



Google

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

NEW QUESTION 1

- (Exam Topic 1)

Which database should HipLocal use for storing user activity?

- A. BigQuery
- B. Cloud SQL
- C. Cloud Spanner
- D. Cloud Datastore

Answer: A

NEW QUESTION 2

- (Exam Topic 1)

HipLocal wants to improve the resilience of their MySQL deployment, while also meeting their business and technical requirements. Which configuration should they choose?

- A. Use the current single instance MySQL on Compute Engine and several read-only MySQL servers on Compute Engine.
- B. Use the current single instance MySQL on Compute Engine, and replicate the data to Cloud SQL in an external master configuration.
- C. Replace the current single instance MySQL instance with Cloud SQL, and configure high availability.
- D. Replace the current single instance MySQL instance with Cloud SQL, and Google provides redundancy without further configuration.

Answer: B

NEW QUESTION 3

- (Exam Topic 1)

For this question refer to the HipLocal case study.

HipLocal wants to reduce the latency of their services for users in global locations. They have created read replicas of their database in locations where their users reside and configured their service to read traffic using those replicas. How should they further reduce latency for all database interactions with the least amount of effort?

- A. Migrate the database to Bigtable and use it to serve all global user traffic.
- B. Migrate the database to Cloud Spanner and use it to serve all global user traffic.
- C. Migrate the database to Firestore in Datastore mode and use it to serve all global user traffic.
- D. Migrate the services to Google Kubernetes Engine and use a load balancer service to better scale the application.

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

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.

Answer: B

Explanation:

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

NEW QUESTION 5

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

How should HipLocal redesign their architecture to ensure that the application scales to support a large increase in users?

- A. Use Google Kubernetes Engine (GKE) to run the application as a microservice
- B. Run the MySQL database on a dedicated GKE node.
- C. Use multiple Compute Engine instances to run MySQL to store state information
- D. Use a GoogleCloud-managed load balancer to distribute the load between instances
- E. Use managed instance groups for scaling.
- F. Use Memorystore to store session information and CloudSQL to store state information
- G. Use a Google Cloud-managed load balancer to distribute the load between instances
- H. Use managed instance groups for scaling.
- I. Use a Cloud Storage bucket to serve the application as a static website, and use another Cloud Storage bucket to store user state information.

Answer: D

NEW QUESTION 6

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

How should HipLocal increase their API development speed while continuing to provide the QA team with a stable testing environment that meets feature requirements?

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

- C. Create health checks for the QA environment, and redeploy the APIs at a later time if the environment is unhealthy.
- D. Redeploy the APIs to App Engine using Traffic Splittin
- E. Do not move QA traffic to the new versions if errors are found.

Answer: B

NEW QUESTION 7

- (Exam Topic 1)

Which service should HipLocal use for their public APIs?

- A. Cloud Armor
- B. Cloud Functions
- C. Cloud Endpoints
- D. Shielded Virtual Machines

Answer: D

NEW QUESTION 8

- (Exam Topic 1)

HipLocal wants to reduce the number of on-call engineers and eliminate manual scaling. Which two services should they choose? (Choose two.)

- A. Use Google App Engine services.
- B. Use serverless Google Cloud Functions.
- C. Use Knative to build and deploy serverless applications.
- D. Use Google Kubernetes Engine for automated deployments.
- E. Use a large Google Compute Engine cluster for deployments.

Answer: BC

NEW QUESTION 9

- (Exam Topic 1)

HipLocal's APIs are showing occasional failures, but they cannot find a pattern. They want to collect some metrics to help them troubleshoot. What should they do?

- A. Take frequent snapshots of all of the VMs.
- B. Install the Stackdriver Logging agent on the VMs.
- C. Install the Stackdriver Monitoring agent on the VMs.
- D. Use Stackdriver Trace to look for performance bottlenecks.

Answer: C

NEW QUESTION 10

- (Exam Topic 2)

Your company needs a database solution that stores customer purchase history and meets the following requirements:

Customers can query their purchase immediately after submission. Purchases can be sorted on a variety of fields. Distinct record formats can be stored at the same time. Which storage option satisfies these requirements?

- A. Firestore in Native mode
- B. Cloud Storage using an object read
- C. Cloud SQL using a SQL SELECT statement
- D. Firestore in Datastore mode using a global query

Answer: A

NEW QUESTION 10

- (Exam Topic 2)

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

Answer: D

Explanation:

Reference:

<https://cloud.google.com/blog/products/application-development/least-privilege-for-cloud-functions-using-cloud>

NEW QUESTION 11

- (Exam Topic 2)

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 SQ
- B. Enable private services access.3) Create a Serverless VPC Access connector.4) Configure Cloud Run to use the connector to connect to Cloud SQL.
- C. 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 V
- D. 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.
- E. 1) Use Cloud SQL and Cloud Run in different projects.2) Configure a private IP address for Cloud SQ
- F. Enable private services access.3) Create a Serverless VPC Access connector.4) Set up a VPN connection between the two project
- G. Configure Cloud Run to use the connector to connect to Cloud SQL.
- H. 1) Install PostgreSQL on a Compute Engine VM, and enable Cloud Run in different projects.2) Configure a private IP address for the V
- I. Enable private services access.3) Create a Serverless VPC Access connector.4) Set up a VPN connection between the two project
- J. Configure Cloud Run to use the connector to access the VM hosting PostgreSQL

Answer: A

Explanation:

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

NEW QUESTION 15

- (Exam Topic 2)

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. Which tool should you use?

- A. Stackdriver Trace
- B. Stackdriver Monitoring
- C. Stackdriver Debug Snapshots
- D. Stackdriver Debug Logpoints

Answer: B

Explanation:

Reference: <https://rominirani.com/gcp-stackdriver-tutorial-debug-snapshots-traces-logging-and-logpoints-1ba49e4780e6>

NEW QUESTION 18

- (Exam Topic 2)

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?

- 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".

NEW QUESTION 22

- (Exam Topic 2)

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
- B. Create a different image for each environment.
- C. 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.
- 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
- E. In your application, read the "DATABASE" environment variable, and use the value to connect to the appropriate database.
- F. When creating the Compute Engine instance, create a metadata item with a key of "DATABASE" and a value for the appropriate database connection string
- G. In your application, query the metadata server for the "DATABASE" value, and use the value to connect to the appropriate database.

Answer: C

NEW QUESTION 24

- (Exam Topic 2)

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
- B. Deploy the application to Cloud Run.
- C. Build the application in C# with Firestore as the database
- D. Deploy the application to App Engine flexible environment.
- E. Build the application in Python with CloudSQL as the database
- F. Deploy the application to App Engine standard environment.

- G. Build the application in Python with Firestore as the databas
- H. Deploy the application to a Compute Engine managed instance group with autoscaling.

Answer: A

NEW QUESTION 29

- (Exam Topic 2)

Your application takes an input from a user and publishes it to the user's contacts. This input is stored in a table in Cloud Spanner. Your application is more sensitive to latency and less sensitive to consistency. How should you perform reads from Cloud Spanner for this application?

- A. Perform Read-Only transactions.
- B. Perform stale reads using single-read methods.
- C. Perform strong reads using single-read methods.
- D. Perform stale reads using read-write transactions.

