Google Certified Professional - Cloud Developer

1. You are responsible for deploying a new API. That API will have three different URL paths:

A. Create one Cloud Function as a backend service exposed using an HTTPS load balancer.

B. Create three Cloud Functions exposed directly.

C. Create one Cloud Function exposed directly.

D. Create three Cloud Functions as three backend services exposed using an HTTPS load balancer.

Answer(s): D

---

2. Your team is developing a Cloud Function triggered by Cloud Storage Events. You want to accelerate testing and development of your Cloud Function while following Google-recommended best practices. What should you do?

A. Create a new Cloud Function that is triggered when Cloud Audit Logs detects the cloudfunctions. functions. sourceCodeSet operation in the original Cloud Function Send mock requests to the new function to evaluate the functionality.

B. Make a copy of the Cloud Function in the Google Cloud Console Use the Cloud console's in-line editor to make source code changes to the new function Modify your web application to call the new function and test the new version in production.

C. Make a copy of the Cloud Function, and rewrite the code to be HTTP-triggered Edit and test the new version by triggering the HTTP endpoint. Send mock requests to the new function to evaluate the functionality.

D. Install the Functions Frameworks library, and configure the Cloud Function on localhost. Make a copy of the function, and make edits to the new version Test the new version using cur1.

Answer(s): D

---

3. Your application stores customers' content in a Cloud Storage bucket, with each object being encrypted with the customer's encryption key. The key for each object in Cloud Storage is entered into your application by the customer. You discover that your application is receiving an HTTP 4xx error when reading the object from Cloud Storage What is a possible cause of this error?

A. You entered the same encryption algorithm specified by the customer when attempting the read operation.

B. You attempted the read operation on the object with the customers base64-encoded key.

C. You attempted the read operation on the object with the base64-encoded SHA256 hash of the customer's key.

D. You attempted the read operation without the base64-encoded SHA256 hash of the encryption key.

Answer(s): B

---

4. You are designing an application that uses a microservices architecture. You are planning to deploy the application in the cloud and on-premises. You want to make sure the application can scale up on demand and also use managed services as much as possible. What should you do?

A. Deploy open source Istio in a multi-cluster deployment on multiple Google Kubernetes Engine (GKE) clusters managed by Anthos.

B. Create a GKE cluster in each environment with Anthos, and use Cloud Run for Anthos to deploy your application to each cluster.

C. Install a GKE cluster in each environment with Anthos, and use Cloud Build to create a Deployment for your application in each cluster.

D. Create a GKE cluster in the cloud and install open-source Kubernetes on-premises. Use an external load balancer service to distribute traffic across the two environments.

Answer(s): B

---

5. You are monitoring a web application that is written in Go and deployed in Google Kubernetes Engine. You notice an increase in CPU and memory utilization. You need to determine which source code is consuming the most CPU and memory resources. What should you do?

A. Download, install, and start the Snapshot Debugger agent in your VM. Take debug snapshots of the functions that take the longest time. Review the call stack frame, and identify the local variables at that level in the stack.

B. Import OpenTelemetry and Trace export packages into your application, and create the trace provider.Review the latency data for your application on the Trace overview page, and identify where bottlenecks are occurring.

C. Import the Cloud Profiler package into your application, and initialize the Profiler agent. Review the generated flame graph in the Google Cloud console to identify time-intensive functions.

D. Create a Cloud Logging query that gathers the web application's logs. Write a Python script that calculates the difference between the timestamps from the beginning and the end of the application's longest functions to identity time-intensive functions.

Answer(s): C

---

6. 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?

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

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

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

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

Answer(s): B

---

7. You have an application running on Google Kubernetes Engine (GKE). The application is currently using a logging library and is outputting to standard output You need to export the logs

to Cloud Logging, and you need the logs to include metadata about each request. You want to use the simplest method to accomplish this.

A. Change your application s logging library to the Cloud Logging library and configure your application to export logs to Cloud Logging

B. Install the Fluent Bit agent on each of your GKE nodes, and have the agent export all logs from /var/log.

C. Update your application to output logs in CSV format, and add the necessary metadata to the CSV.

D. Update your application to output logs in JSON format, and add the necessary metadata to the JSON

Answer(s): C

---

8. You are developing an online gaming platform as a microservices application on Google Kubernetes Engine (GKE). Users on social media are complaining about long loading times for certain URL requests to the application. You need to investigate performance bottlenecks in the application and identify. which HTTP requests have a significantly high latency span in user requests What should you do?

A. Configure GKE workload metrics using kubect1 Select all Pods to send their metrics to Cloud Monitoring. Create a custom dashboard of application metrics in Cloud Monitoring to determine performance bottlenecks of your GKE cluster

B. Update your microservices to log HTTP request methods and URL paths to STDOUT Use the logs router to send container logs to Cloud Logging Create filters in Cloud Logging to evaluate the latency of user requests across different methods and URL paths.

C. Instrument your microservices by installing the OpenTelemetry tracing package Update your application code to send traces to Trace for inspection and analysis Create an analysis report on Trace to analyze user requests

D. Install tcpdump on your GKE nodes. Run tcpdump to capture network traffic over an extended period of time to collect data Analyze the data files using Wireshark to determine the cause of high latency

Answer(s): A

---

9. 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.

