



**Free Questions for Professional-Cloud-Developer by
go4braindumps**

Shared by Farrell on 29-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You recently deployed a Go application on Google Kubernetes Engine (GKE). The operations team has noticed that the application's CPU usage is high even when there is low production traffic. The operations team has asked you to optimize your application's CPU resource consumption. You want to determine which Go functions consume the largest amount of CPU. What should you do?

Options:

- A-** Deploy a Fluent Bit daemonset on the GKE cluster to log data in Cloud Logging. Analyze the logs to get insights into your application code's performance.
- B-** Create a custom dashboard in Cloud Monitoring to evaluate the CPU performance metrics of your application.
- C-** Connect to your GKE nodes using SSH. Run the top command on the shell to extract the CPU utilization of your application.
- D-** Modify your Go application to capture profiling data. Analyze the CPU metrics of your application in flame graphs in Profiler.

Answer:

D

Explanation:

<https://cloud.google.com/profiler/docs/about-profiler>

Cloud Profiler is a statistical, low-overhead profiler that continuously gathers CPU usage and memory-allocation information from your production applications. It attributes that information to the source code that generated it, helping you identify the parts of your application that are consuming the most resources, and otherwise illuminating your applications performance characteristics.

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

Question 2

Question Type: MultipleChoice

You recently joined a new team that has a Cloud Spanner database instance running in production. Your manager has asked you to optimize the Spanner instance to reduce cost while maintaining high reliability and availability of the database. What should you do?

Options:

- A-** Use Cloud Logging to check for error logs, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- B-** Use Cloud Trace to monitor the requests per sec of incoming requests to Spanner, and reduce Spanner processing units by small increments until you find the minimum capacity required.

- C-** Use Cloud Monitoring to monitor the CPU utilization, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- D-** Use Snapshot Debugger to check for application errors, and reduce Spanner processing units by small increments until you find the minimum capacity required.

Answer:

C

Explanation:

https://cloud.google.com/spanner/docs/compute-capacity#increasing_and_decreasing_compute_capacity

Question 3

Question Type: MultipleChoice

You work at a rapidly growing financial technology startup. You manage the payment processing application written in Go and hosted on Cloud Run in the Singapore region (asia-southeast1). The payment processing application processes data stored in a Cloud Storage bucket that is also located in the Singapore region.

The startup plans to expand further into the Asia Pacific region. You plan to deploy the Payment Gateway in Jakarta, Hong Kong, and Taiwan over the next six months. Each location has data residency requirements that require customer data to reside in the country where the transaction was made. You want to minimize the cost of these deployments. What should you do?

Options:

- A-** Create a Cloud Storage bucket in each region, and create a Cloud Run service of the payment processing application in each region.
- B-** Create a Cloud Storage bucket in each region, and create three Cloud Run services of the payment processing application in the Singapore region.
- C-** Create three Cloud Storage buckets in the Asia multi-region, and create three Cloud Run services of the payment processing application in the Singapore region.
- D-** Create three Cloud Storage buckets in the Asia multi-region, and create three Cloud Run revisions of the payment processing application in the Singapore region.

Answer:

A

Question 4

Question Type: MultipleChoice

You are deploying a microservices application to Google Kubernetes Engine (GKE) that will broadcast livestreams. You expect unpredictable traffic patterns and large variations in the number of concurrent users. Your application must meet the following requirements:

- * Scales automatically during popular events and maintains high availability
- * Is resilient in the event of hardware failures

How should you configure the deployment parameters? (Choose two.)

Options:

- A-** Distribute your workload evenly using a multi-zonal node pool.
- B-** Distribute your workload evenly using multiple zonal node pools.
- C-** Use cluster autoscaler to resize the number of nodes in the node pool, and use a Horizontal Pod Autoscaler to scale the workload.
- D-** Create a managed instance group for Compute Engine with the cluster nodes. Configure autoscaling rules for the managed instance group.
- E-** Create alerting policies in Cloud Monitoring based on GKE CPU and memory utilization. Ask an on-duty engineer to scale the workload by executing a script when CPU and memory usage exceed predefined thresholds.

Answer:

A, C

Question 5

Question Type: MultipleChoice

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.

a. You want to maximize the performance of the retrieval process. How should you retrieve the metadata?

Options:

- A- Use the patch method.
- B- Use the compose method.
- C- Use the copy method.
- D- Use the fields request parameter.

Answer:

D

Explanation:

Question 6

Question Type: MultipleChoice

Your company uses Cloud Logging to manage large volumes of log data

a. You need to build a real-time log analysis architecture that pushes logs to a third-party application for processing. What should you do?

Options:

- A- Create a Cloud Logging log export to Pub/Sub.
- B- Create a Cloud Logging log export to BigQuery.
- C- Create a Cloud Logging log export to Cloud Storage.
- D- Create a Cloud Function to read Cloud Logging log entries and send them to the third-party application.