Answer: D

Explanation:

Reference: <https://cloud.google.com/solutions/best-practices-cloud-spanner-gaming-database>

NEW QUESTION 31

- (Exam Topic 2)

Your company's corporate policy states that there must be a copyright comment at the very beginning of all source files. You want to write a custom step in Cloud Build that is triggered by each source commit. You need the trigger to validate that the source contains a copyright and add one for subsequent steps if not there. What should you do?

- A. Build a new Docker container that examines the files in /workspace and then checks and adds a copyright for each source fil
- B. Changed files are explicitly committed back to the source repository.
- C. Build a new Docker container that examines the files in /workspace and then checks and adds a copyright for each source fil
- D. Changed files do not need to be committed back to the source repository.
- E. Build a new Docker container that examines the files in a Cloud Storage bucket and then checks and adds a copyright for each source fil
- F. Changed files are written back to the Cloud Storage bucket.
- G. Build a new Docker container that examines the files in a Cloud Storage bucket and then checks and adds a copyright for each source fil
- H. Changed files are explicitly committed back to the source repository.

Answer: A

Explanation:

https://cloud.google.com/build/docs/configuring-builds/pass-data-between-steps#passing_data_using_workspace To pass data between build steps, store the assets produced by the build step in /workspace and these assets will be available to any subsequent build steps.

NEW QUESTION 34

- (Exam Topic 2)

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:

Reference: <https://cloud.google.com/logging/docs/agent/logging/configuration>

NEW QUESTION 35

- (Exam Topic 2)

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.

Answer: C

Explanation:

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

NEW QUESTION 37

- (Exam Topic 2)

Your company is planning to migrate their on-premises Hadoop environment to the cloud. Increasing storage cost and maintenance of data stored in HDFS is a major concern for your company. You also want to make minimal changes to existing data analytics jobs and existing architecture. How should you proceed with the migration?

- A. Migrate your data stored in Hadoop to BigQuer
- B. Change your jobs to source their information from BigQuery instead of the on-premises Hadoop environment.
- C. Create Compute Engine instances with HDD instead of SSD to save cost
- D. Then perform a full migration of your existing environment into the new one in Compute Engine instances.
- E. Create a Cloud Dataproc cluster on Google Cloud Platform, and then migrate your Hadoop environment to the new Cloud Dataproc cluste
- F. Move your HDFS data into larger HDD disks to save on storage costs.
- G. Create a Cloud Dataproc cluster on Google Cloud Platform, and then migrate your Hadoop code objects to the new cluste
- H. Move your data to Cloud Storage and leverage the Cloud Dataproc connector to run jobs on that data.

Answer: D

NEW QUESTION 38

- (Exam Topic 2)

You are running an application on App Engine that you inherited. You want to find out whether the application is using insecure binaries or is vulnerable to XSS attacks. Which service should you use?

- A. Cloud Amor
- B. Stackdriver Debugger
- C. Cloud Security Scanner
- D. Stackdriver Error Reporting

Answer: C

Explanation:

Reference: <https://cloud.google.com/security-scanner>

NEW QUESTION 42

- (Exam Topic 2)

Your application requires service accounts to be authenticated to GCP products via credentials stored on its host Compute Engine virtual machine instances. You want to distribute these credentials to the host instances as securely as possible. What should you do?

- A. Use HTTP signed URLs to securely provide access to the required resources.
- B. Use the instance's service account Application Default Credentials to authenticate to the required resources.
- C. Generate a P12 file from the GCP Console after the instance is deployed, and copy the credentials to the host instance before starting the application.
- D. Commit the credential JSON file into your application's source repository, and have your CI/CD process package it with the software that is deployed to the instance.

Answer: B

Explanation:

Reference: <https://cloud.google.com/compute/docs/api/how-tos/authorization>

NEW QUESTION 47

- (Exam Topic 2)

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 developmen
- D. Send your source code updates as pull requests.
- E. Use a Cloud Storage bucket to store the source code that you need to edi
- F. Mount the bucket to a public computer as a drive, and use a code editor to update the cod
- G. Turn on versioning for the bucket, and point it to the team's Git repository.

Answer: C

Explanation:

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

NEW QUESTION 49

- (Exam Topic 2)

You are developing an application that needs to store files belonging to users in Cloud Storage. You want each user to have their own subdirectory in Cloud Storage. When a new user is created, the corresponding empty subdirectory should also be created. What should you do?

- A. Create an object with the name of the subdirectory ending with a trailing slash ('/') that is zero bytes in length.
- B. Create an object with the name of the subdirectory, and then immediately delete the object within that subdirectory.
- C. Create an object with the name of the subdirectory that is zero bytes in length and has WRITER access control list permission.
- D. Create an object with the name of the subdirectory that is zero bytes in lengt
- E. Set the Content-Type metadata to CLOUDSTORAGE_FOLDER.

Answer: A

Explanation:

<https://cloud.google.com/storage/docs/folders>

If you create an empty folder using the Google Cloud console, Cloud Storage creates a zero-byte object as a placeholder. For example, if you create a folder called folder in a bucket called my-bucket, a zero- byte object called gs://my-bucket/folder/ is created. This placeholder is discoverable by other tools when listing

the objects in the bucket, for example when using the `gsutil ls` command.

NEW QUESTION 52

- (Exam Topic 2)

You configured your Compute Engine instance group to scale automatically according to overall CPU usage. However, your application's response latency increases sharply before the cluster has finished adding up instances. You want to provide a more consistent latency experience for your end users by changing the configuration of the instance group autoscaler. Which two configuration changes should you make? (Choose two.)

- A. Add the label "AUTOSCALE" to the instance group template.
- B. Decrease the cool-down period for instances added to the group.
- C. Increase the target CPU usage for the instance group autoscaler.
- D. Decrease the target CPU usage for the instance group autoscaler.
- E. Remove the health-check for individual VMs in the instance group.

Answer: AC

NEW QUESTION 57

- (Exam Topic 2)

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.

Answer: C

NEW QUESTION 58

- (Exam Topic 2)

Your company's product team has a new requirement based on customer demand to autoscale your stateless and distributed service running in a Google Kubernetes Engine (GKE) cluster. You want to find a solution that minimizes changes because this feature will go live in two weeks. What should you do?

- A. Deploy a Vertical Pod Autoscaler, and scale based on the CPU load.
- B. Deploy a Vertical Pod Autoscaler, and scale based on a custom metric.
- C. Deploy a Horizontal Pod Autoscaler, and scale based on the CPU load.
- D. Deploy a Horizontal Pod Autoscaler, and scale based on a custom metric.

Answer: C

Explanation:

<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.

NEW QUESTION 61

- (Exam Topic 2)

You are building a new API. You want to minimize the cost of storing and reduce the latency of serving images. Which architecture should you use?

- A. App Engine backed by Cloud Storage
- B. Compute Engine backed by Persistent Disk
- C. Transfer Appliance backed by Cloud Filestore
- D. Cloud Content Delivery Network (CDN) backed by Cloud Storage

Answer: B

NEW QUESTION 62

- (Exam Topic 2)

Your company has deployed a new API to a Compute Engine instance. During testing, the API is not behaving as expected. You want to monitor the application over 12 hours to diagnose the problem within the application code without redeploying the application. Which tool should you use?