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(s): C

---

10. You are developing an application hosted on Google Cloud that uses a MySQL relational database schema.

A. Configure Cloud SQL to host the database, and import the schema into Cloud SQL.

B. Deploy MySQL from the Google Cloud Marketplace to the database using a client, and import the schema.

C. Configure Bigtable to host the database, and import the data into Bigtable.

D. Configure Cloud Spanner to host the database, and import the schema into Cloud Spanner.

E. Configure Firestore to host the database, and import the data into Firestore.

Answer(s): A

---

11. You are deploying your application to a Compute Engine virtual machine instance. Your application is configured to write its log files to disk. You want to view the logs in Stackdriver Logging without changing the application code.

A. Use a Stackdriver Logging Library to log directly from the application to Stackdriver Logging.

B. Change the application to log to /var/log so that its logs are automatically sent to Stackdriver Logging.

C. Install the Stackdriver Logging Agent and configure it to send the application logs.

D. Provide the log file folder path in the metadata of the instance to configure it to send the application logs.

Answer(s): C

---

12. You are a developer at a financial institution You use Cloud Shell to interact with Google Cloud services. User data is currently stored on an ephemeral disk however a recently passed regulation mandates that you can no longer store sensitive information on an ephemeral disk. You need to implement a new storage solution for your user data You want to minimize code changes Where should you store your user data'?

A. Store user data on a persistent disk in a Compute Engine instance

B. Store user data in a Cloud Storage bucket

C. Store user data on a Cloud Shell home disk and log in at least every 120 days to prevent its deletion

D. Store user data m BigQuery tables

Answer(s): A

---

13. Your team is setting up a build pipeline for an application that will run in Google Kubernetes Engine (GKE).

A. Cloud Build. Artifact Registry and Binary Authorization

B. Google Cloud Deploy. Artifact Registry, and Google Cloud Armor

C. Cloud Build, Cloud Storage, and Binary Authorization

D. Google Cloud Deploy, Cloud Storage and Google Cloud Armor

Answer(s): B

---

14. 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.

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(s): D

---

15. Your company has deployed a new API to App Engine Standard environment. During testing, the API is not behaving as expected. You want to monitor the application over time to diagnose the problem within the application code without redeploying the application.

A. Stackdriver Trace

B. Stackdriver Monitoring

C. Stackdriver Debug Snapshots

D. Stackdriver Debug Logpoints

Answer(s): B

---

16. You are monitoring a web application that is written in Go and deployed in Google Kubemetes Engine. You notice an increase in CPU and memory utilization. You need to determine which function is consuming the most CPU and memory resources. What should you do?

A. Import the Cloud Profiler package into your application, and initialize the Profiler agent. Review the generated flame graph in the Google Cloud console to identify time-intensive functions.

B. Create a Cloud Logging query that gathers the web application's logs. Write a Python script that calculates the difference between the timestamps from the beginning and the end of the application's longest functions to identity time-intensive functions.

C. Import OpenTelemetry and Trace export packages into your application, and create the trace provider.Review the latency data for yourapplication on the Trace overview page, and identify which functions cause the most latercy.

D. Add print commands to the application source code to log when each function is called, and redeploy the application.

Answer(s): A

---

17. For this question, refer to the HipLocal case study.

A. Create health checks for the QA environment, and redeploy the APIs at a later time if the environment is unhealthy.

B. Include performance tests in their code, and prevent deployments to QA until all tests have a passing status.

C. Include unit tests in their code, and prevent deployments to QA until all tests have a passing status.

D. Redeploy the APIs to App Engine using Traffic Splitting. Do not move QA traffic to the new versions if errors are found.

Answer(s): B

---

18. 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, query the metadata server for the "DATABASE" value, 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, read the "DATABASE" environment variable, and use the value to connect to the appropriate database.

Answer(s): C

---

19. You have a container deployed on Google Kubernetes Engine. The container can sometimes be slow to launch, so you have implemented a liveness probe. You notice that the liveness probe occasionally fails on launch. What should you do?

A. Add a startup probe.

B. Increase the initial delay for the liveness probe.

C. Increase the CPU limit for the container.

D. Add a readiness probe.

Answer(s): B

---

20. You are developing an online gaming platform as a microservices application on Google Kubernetes Engine (GKE). Users on social media are complaining about long loading times for certain URL requests to the application. You need to investigate performance bottlenecks in the application and identify which HTTP requests have a significantly high latency span in user requests. What should you do9

A. Install tcpdiimp on your GKE nodes. Run tcpdunm-- to capture network traffic over an extended period of time to collect data. Analyze the data files using Wireshark to determine the cause of high latency

B. Update your microservices lo log HTTP request methods and URL paths to STDOUT Use the logs router to send container logs to Cloud Logging. Create fillers in Cloud Logging to evaluate the latency of user requests across different methods and URL paths.

C. Configure GKE workload metrics using kubect1. Select all Pods to send their metrics to Cloud Monitoring. Create a custom dashboard of application metrics in Cloud Monitoring to determine

performance bottlenecks of your GKE cluster.

D. Instrument your microservices by installing the Open Telemetry tracing package. Update your application code to send traces to Trace for inspection and analysis. Create an analysis report on Trace to analyze user requests

Answer(s): C