Answer:

A

Question 7

Question Type: MultipleChoice

You need to deploy resources from your laptop to Google Cloud using Terraform. Resources in your Google Cloud environment must be created using a service account. Your Cloud Identity has the roles/iam.serviceAccountTokenCreator Identity and Access Management (IAM) role and the necessary permissions to deploy the resources using Terraform. You want to set up your development environment to deploy the desired resources following Google-recommended best practices. What should you do?

Options:

- A-** 1) Download the service account's key file in JSON format, and store it locally on your laptop.
2) Set the `GOOGLE_APPLICATION_CREDENTIALS` environment variable to the path of your downloaded key file.
- B-** 1) Run the following command from a command line: `gcloud config set auth/impersonate_service_account service-account-name@project.iam.gserviceaccount.com`.
2) Set the `GOOGLE_OAUTH_ACCESS_TOKEN` environment variable to the value that is returned by the `gcloud auth print-access-token` command.
- C-** 1) Run the following command from a command line: `gcloud auth application-default login`.
2) In the browser window that opens, authenticate using your personal credentials.
- D-** 1) Store the service account's key file in JSON format in Hashicorp Vault.

2) Integrate Terraform with Vault to retrieve the key file dynamically, and authenticate to Vault using a short-lived access token.

Answer:

D

Explanation:

<https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#file-system>

Whenever possible, avoid storing service account keys on a file system. If you can't avoid storing keys on disk, make sure to restrict access to the key file, configure file access auditing, and encrypt the underlying disk.

<https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#software-keystore>

In situations where using a hardware-based key store isn't viable, use a software-based key store to manage service account keys. Similar to hardware-based options, a software-based key store lets users or applications use service account keys without revealing the private key. Software-based key store solutions can help you control key access in a fine-grained manner and can also ensure that each key access is logged.

Question 8

Question Type: MultipleChoice

You are using Cloud Run to host a web application. You need to securely obtain the application project ID and region where the application is running and display this information to users. You want to use the most performant approach. What should you do?

Options:

- A-** Use HTTP requests to query the available metadata server at the `http://metadata.google.internal/` endpoint with the `Metadata-Flavor: Google` header.
- B-** In the Google Cloud console, navigate to the Project Dashboard and gather configuration details. Navigate to the Cloud Run "Variables & Secrets" tab, and add the desired environment variables in Key:Value format.
- C-** In the Google Cloud console, navigate to the Project Dashboard and gather configuration details. Write the application configuration information to Cloud Run's in-memory container filesystem.
- D-** Make an API call to the Cloud Asset Inventory API from the application and format the request to include instance metadata.

Answer:

B

Question 9

Question Type: MultipleChoice

You have an application deployed in Google Kubernetes Engine (GKE) that reads and processes Pub/Sub messages. Each Pod handles a fixed number of messages per minute. The rate at which messages are published to the Pub/Sub topic varies considerably throughout the day and week, including occasional large batches of messages published at a single moment.

You want to scale your GKE Deployment to be able to process messages in a timely manner. What GKE feature should you use to automatically adapt your workload?

Options:

- A- Vertical Pod Autoscaler in Auto mode
- B- Vertical Pod Autoscaler in Recommendation mode
- C- Horizontal Pod Autoscaler based on an external metric
- D- Horizontal Pod Autoscaler based on resources utilization

Answer:

C

Explanation:

<https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/>

Question 10

Question Type: MultipleChoice

You are planning to add unit tests to your application. You need to be able to assert that published Pub/Sub messages are processed by your subscriber in order. You want the unit tests to be cost-effective and reliable. What should you do?

Options:

- A- Implement a mocking framework.
- B- Create a topic and subscription for each tester.
- C- Add a filter by tester to the subscription.
- D- Use the Pub/Sub emulator.

Answer:

D

Explanation:

<https://cloud.google.com/pubsub/docs/emulator>, 'Testing apps locally with the emulator'.

Question 11

Question Type: MultipleChoice

Your team is responsible for maintaining an application that aggregates news articles from many different sources. Your monitoring dashboard contains publicly accessible real-time reports and runs on a Compute Engine instance as a web application. External stakeholders and analysts need to access these reports via a secure channel without authentication. How should you configure this secure channel?

Options:

- A-** Add a public IP address to the instance. Use the service account key of the instance to encrypt the traffic.
- B-** Use Cloud Scheduler to trigger Cloud Build every hour to create an export from the reports. Store the reports in a public Cloud Storage bucket.
- C-** Add an HTTP(S) load balancer in front of the monitoring dashboard. Configure Identity-Aware Proxy to secure the communication channel.
- D-** Add an HTTP(S) load balancer in front of the monitoring dashboard. Set up a Google-managed SSL certificate on the load balancer for traffic encryption.

Answer:

D

Explanation:

<https://cloud.google.com/load-balancing/docs/ssl-certificates/google-managed-certs>

To Get Premium Files for Professional-Cloud-Developer Visit

<https://www.p2pexams.com/products/professional-cloud-developer>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-cloud-developer>