- A. Cloud Trace
- B. Cloud Monitoring
- C. Cloud Debugger logpoints
- D. Cloud Debugger snapshots

Answer: C

Explanation:

<https://cloud.google.com/debugger/docs/using/logpoints>

Logpoints allow you to inject logging into running services without restarting or interfering with the normal function of the service

NEW QUESTION 65

- (Exam Topic 2)

Your security team is auditing all deployed applications running in Google Kubernetes Engine. After completing the audit, your team discovers that some of the applications send traffic within the cluster in clear text. You need to ensure that all application traffic is encrypted as quickly as possible while minimizing changes to your applications and maintaining support from Google. What should you do?

- A. Use Network Policies to block traffic between applications.
- B. Install Istio, enable proxy injection on your application namespace, and then enable mTLS.
- C. Define Trusted Network ranges within the application, and configure the applications to allow traffic only from those networks.
- D. Use an automated process to request SSL Certificates for your applications from Let's Encrypt and add them to your applications.

Answer: D

NEW QUESTION 66

- (Exam Topic 2)

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 serve
- C. Use your deployment tool to query the instance metadata and configure the application based on the "environment" value.
- D. Deploy your chosen deployment tool on an instance in each projec
- E. 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.
- F. During each instance launch, configure an instance custom-metadata key named "environment" whose value is the environment the instance serve
- G. Use your deployment tool to query the instance metadata, and configure the application based on the "environment" value.

Answer: B

Explanation:

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

NEW QUESTION 71

- (Exam Topic 2)

You are writing from a Go application to a Cloud Spanner database. You want to optimize your application's performance using Google-recommended best practices. What should you do?

- A. Write to Cloud Spanner using Cloud Client Libraries.
- B. Write to Cloud Spanner using Google API Client Libraries
- C. Write to Cloud Spanner using a custom gRPC client library.
- D. Write to Cloud Spanner using a third-party HTTP client library.

Answer: A

Explanation:

<https://cloud.google.com/apis/docs/cloud-client-libraries>

"Cloud Client Libraries are the recommended option for accessing Cloud APIs programmatically, where available. Cloud Client Libraries use the latest client library models"

<https://cloud.google.com/apis/docs/client-libraries-explained> <https://cloud.google.com/go/docs/reference>

NEW QUESTION 73

- (Exam Topic 2)

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. What should you do?

- A. Install the Stackdriver Logging Agent and configure it to send the application logs.
- B. Use a Stackdriver Logging Library to log directly from the application to Stackdriver Logging.
- C. Provide the log file folder path in the metadata of the instance to configure it to send the application logs.
- D. Change the application to log to `/var/log` so that its logs are automatically sent to Stackdriver Logging.

Answer: A

NEW QUESTION 74

- (Exam Topic 2)

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>

NEW QUESTION 76

- (Exam Topic 2)

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?

- 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: D

Explanation:

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

NEW QUESTION 80

- (Exam Topic 2)

You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. The application exposes an HTTP-based health check at /healthz. You want to use this health check endpoint to determine whether traffic should be routed to the pod by the load balancer.

Which code snippet should you include in your Pod configuration?

- A.


```
livenessProbe:
  httpGet:
    path: /healthz
    port: 80
```
- B.


```
readinessProbe:
  httpGet:
    path: /healthz
    port: 80
```
- C.


```
loadbalancerHealthCheck:
  httpGet:
    path: /healthz
    port: 80
```
- D.


```
healthCheck:
  httpGet:
    path: /healthz
    port: 80
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Explanation:

For the GKE ingress controller to use your readinessProbes as health checks, the Pods for an Ingress must exist at the time of Ingress creation. If your replicas are scaled to 0, the default health check will apply.

NEW QUESTION 82

- (Exam Topic 2)

You are developing a Java Web Server that needs to interact with Google Cloud services via the Google Cloud API on the user's behalf. Users should be able to authenticate to the Google Cloud API using their Google Cloud identities. Which workflow should you implement in your web application?

- A. 1) When a user arrives at your application, prompt them for their Google username and password. 2) Store an SHA password hash in your application's

database along with the user's username.3) The application authenticates to the Google Cloud API using HTTPs requests with the user's username and password hash in the Authorization request header.

B. 1) When a user arrives at your application, prompt them for their Google username and password.2) Forward the user's username and password in an HTTPS request to the Google Cloud authorization server, and request an access token.3) The Google server validates the user's credentials and returns an access token to the application.4) The application uses the access token to call the Google Cloud API.

C. 1) When a user arrives at your application, route them to a Google Cloud consent screen with a list of requested permissions that prompts the user to sign in with SSO to their Google Account.2) After the user signs in and provides consent, your application receives an authorization code from a Google server.3) The Google server returns the authorization code to the user, which is stored in the browser's cookies.4) The user authenticates to the Google Cloud API using the authorization code in the cookie.

D. 1) When a user arrives at your application, route them to a Google Cloud consent screen with a list of requested permissions that prompts the user to sign in with SSO to their Google Account.2) After the user signs in and provides consent, your application receives an authorization code from a Google server.3) The application requests a Google Server to exchange the authorization code with an access token.4) The Google server responds with the access token that is used by the application to call the Google Cloud API.

Answer: D

Explanation:

<https://developers.google.com/identity/protocols/oauth2#webserver>

The Google OAuth 2.0 endpoint supports web server applications that use languages and frameworks such as PHP, Java, Python, Ruby, and ASP.NET. The authorization sequence begins when your application redirects a browser to a Google URL; the URL includes query parameters that indicate the type of access being requested. Google handles the user authentication, session selection, and user consent. The result is an authorization code, which the application can exchange for an access token and a refresh token.

NEW QUESTION 84

- (Exam Topic 2)

Your application is deployed in a Google Kubernetes Engine (GKE) cluster. When a new version of your application is released, your CI/CD tool updates the spec.template.spec.containers[0].image value to reference the Docker image of your new application version. When the Deployment object applies the change, you want to deploy at least 1 replica of the new version and maintain the previous replicas until the new replica is healthy.

Which change should you make to the GKE Deployment object shown below?

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: ecommerce-frontend-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: ecommerce-frontend
  template:
    metadata:
      labels:
        app: ecommerce-frontend
    spec:
      containers:
      - name: ecommerce-frontend-webapp
        image: ecommerce-frontend-webapp:1.7.9
        ports:
        - containerPort: 80
```

- A. Set the Deployment strategy to RollingUpdate with maxSurge set to 0, maxUnavailable set to 1.
- B. Set the Deployment strategy to RollingUpdate with maxSurge set to 1, maxUnavailable set to 0.
- C. Set the Deployment strategy to Recreate with maxSurge set to 0, maxUnavailable set to 1.
- D. Set the Deployment strategy to Recreate with maxSurge set to 1, maxUnavailable set to 0.

Answer: D

NEW QUESTION 86

- (Exam Topic 2)

You work for an organization that manages an online ecommerce website. Your company plans to expand across the world; however, the estore currently serves one specific region. You need to select a SQL database and configure a schema that will scale as your organization grows. You want to create a table that stores all customer transactions and ensure that the customer (CustomerId) and the transaction (TransactionId) are unique. What should you do?

