently, the audit service and the authentication system run in the same Compute Engine virtual machine. You plan to use the following Google Cloud tools in the new architecture:

- Multiple Compute Engine machines, each running an instance of the authentication service
- Multiple Compute Engine machines, each running an instance of the audit service
- Pub/Sub to send the events from the authentication services.

How should you set up the topics and subscriptions to ensure that the system can handle a large volume of messages and can scale efficiently?

- A. Create one Pub/Sub topic. Create one pull subscription to allow the audit services to share the messages.
- B. Create one Pub/Sub topic. Create one pull subscription per audit service instance to allow the services to share the messages.
- C. Create one Pub/Sub topic. Create one push subscription with the endpoint pointing to a load balancer in front of the audit services.
- D. Create one Pub/Sub topic per authentication service. Create one pull subscription per topic to be used by one audit service.
- E. Create one Pub/Sub topic per authentication service. Create one push subscription per topic, with the endpoint pointing to one audit service.

Answer: A

Explanation:

<https://cloud.google.com/pubsub/docs/subscriber>

Multiple subscribers can make pull calls to the same "shared" subscription. Each subscriber will receive a subset of the messages.

QUESTION 12

You are developing a marquee stateless web application that will run on Google Cloud. The rate of the incoming user traffic is expected to be unpredictable, with no traffic on some days and large spikes on other days. You need the application to automatically scale up and down, and you need to minimize the cost associated with running the application. What should you do?

- A. Build the application in Python with Firestore as the database. Deploy the application to Cloud Run.
- B. Build the application in C# with Firestore as the database. Deploy the application to App Engine flexible environment.
- C. Build the application in Python with CloudSQL as the database. Deploy the application to App Engine standard environment.
- D. Build the application in Python with Firestore as the database. Deploy the application to a Compute Engine managed instance group with autoscaling.

Answer: A

QUESTION 13

You have written a Cloud Function that accesses other Google Cloud resources. You want to

secure the environment using the principle of least privilege. What should you do?

- A. Create a new service account that has Editor authority to access the resources. The deployer is given permission to get the access token.
- B. Create a new service account that has a custom IAM role to access the resources. The deployer is given permission to get the access token.
- C. Create a new service account that has Editor authority to access the resources. The deployer is given permission to act as the new service account.
- D. Create a new service account that has a custom IAM role to access the resources. The deployer is given permission to act as the new service account.

Answer: D

Explanation:

<https://cloud.google.com/functions/docs/securing/function-identity#individual>

In order to deploy a function with a user-managed service account, the deployer must have the iam.serviceAccounts.actAs permission on the service account being deployed.

QUESTION 14

You are a SaaS provider deploying dedicated blogging software to customers in your Google Kubernetes Engine (GKE) cluster. You want to configure a secure multi-tenant platform to ensure that each customer has access to only their own blog and can't affect the workloads of other customers. What should you do?

- A. Enable Application-layer Secrets on the GKE cluster to protect the cluster.
- B. Deploy a namespace per tenant and use Network Policies in each blog deployment.
- C. Use GKE Audit Logging to identify malicious containers and delete them on discovery.
- D. Build a custom image of the blogging software and use Binary Authorization to prevent untrusted image deployments.

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/multitenancy-overview>

QUESTION 15

You have decided to migrate your Compute Engine application to Google Kubernetes Engine. You need to build a container image and push it to Artifact Registry using Cloud Build. What should you do? (Choose two.)

- A. Run `gcloud builds submit` in the directory that contains the application source code.
- B. Run `gcloud run deploy app-name --image gcr.io/$PROJECT_ID/app-name` in the directory that contains the application source code.
- C. Run `gcloud container images add-tag gcr.io/$PROJECT_ID/app-name gcr.io/$PROJECT_ID/app-name:latest` in the directory that contains the application source code.
- D. In the application source directory, create a file named `cloudbuild.yaml` that contains the following contents:

```
steps:
- name: 'gcr.io/cloud-builders/docker'
  args: ['build', '-t', 'gcr.io/$PROJECT_ID/app-name', '.']
- name: 'gcr.io/cloud-builders/docker'
  args: ['push', 'gcr.io/$PROJECT_ID/app-name']
```


- E. In the application source directory, create a file named `cloudbuild.yaml` that contains the following contents:

```
steps:
- name: 'gcr.io/cloud-builders/gcloud'
  args: ['app', 'deploy']
  timeout: '1600s'
```

Answer: AD

Explanation:

<https://cloud.google.com/build/docs/building/build-containers#store-images>

QUESTION 16

You are developing an internal application that will allow employees to organize community events within your company. You deployed your application on a single Compute Engine instance. Your company uses Google Workspace (formerly G Suite), and you need to ensure that the company employees can authenticate to the application from anywhere. What should you do?