- A. Create a Cloud SQL table that has TransactionId and CustomerId configured as primary key
- B. Use an incremental number for the TransactionId.
- C. Create a Cloud SQL table that has TransactionId and CustomerId configured as primary key
- D. Use a random string (UUID) for the Transactionid.
- E. Create a Cloud Spanner table that has TransactionId and CustomerId configured as primary key

- F. Use a random string (UUID) for the TransactionId.
- G. Create a Cloud Spanner table that has TransactionId and CustomerId configured as primary key
- H. Use an incremental number for the TransactionId.

Answer: C

NEW QUESTION 91

- (Exam Topic 2)

You recently developed a new application. You want to deploy the application on Cloud Run without a Dockerfile. Your organization requires that all container images are pushed to a centrally managed container repository. How should you build your container using Google Cloud services? (Choose two.)

- A. Push your source code to Artifact Registry.
- B. Submit a Cloud Build job to push the image.
- C. Use the pack build command with pack CLI.
- D. Include the --source flag with the gcloud run deploy CLI command.
- E. Include the --platform=kubernetes flag with the gcloud run deploy CLI command.

Answer: AC

Explanation:

<https://cloud.google.com/run/docs/deploying#images> <https://cloud.google.com/blog/products/containers-kubernetes/google-cloud-now-supports-buildpacks>

NEW QUESTION 94

- (Exam Topic 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 Ru
- B. Use traffic splitting to direct a subset of user traffic to the new version based on the revision tag.
- C. Deploy your application on Google Kubernetes Engine with Anthos Service Mes
- D. Use traffic splitting to direct a subset of user traffic to the new version based on the user-agent header.
- E. Deploy your application on App Engin
- F. Use traffic splitting to direct a subset of user traffic to the new version based on the IP address.
- G. Deploy your application on Compute Engine
- H. Use Traffic Director to direct a subset of user traffic to the new version based on predefined weights.

Answer: B

NEW QUESTION 95

- (Exam Topic 2)

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 image
- B. Review the vulnerability results before each deployment.
- C. Enable Container Analysis, and upload new container images to Artifact Registr
- D. Review the vulnerability results before each deployment.
- E. Enable Container Analysis, and upload new container images to Artifact Registr
- F. Review the critical vulnerability results before each deployment.
- G. Use the Container Analysis REST API to call Container Analysis to scan new container image
- H. Review the vulnerability results before each deployment.

Answer: B

Explanation:

<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).

NEW QUESTION 98

- (Exam Topic 2)

You migrated some of your applications to Google Cloud. You are using a legacy monitoring platform deployed on-premises for both on-premises and cloud-deployed applications. You discover that your notification system is responding slowly to time-critical problems in the cloud applications. What should you do?

- A. Replace your monitoring platform with Cloud Monitoring.
- B. Install the Cloud Monitoring agent on your Compute Engine instances.
- C. Migrate some traffic back to your old platfor
- D. Perform A/B testing on the two platforms concurrently.
- E. Use Cloud Logging and Cloud Monitoring to capture logs, monitor, and send alert
- F. Send them to your existing platform.

Answer: D

NEW QUESTION 100

- (Exam Topic 2)

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?

- 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.

NEW QUESTION 102

- (Exam Topic 2)

Your teammate has asked you to review the code below, which is adding a credit to an account balance in Cloud Datastore. Which improvement should you suggest your teammate make?

```
public Entity creditAccount(long accountId, long
creditAmount) {
    Entity account = datastore.get
(keyFactory.newKey(accountId));
    account = Entity.builder(account).set(
        "balance", account.getLong("balance")
+ creditAmount).build()
    datastore.put(account);
    return account;
}
```

- A. Get the entity with an ancestor query.
- B. Get and put the entity in a transaction.
- C. Use a strongly consistent transactional database.
- D. Don't return the account entity from the function.

Answer: A

NEW QUESTION 104

- (Exam Topic 2)

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 V
- B. Take debug snapshots of the functions that take the longest tim
- C. Review the call stack frame, and identify the local variables at that level in the stack.
- D. Import the Cloud Profiler package into your application, and initialize the Profiler agen
- E. Review the generated flame graph in the Google Cloud console to identify time-intensive functions.
- F. 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.
- G. Create a Cloud Logging query that gathers the web application's log
- H. 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: B

NEW QUESTION 105

- (Exam Topic 2)

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

NEW QUESTION 108

- (Exam Topic 2)

One of your deployed applications in Google Kubernetes Engine (GKE) is having intermittent performance issues. Your team uses a third-party logging solution. You want to install this solution on each node in your GKE cluster so you can view the logs. What should you do?

- A. Deploy the third-party solution as a DaemonSet
- B. Modify your container image to include the monitoring software
- C. Use SSH to connect to the GKE node, and install the software manually
- D. Deploy the third-party solution using Terraform and deploy the logging Pod as a Kubernetes Deployment

Answer: A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset#usage_patterns DaemonSets are useful for deploying ongoing background tasks that you need to run on all or certain nodes, and which do not require user intervention. Examples of such tasks include storage daemons like ceph, log collection daemons like fluent-bit, and node monitoring daemons like collectd.

NEW QUESTION 112

- (Exam Topic 2)

You are developing a new web application using Cloud Run and committing code to Cloud Source Repositories. You want to deploy new code in the most efficient way possible. You have already created a Cloud Build YAML file that builds a container and runs the following command: `gcloud run deploy`. What should you do next?

- A. Create a Pub/Sub topic to be notified when code is pushed to the repository
- B. Create a Pub/Sub trigger that runs the build file when an event is published to the topic.
- C. Create a build trigger that runs the build file in response to a repository code being pushed to the development branch.
- D. Create a webhook build trigger that runs the build file in response to HTTP POST calls to the webhook URL.
- E. Create a Cron job that runs the following command every 24 hours: `gcloud builds submit`.

Answer: B

Explanation:

<https://cloud.google.com/build/docs/triggers>

Cloud Build uses build triggers to enable CI/CD automation. You can configure triggers to listen for incoming events, such as when a new commit is pushed to a repository or when a pull request is initiated, and then automatically execute a build when new events come in. You can also configure triggers to build code on any changes to your source repository or only on changes that match certain criteria.

NEW QUESTION 116

- (Exam Topic 2)

Your team develops services that run on Google Cloud. You need to build a data processing service and will use Cloud Functions. The data to be processed by the function is sensitive. You need to ensure that invocations can only happen from authorized services and follow Google-recommended best practices for securing functions. What should you do?

- A. Enable Identity-Aware Proxy in your projec
- B. Secure function access using its permissions.
- C. Create a service account with the Cloud Functions Viewer rol
- D. Use that service account to invoke the function.
- E. Create a service account with the Cloud Functions Invoker rol
- F. Use that service account to invoke the function.
- G. Create an OAuth 2.0 client ID for your calling service in the same project as the function you want to secur
- H. Use those credentials to invoke the function.

Answer: C

Explanation:

Reference:

<https://medium.com/google-cloud/how-to-securely-invoke-a-cloud-function-from-google-kubernetes-engine-run>

NEW QUESTION 117

- (Exam Topic 2)

Your team is developing an ecommerce platform for your company. Users will log in to the website and add items to their shopping cart. Users will be automatically logged out after 30 minutes of inactivity. When users log back in, their shopping cart should be saved. How should you store users' session and shopping cart information while following Google-recommended best practices?

- A. Store the session information in Pub/Sub, and store the shopping cart information in Cloud SQL.
- B. Store the shopping cart information in a file on Cloud Storage where the filename is the SESSION ID.
- C. Store the session and shopping cart information in a MySQL database running on multiple Compute Engine instances.
- D. Store the session information in Memorystore for Redis or Memorystore for Memcached, and store the shopping cart information in Firestore.

Answer: D

NEW QUESTION 121

- (Exam Topic 2)

You are evaluating developer tools to help drive Google Kubernetes Engine adoption and integration with your development environment, which includes VS Code and IntelliJ. What should you do?

- A. Use Cloud Code to develop applications.
- B. Use the Cloud Shell integrated Code Editor to edit code and configuration files.
- C. Use a Cloud Notebook instance to ingest and process data and deploy models.
- D. Use Cloud Shell to manage your infrastructure and applications from the command line.

Answer: A

Explanation:

Reference: <https://cloud.google.com/code>

NEW QUESTION 124

- (Exam Topic 2)

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 minute
- D. After the file is read, the application disconnects from the database.
- E. Write a shutdown script that publishes a message to the Pub/Sub topic announcing that a shutdown is in progres
- F. After the application reads the message, it disconnects from the database.

Answer: D

NEW QUESTION 126

- (Exam Topic 2)

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: ['-']
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/gsutil/commands/cp>

NEW QUESTION 130

- (Exam Topic 2)

You are using Cloud Build to create a new Docker image on each source code commit to a Cloud Source Repositories repository. Your application is built on every commit to the master branch. You want to release specific commits made to the master branch in an automated method. What should you do?

- A. Manually trigger the build for new releases.
- B. Create a build trigger on a Git tag patter
- C. Use a Git tag convention for new releases.
- D. Create a build trigger on a Git branch name patter
- E. Use a Git branch naming convention for new releases.
- F. Commit your source code to a second Cloud Source Repositories repository with a second Cloud Build trigge
- G. Use this repository for new releases only.

Answer: C

Explanation:

Reference: <https://docs.docker.com/docker-hub/builds/>

NEW QUESTION 135

- (Exam Topic 2)

Your application is running on Compute Engine and is showing sustained failures for a small number of requests. You have narrowed the cause down to a single

Compute Engine instance, but the instance is unresponsive to SSH. What should you do next?

- A. Reboot the machine.
- B. Enable and check the serial port output.
- C. Delete the machine and create a new one.
- D. Take a snapshot of the disk and attach it to a new machine.

Answer: A

NEW QUESTION 140

- (Exam Topic 2)

Your team develops services that run on Google Cloud. You want to process messages sent to a Pub/Sub topic, and then store them. Each message must be processed exactly once to avoid duplication of data and any data conflicts. You need to use the cheapest and most simple solution. What should you do?

- A. Process the messages with a Dataproc job, and write the output to storage.
- B. Process the messages with a Dataflow streaming pipeline using Apache Beam's PubSubIO package, and write the output to storage.
- C. Process the messages with a Cloud Function, and write the results to a BigQuery location where you can run a job to deduplicate the data.
- D. Retrieve the messages with a Dataflow streaming pipeline, store them in Cloud Bigtable, and use another Dataflow streaming pipeline to deduplicate messages.

Answer: B

Explanation:

<https://cloud.google.com/dataflow/docs/concepts/streaming-with-cloud-pubsub>

NEW QUESTION 143

- (Exam Topic 2)

You plan to make a simple HTML application available on the internet. This site keeps information about FAQs for your application. The application is static and contains images, HTML, CSS, and Javascript. You want to make this application available on the internet with as few steps as possible. What should you do?

- A. Upload your application to Cloud Storage.
- B. Upload your application to an App Engine environment.
- C. Create a Compute Engine instance with Apache web server installed.
- D. Configure Apache web server to host the application.
- E. Containerize your application first.
- F. Deploy this container to Google Kubernetes Engine (GKE) and assign an external IP address to the GKE pod hosting the application.

Answer: A

Explanation:

Reference: <https://cloud.google.com/storage/docs/hosting-static-website>

NEW QUESTION 147

- (Exam Topic 2)

You manage an ecommerce application that processes purchases from customers who can subsequently cancel or change those purchases. You discover that order volumes are highly variable and the backend order-processing system can only process one request at a time. You want to ensure seamless performance for customers regardless of usage volume. It is crucial that customers' order update requests are performed in the sequence in which they were generated. What should you do?

- A. Send the purchase and change requests over WebSockets to the backend.
- B. Send the purchase and change requests as REST requests to the backend.
- C. Use a Pub/Sub subscriber in pull mode and use a data store to manage ordering.
- D. Use a Pub/Sub subscriber in push mode and use a data store to manage ordering.

Answer: C

Explanation:

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

NEW QUESTION 151

- (Exam Topic 2)

You need to deploy a new European version of a website hosted on Google Kubernetes Engine. The current and new websites must be accessed via the same HTTP(S) load balancer's external IP address, but have different domain names. What should you do?

- A. Define a new Ingress resource with a host rule matching the new domain
- B. Modify the existing Ingress resource with a host rule matching the new domain
- C. Create a new Service of type LoadBalancer specifying the existing IP address as the loadBalancerIP
- D. Generate a new Ingress resource and specify the existing IP address as the `kubernetes.io/ingress.global-static-ip-name` annotation value

Answer: B

Explanation:

<https://kubernetes.io/docs/concepts/services-networking/ingress/#name-based-virtual-hosting> Name-based virtual hosts support routing HTTP traffic to multiple host names at the same IP address.

NEW QUESTION 156

- (Exam Topic 2)

You are developing a corporate tool on Compute Engine for the finance department, which needs to authenticate users and verify that they are in the finance department. All company employees use G Suite. What should you do?

- A. Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department.
- B. Verify the provided JSON Web Token within the application.
- C. Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department.
- D. Issue client-side certificates to everybody in the finance team and verify the certificates in the application.
- E. Configure Cloud Armor Security Policies to restrict access to only corporate IP address range.
- F. Verify the provided JSON Web Token within the application.
- G. Configure Cloud Armor Security Policies to restrict access to only corporate IP address range.
- H. Issue client side certificates to everybody in the finance team and verify the certificates in the application.

Answer: A

Explanation:

https://cloud.google.com/iap/docs/signed-headers-howto#securing_iap_headers (<https://cloud.google.com/endpoints/docs/openapi/authenticating-users-google-id>).
<https://cloud.google.com/armor/docs/security-policy-overview#:~:text=Google%20Cloud%20Armor%20security%20policies%20protect%20your%20application%20by%20providing%20Layer%207%20filtering%20and%20by%20scrubbing%20incoming%20requests%20for%20common%20web%20attacks%20or%20other%20Layer%207%20attributes%20to%20potentially%20block%20traffic%20before%20it%20reaches%20your%20load%20balanced%20backend%20services%20or%20backend%20buckets>"

NEW QUESTION 158

- (Exam Topic 2)

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

NEW QUESTION 161

- (Exam Topic 2)

You are using Cloud 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.