- A. Add a public IP address to your instance, and restrict access to the instance using firewall rules. Allow your company's proxy as the only source IP address.
- B. Add an HTTP(S) load balancer in front of the instance, and set up Identity-Aware Proxy (IAP). Configure the IAP settings to allow your company domain to access the website.
- C. Set up a VPN tunnel between your company network and your instance's VPC location on Google Cloud. Configure the required firewall rules and routing information to both the on-premises and Google Cloud networks.
- D. Add a public IP address to your instance, and allow traffic from the internet. Generate a random hash, and create a subdomain that includes this hash and points to your instance. Distribute this DNS address to your company's employees.

Answer: B

Explanation:

https://cloud.google.com/iap/docs/concepts-overview#how_iap_works

When an application or resource is protected by IAP, it can only be accessed through the proxy by principals, also known as users, who have the correct Identity and Access Management (IAM) role. When you grant a user access to an application or resource by IAP, they're subject to the fine-grained access controls implemented by the product in use without requiring a VPN. When a user tries to access an IAP-secured resource, IAP performs authentication and authorization checks.

QUESTION 17

Your development team is using Cloud Build to promote a Node.js application built on App Engine from your staging environment to production. The application relies on several directories of photos stored in a Cloud Storage bucket named `webphotos-staging` in the staging environment. After the promotion, these photos must be available in a Cloud Storage bucket named `webphotos-prod` in the production environment. You want to automate the process where possible. What should you do?

- A. Manually copy the photos to `webphotos-prod`.
- B. Add a startup script in the application's `app.yaml` file to move the photos from `webphotos-staging` to `webphotos-prod`.
- C. Add a build step in the `cloudbuild.yaml` file before the promotion step with the arguments:


```
- name: gcr.io/cloud-builders/gsutil
  args: ['cp', '-r', 'gs://webphotos-staging',
        'gs://webphotos-prod']
  waitFor: ['-']
```

D. Add a build step in the cloudbuild.yaml file before the promotion step with the arguments:

```
- name: gcr.io/cloud-builders/gcloud
  args: ['cp', '-A', 'gs://webphotos-staging',
        'gs://webphotos-prod']
  waitFor: ['-']
```

Answer: C

Explanation:

You should add a build step in the cloudbuild.yaml file before the promotion step with the arguments shown above. This build step will use the gsutil tool to copy the photos from the webphotos-staging bucket to the webphotos-prod bucket. The -r flag tells gsutil to copy all files in the bucket recursively, and the waitFor parameter tells Cloud Build to wait for this step to complete before continuing with the promotion step.

QUESTION 18

You are developing a web application that will be accessible over both HTTP and HTTPS and will run on Compute Engine instances. On occasion, you will need to SSH from your remote laptop into one of the Compute Engine instances to conduct maintenance on the app. How should you configure the instances while following Google-recommended best practices?

- A. Set up a backend with Compute Engine web server instances with a private IP address behind a TCP proxy load balancer.
- B. Configure the firewall rules to allow all ingress traffic to connect to the Compute Engine web servers, with each server having a unique external IP address.
- C. Configure Cloud Identity-Aware Proxy API for SSH access. Then configure the Compute Engine servers with private IP addresses behind an HTTP(s) load balancer for the application web traffic.
- D. Set up a backend with Compute Engine web server instances with a private IP address behind an HTTP(S) load balancer. Set up a bastion host with a public IP address and open firewall ports. Connect to the web instances using the bastion host.

Answer: C

Explanation:

This document describes how to connect to a virtual machine (VM) instance through its internal IP address, using Identity-Aware Proxy (IAP) TCP forwarding.
<https://cloud.google.com/compute/docs/connect/ssh-using-iap>

QUESTION 19

You have a mixture of packaged and internally developed applications hosted on a Compute Engine instance that is running Linux. These applications write log records as text in local files. You want the logs to be written to Cloud Logging. What should you do?

- A. Pipe the content of the files to the Linux Syslog daemon.
- B. Install a Google version of fluentd on the Compute Engine instance.
- C. Install a Google version of collectd on the Compute Engine instance.
- D. Using cron, schedule a job to copy the log files to Cloud Storage once a day.

Answer: B

Explanation:

<https://cloud.google.com/stackdriver/docs/solutions/agents/ops-agent>

The Ops Agent is the primary agent for collecting telemetry from your Compute Engine instances. Combining logging and metrics into a single agent, the Ops Agent uses Fluent Bit for logs, which supports high-throughput logging, and the OpenTelemetry Collector for metrics.