Answer: D

NEW QUESTION 165

- (Exam Topic 2)

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 rule.
- B. Allow your company's proxy as the only source IP address.
- C. 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.
- D. Set up a VPN tunnel between your company network and your instance's VPC location on Google Cloud.
- E. Configure the required firewall rules and routing information to both the on-premises and Google Cloud networks.
- F. Add a public IP address to your instance, and allow traffic from the internet.
- G. Generate a random hash, and create a subdomain that includes this hash and points to your instance.
- H. Distribute this DNS address to your company's employees.

Answer: B

Explanation:

<https://cloud.google.com/blog/topics/developers-practitioners/control-access-your-web-sites-identity-aware-proxy>

NEW QUESTION 166

- (Exam Topic 2)

Your company uses Cloud Logging to manage large volumes of log data. You need to build a real-time log analysis architecture that pushes logs to a third-party application for processing. What should you do?

- 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: B

NEW QUESTION 171

- (Exam Topic 2)

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
- B. Create a view with an appropriate WHERE clause to select records from a particular dataset for the specific department
- C. Authorize this view to access records from your Master dataset
- D. Give employees the permission to this department-specific dataset.
- E. Create a separate dataset for each department
- F. Create a data pipeline for each department to copy appropriate information from the Master dataset to the specific dataset for the department
- G. Give employees the permission to this department-specific dataset.
- H. Create a dataset named Master dataset
- I. Create a separate view for each department in the Master dataset
- J. Give employees access to the specific view for their department.
- K. Create a dataset named Master dataset
- L. Create a separate table for each department in the Master dataset
- M. Give employees access to the specific table for their department.

Answer: B

NEW QUESTION 176

- (Exam Topic 2)

You recently migrated an on-premises monolithic application to a microservices application on Google Kubernetes Engine (GKE). The application has dependencies on backend services on-premises, including a CRM system and a MySQL database that contains personally identifiable information (PII). The backend services must remain on-premises to meet regulatory requirements.

You established a Cloud VPN connection between your on-premises data center and Google Cloud. You notice that some requests from your microservices application on GKE to the backend services are failing due to latency issues caused by fluctuating bandwidth, which is causing the application to crash. How should you address the latency issues?

- A. Use Memorystore to cache frequently accessed PII data from the on-premises MySQL database
- B. Use Istio to create a service mesh that includes the microservices on GKE and the on-premises services
- C. Increase the number of Cloud VPN tunnels for the connection between Google Cloud and the on-premises services
- D. Decrease the network layer packet size by decreasing the Maximum Transmission Unit (MTU) value from its default value on Cloud VPN

Answer: C

Explanation:

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/choosing-networks-routing#route-alignment>

NEW QUESTION 181

- (Exam Topic 2)

You have an application written in Python running in production on Cloud Run. Your application needs to read/write data stored in a Cloud Storage bucket in the same project. You want to grant access to your application following the principle of least privilege. What should you do?

- A. Create a user-managed service account with a custom Identity and Access Management (IAM) role.
- B. Create a user-managed service account with the Storage Admin Identity and Access Management (IAM) role.
- C. Create a user-managed service account with the Project Editor Identity and Access Management (IAM) role.
- D. Use the default service account linked to the Cloud Run revision in production.

Answer: A

Explanation:

<https://cloud.google.com/iam/docs/understanding-roles#storage.admin>

NEW QUESTION 183

- (Exam Topic 2)

You have an application deployed in production. When a new version is deployed, you want to ensure that all production traffic is routed to the new version of your application. You also want to keep the previous version deployed so that you can revert to it if there is an issue with the new version.

Which deployment strategy should you use?

- A. Blue/green deployment
- B. Canary deployment

- C. Rolling deployment
- D. Recreate deployment

Answer: A

NEW QUESTION 185

- (Exam Topic 2)

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?

- 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

NEW QUESTION 187

- (Exam Topic 2)

Your App Engine standard configuration is as follows: service: production
instance_class: B1

You want to limit the application to 5 instances. Which code snippet should you include in your configuration?

- A. manual_scaling:instances: 5min_pending_latency: 30ms
- B. manual_scaling:max_instances: 5idle_timeout: 10m
- C. basic_scaling:instances: 5min_pending_latency: 30ms
- D. basic_scaling:max_instances: 5idle_timeout: 10m

Answer: C

NEW QUESTION 192

- (Exam Topic 2)

You are developing an application that consists of several microservices running in a Google Kubernetes Engine cluster. One microservice needs to connect to a third-party database running on-premises. You need to store credentials to the database and ensure that these credentials can be rotated while following security best practices. What should you do?

- A. Store the credentials in a sidecar container proxy, and use it to connect to the third-party database.
- B. Configure a service mesh to allow or restrict traffic from the Pods in your microservice to the database.
- C. Store the credentials in an encrypted volume mount, and associate a Persistent Volume Claim with the client Pod.
- D. Store the credentials as a Kubernetes Secret, and use the Cloud Key Management Service plugin to handle encryption and decryption.

Answer: D

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/how-to/encrypting-secrets>

By default, Google Kubernetes Engine (GKE) encrypts customer content stored at rest, including Secrets. GKE handles and manages this default encryption for you without any additional action on your part.

Application-layer secrets encryption provides an additional layer of security for sensitive data, such as Secrets, stored in etcd. Using this functionality, you can use a key managed with Cloud KMS to encrypt data at the application layer. This encryption protects against attackers who gain access to an offline copy of etcd.

NEW QUESTION 193

- (Exam Topic 2)

You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. Your application can scale horizontally, and each instance of your application needs to have a stable network identity and its own persistent disk.

Which GKE object should you use?

- A. Deployment
- B. StatefulSet
- C. ReplicaSet
- D. ReplicaController

Answer: B

Explanation:

Reference: <https://livebook.manning.com/book/kubernetes-in-action/chapter-10/46>

NEW QUESTION 195

- (Exam Topic 2)

You are developing a new application that has the following design requirements: Creation and changes to the application infrastructure are versioned and auditable.

The application and deployment infrastructure uses Google-managed services as much as possible. The application runs on a serverless compute platform. How should you design the application's architecture?

- A. * 1. Store the application and infrastructure source code in a Git repository.* 2. Use Cloud Build to deploy the application infrastructure with Terraform.* 3. Deploy the application to a Cloud Function as a pipeline step.
- B. * 1. Deploy Jenkins from the Google Cloud Marketplace, and define a continuous integration pipeline in Jenkins.* 2. Configure a pipeline step to pull the application source code from a Git repository.* 3. Deploy the application source code to App Engine as a pipeline step.
- C. * 1. Create a continuous integration pipeline on Cloud Build, and configure the pipeline to deploy the application infrastructure using Deployment Manager templates.* 2. Configure a pipeline step to create a container with the latest application source code.* 3. Deploy the container to a Compute Engine instance as a pipeline step.
- D. * 1. Deploy the application infrastructure using gcloud commands.* 2. Use Cloud Build to define a continuous integration pipeline for changes to the application source code.* 3. Configure a pipeline step to pull the application source code from a Git repository, and create a containerized application.* 4. Deploy the new container on Cloud Run as a pipeline step.

Answer: D

Explanation:

Reference: <https://cloud.google.com/docs/ci-cd>

NEW QUESTION 197

- (Exam Topic 2)

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?