QUESTION 20

You want to create `fully baked` or `golden` Compute Engine images for your application. You need to bootstrap your application to connect to the appropriate database according to the environment the application is running on (test, staging, production). What should you do?

- A. Embed the appropriate database connection string in the image. Create a different image for each environment.
- B. When creating the Compute Engine instance, add a tag with the name of the database to be connected. In your application, query the Compute Engine API to pull the tags for the current instance, and use the tag to construct the appropriate database connection string.
- C. When creating the Compute Engine instance, create a metadata item with a key of "DATABASE" and a value for the appropriate database connection string. In your application, read the "DATABASE" environment variable, and use the value to connect to the appropriate database.
- D. When creating the Compute Engine instance, create a metadata item with a key of "DATABASE" and a value for the appropriate database connection string. In your application, query the metadata server for the "DATABASE" value, and use the value to connect to the appropriate database.

Answer: D

Explanation:

<https://cloud.google.com/compute/docs/metadata/querying-metadata>

QUESTION 21

You are developing a microservice-based application that will be deployed on a Google Kubernetes Engine cluster. The application needs to read and write to a Spanner database. You want to follow security best practices while minimizing code changes. How should you configure your application to retrieve Spanner credentials?

- A. Configure the appropriate service accounts, and use Workload Identity to run the pods.
- B. Store the application credentials as Kubernetes Secrets, and expose them as environment variables.
- C. Configure the appropriate routing rules, and use a VPC-native cluster to directly connect to the database.
- D. Store the application credentials using Cloud Key Management Service, and retrieve them whenever a database connection is made.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/containers-kubernetes/introducing-workload-identity-better-authentication-for-your-gke-applications>

A Cloud IAM service account is an identity that an application can use to make requests to Google APIs. As an application developer, you could generate individual IAM service accounts for each application, and then download and store the keys as a Kubernetes secret that you manually rotate. Not only is this process burdensome, but service account keys only expire every 10 years (or until you manually rotate them). In the case of a breach or compromise, an

unaccounted-for key could mean prolonged access for an attacker. This potential blind spot, plus the management overhead of key inventory and rotation, makes using service account keys as secrets a less than ideal method for authenticating GKE workloads.

QUESTION 22

You are deploying your application on a Compute Engine instance that communicates with Cloud SQL. You will use Cloud SQL Proxy to allow your application to communicate to the database using the service account associated with the application's instance. You want to follow the Google-recommended best practice of providing minimum access for the role assigned to the service account. What should you do?

- A. Assign the Project Editor role.
- B. Assign the Project Owner role.
- C. Assign the Cloud SQL Client role.
- D. Assign the Cloud SQL Editor role.

Answer: C

Explanation:

<https://cloud.google.com/sql/docs/mysql/roles-and-permissions>

QUESTION 23

Your team develops stateless services that run on Google Kubernetes Engine (GKE). You need to deploy a new service that will only be accessed by other services running in the GKE cluster. The service will need to scale as quickly as possible to respond to changing load. What should you do?

- A. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
- B. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a NodePort Service.
- C. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
- D. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

Answer: C

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/service#services_of_type_clusterip

When you create a Service of type ClusterIP, Kubernetes creates a stable IP address that is accessible from nodes in the cluster.

<https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler>

The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload's CPU or memory consumption, or in response to custom metrics reported from within Kubernetes or external metrics from sources outside of your cluster.

QUESTION 24

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.

Answer: C

Explanation:

<https://cloud.google.com/endpoints/docs/openapi/versioning-an-api#backwards-incompatible>

When you make changes to your API that breaks your customers' client code, as a best practice, increment the major version number of your API. Endpoints can run more than one major version of an API concurrently. By providing both versions of the API, your customers can pick which version they want to use and control when they migrate to the new version.

QUESTION 25

You recently developed an application. You need to call the Cloud Storage API from a Compute Engine instance that doesn't have a public IP address. What should you do?

- A. Use Carrier Peering
- B. Use VPC Network Peering
- C. Use Shared VPC networks
- D. Use Private Google Access

Answer: D

Explanation:

<https://cloud.google.com/vpc/docs/private-google-access>

VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the external IP addresses of Google APIs and services. The source IP address of the packet can be the primary internal IP address of the network interface or an address in an alias IP range that is assigned to the interface. If you disable Private Google Access, the VM instances can no longer reach Google APIs and services; they can only send traffic within the VPC network.

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



Microsoft



ORACLE



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14