- 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

NEW QUESTION 200

- (Exam Topic 2)

Your company has created an application that uploads a report to a Cloud Storage bucket. When the report is uploaded to the bucket, you want to publish a message to a Cloud Pub/Sub topic. You want to implement a solution that will take a small amount of effort to implement. What should you do?

- A. Configure the Cloud Storage bucket to trigger Cloud Pub/Sub notifications when objects are modified.
- B. Create an App Engine application to receive the file; when it is received, publish a message to the Cloud Pub/Sub topic.
- C. Create a Cloud Function that is triggered by the Cloud Storage bucket.
- D. In the Cloud Function, publish a message to the Cloud Pub/Sub topic.
- E. Create an application deployed in a Google Kubernetes Engine cluster to receive the file; when it is received, publish a message to the Cloud Pub/Sub topic.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/pubsub-notifications>

NEW QUESTION 203

- (Exam Topic 2)

You have deployed an HTTP(s) Load Balancer with the gcloud commands shown below.

```

export NAME=load-balancer

# create network
gcloud compute networks create ${NAME}

# add instance
gcloud compute instances create ${NAME}-backend-instance-1 --subnet ${NAME} --no address

# create the instance group
gcloud compute instance-groups unmanaged create ${NAME}-i
gcloud compute instance-groups unmanaged set-named-ports ${NAME}-i --named-ports http:80
gcloud compute instance-groups unmanaged add-instances ${NAME}-i --instances ${NAME}-instance-1

# configure health checks
gcloud compute health-checks create http ${NAME}-http-hc --port 80

# create backend service
gcloud compute backend-services create ${NAME}-http-bes --health-checks ${NAME}-http-hc --protocol HTTP --port-name http
--global
gcloud compute backend-services add-backend ${NAME}-http-bes --instance-group ${NAME}-i --balancing-mode RATE --max-rate
100000 --capacity-scaler 1.0 --global --instance-group-zone us-east1-d

# create url maps and forwarding rule
gcloud compute url-maps create ${NAME}-http-urlmap --default-service ${NAME}-http-bes
gcloud compute target-http-proxies create ${NAME}-http-proxy --url-map ${NAME}-http-urlmap
gcloud compute forwarding-rules create ${NAME}-http-fw --global --ip-protocol ICP --target-http-proxy ${NAME}-http-proxy
--ports 80

```

Health checks to port 80 on the Compute Engine virtual machine instance are failing and no traffic is sent to your instances. You want to resolve the problem. Which commands should you run?

- A. gcloud compute instances add-access-config \${NAME}-backend-instance-1
- B. gcloud compute instances add-tags \${NAME}-backend-instance-1 --tags http-server
- C. gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --source-ranges 130.211.0.0/22,35.191.0.0/16 --direction INGRESS
- D. gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --destination-ranges 130.211.0.0/22,35.191.0.0/16 --direction EGRESS

Answer: C

Explanation:

Reference: <https://cloud.google.com/vpc/docs/special-configurations>

NEW QUESTION 207

- (Exam Topic 2)

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.

Answer: D

Explanation:

<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.

NEW QUESTION 210

- (Exam Topic 2)

You are deploying your applications on Compute Engine. One of your Compute Engine instances failed to launch. What should you do? (Choose two.)

- A. Determine whether your file system is corrupted.
- B. Access Compute Engine as a different SSH user.
- C. Troubleshoot firewall rules or routes on an instance.
- D. Check whether your instance boot disk is completely full.
- E. Check whether network traffic to or from your instance is being dropped.

Answer: AD

Explanation:

<https://cloud.google.com/compute/docs/troubleshooting/vm-startup>

NEW QUESTION 211

- (Exam Topic 2)

You are deploying your application to a Compute Engine virtual machine instance with the Stackdriver Monitoring Agent installed. Your application is a unix process on the instance. You want to be alerted if the unix process has not run for at least 5 minutes. You are not able to change the application to generate metrics or logs. Which alert condition should you configure?

- A. Uptime check
- B. Process health
- C. Metric absence
- D. Metric threshold

Answer: B

Explanation:

Reference: <https://cloud.google.com/monitoring/alerts/concepts-indepth>

NEW QUESTION 213

- (Exam Topic 2)

Your service adds text to images that it reads from Cloud Storage. During busy times of the year, requests to Cloud Storage fail with an HTTP 429 "Too Many Requests" status code. How should you handle this error?

- A. Add a cache-control header to the objects.
- B. Request a quota increase from the GCP Console.
- C. Retry the request with a truncated exponential backoff strategy.
- D. Change the storage class of the Cloud Storage bucket to Multi-regional.

Answer: C

Explanation:

Reference: <https://developers.google.com/gmail/api/v1/reference/quota>

NEW QUESTION 214

- (Exam Topic 2)

Your application is deployed on hundreds of Compute Engine instances in a managed instance group (MIG) in multiple zones. You need to deploy a new instance template to fix a critical vulnerability immediately but must avoid impact to your service. What setting should be made to the MIG after updating the instance template?

- A. Set the Max Surge to 100%.
- B. Set the Update mode to Opportunistic.
- C. Set the Maximum Unavailable to 100%.
- D. Set the Minimum Wait time to 0 seconds.

Answer: B

Explanation:

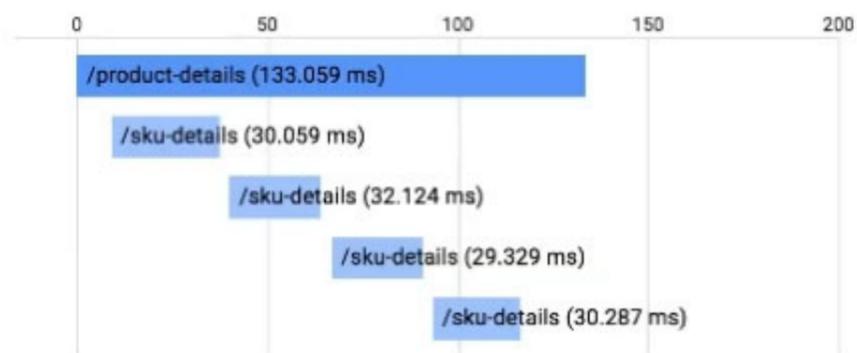
<https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups#type> Alternatively, if an automated update is potentially too disruptive, you can choose to perform an opportunistic update. The MIG applies an opportunistic update only when you manually initiate the update on selected instances or when new instances are created. New instances can be created when you or another service, such as an autoscaler, resizes the MIG. Compute Engine does not actively initiate requests to apply opportunistic updates on existing instances.

NEW QUESTION 217

- (Exam Topic 2)

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?

Timeline



- 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.

Answer: C

NEW QUESTION 219

- (Exam Topic 2)

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

Answer: A

Explanation:

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

NEW QUESTION 220

- (Exam Topic 2)

You want to use the Stackdriver Logging Agent to send an application's log file to Stackdriver from a Compute Engine virtual machine instance. After installing the Stackdriver Logging Agent, what should you do first?

- A. Enable the Error Reporting API on the project.
- B. Grant the instance full access to all Cloud APIs.
- C. Configure the application log file as a custom source.
- D. Create a Stackdriver Logs Export Sink with a filter that matches the application's log entries.

Answer: B

NEW QUESTION 222

- (Exam Topic 2)

Your operations team has asked you to create a script that lists the Cloud Bigtable, Memorystore, and Cloud SQL databases running within a project. The script should allow users to submit a filter expression to limit the results presented. How should you retrieve the data?

- A. Use the HBase API, Redis API, and MySQL connection to retrieve database list
- B. Combine the results, and then apply the filter to display the results
- C. Use the HBase API, Redis API, and MySQL connection to retrieve database list
- D. Filter the results individually, and then combine them to display the results
- E. Run `gcloud bigtable instances list`, `gcloud redis instances list`, and `gcloud sql databases list`
- F. Use a filter within the application, and then display the results
- G. Run `gcloud bigtable instances list`, `gcloud redis instances list`, and `gcloud sql databases list`
- H. Use `--filter` flag with each command, and then display the results

Answer: D

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/topic/filters>

Most `gcloud` commands return a list of resources on success. By default they are pretty-printed on the standard output. The `--format=NAME[ATTRIBUTES](PROJECTION)` and `--filter=EXPRESSION` flags along with projections can be used to format and change the default output to a more meaningful result. Use the `--format` flag to change the default output format of a command. For details run `$ gcloud topic formats`.

NEW QUESTION 226

- (Exam Topic 2)

You are planning to migrate a MySQL database to the managed Cloud SQL database for Google Cloud. You have Compute Engine virtual machine instances that will connect with this Cloud SQL instance. You do not want to whitelist IPs for the Compute Engine instances to be able to access Cloud SQL. What should you do?

- A. Enable private IP for the Cloud SQL instance.
- B. Whitelist a project to access Cloud SQL, and add Compute Engine instances in the whitelisted project.
- C. Create a role in Cloud SQL that allows access to the database from external instances, and assign the Compute Engine instances to that role.
- D. Create a CloudSQL instance on one project
- E. Create Compute engine instances in a different project. Create a VPN between these two projects to allow internal access to CloudSQL.

Answer: C

Explanation:

Reference: <https://cloud.google.com/sql/docs/mysql/connect-external-app>

NEW QUESTION 231

- (Exam Topic 2)

You manage your company's ecommerce platform's payment system, which runs on Google Cloud. Your company must retain user logs for 1 year for internal auditing purposes and for 3 years to meet compliance requirements. You need to store new user logs on Google Cloud to minimize on-premises storage usage and ensure that they are easily searchable. You want to minimize effort while ensuring that the logs are stored correctly. What should you do?

- A. Store the logs in a Cloud Storage bucket with bucket lock turned on.
- B. Store the logs in a Cloud Storage bucket with a 3-year retention period.
- C. Store the logs in Cloud Logging as custom logs with a custom retention period.
- D. Store the logs in a Cloud Storage bucket with a 1-year retention period
- E. After 1 year, move the logs to another bucket with a 2-year retention period.

Answer: C

Explanation:

<https://cloud.google.com/logging/docs/buckets#custom-retention>

NEW QUESTION 235

- (Exam Topic 2)

You are creating and running containers across different projects in Google Cloud. The application you are developing needs to access Google Cloud services from within Google Kubernetes Engine (GKE).

What should you do?

- A. Assign a Google service account to the GKE nodes.
- B. Use a Google service account to run the Pod with Workload Identity.
- C. Store the Google service account credentials as a Kubernetes Secret.
- D. Use a Google service account with GKE role-based access control (RBAC).

Answer: B

Explanation:

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

NEW QUESTION 237

- (Exam Topic 2)

Your analytics system executes queries against a BigQuery dataset. The SQL query is executed in batch and passes the contents of a SQL file to the BigQuery CLI. Then it redirects the BigQuery CLI output to another process. However, you are getting a permission error from the BigQuery CLI when the queries are executed. You want to resolve the issue. What should you do?

- A. Grant the service account BigQuery Data Viewer and BigQuery Job User roles.
- B. Grant the service account BigQuery Data Editor and BigQuery Data Viewer roles.
- C. Create a view in BigQuery from the SQL query and SELECT* from the view in the CLI.
- D. Create a new dataset in BigQuery, and copy the source table to the new dataset Query the new dataset and table from the CLI.

Answer: B

NEW QUESTION 238

- (Exam Topic 2)

You are developing an application that will allow users to read and post comments on news articles. You want to configure your application to store and display user-submitted comments using Firestore. How should you design the schema to support an unknown number of comments and articles?

- A. Store each comment in a subcollection of the article.
- B. Add each comment to an array property on the article.
- C. Store each comment in a document, and add the comment's key to an array property on the article.
- D. Store each comment in a document, and add the comment's key to an array property on the user profile.

Answer: D

NEW QUESTION 243

- (Exam Topic 2)

You are using Cloud Build to build a Docker image. You need to modify the build to execute unit and run integration tests. When there is a failure, you want the build history to clearly display the stage at which the build failed.

What should you do?

- A. Add RUN commands in the Dockerfile to execute unit and integration tests.
- B. Create a Cloud Build build config file with a single build step to compile unit and integration tests.
- C. Create a Cloud Build build config file that will spawn a separate cloud build pipeline for unit and integration tests.
- D. Create a Cloud Build build config file with separate cloud builder steps to compile and execute unit and integration tests.

Answer: D

NEW QUESTION 247

- (Exam Topic 2)

Your development team has been asked to refactor an existing monolithic application into a set of composable microservices. Which design aspects should you implement for the new application? (Choose two.)

- A. Develop the microservice code in the same programming language used by the microservice caller.
- B. Create an API contract agreement between the microservice implementation and microservice caller.
- C. Require asynchronous communications between all microservice implementations and microservice callers.
- D. Ensure that sufficient instances of the microservice are running to accommodate the performance requirements.
- E. Implement a versioning scheme to permit future changes that could be incompatible with the current interface.

Answer: BE

NEW QUESTION 252

- (Exam Topic 2)

Your application performs well when tested locally, but it runs significantly slower when you deploy it to App Engine standard environment. You want to diagnose the problem. What should you do?

- A. File a ticket with Cloud Support indicating that the application performs faster locally.
- B. Use Stackdriver Debugger Snapshots to look at a point-in-time execution of the application.
- C. Use Stackdriver Trace to determine which functions within the application have higher latency.
- D. Add logging commands to the application and use Stackdriver Logging to check where the latency problem occurs.

Answer: D

NEW QUESTION 257

- (Exam Topic 2)

You are trying to connect to your Google Kubernetes Engine (GKE) cluster using kubectl from Cloud Shell. You have deployed your GKE cluster with a public endpoint. From Cloud Shell, you run the following command:

```
gcloud container clusters get-credentials <cluster-name> \
  --zone <zone> --project <project-name> \
```

You notice that the kubectl commands time out without returning an error message. What is the most likely cause of this issue?

- A. Your user account does not have privileges to interact with the cluster using kubectl.
- B. Your Cloud Shell external IP address is not part of the authorized networks of the cluster.
- C. The Cloud Shell is not part of the same VPC as the GKE cluster.
- D. A VPC firewall is blocking access to the cluster's endpoint.

Answer: B

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/private-clusters#cloud_shell

If you want to use Cloud Shell to access the cluster, you must add the public IP address of your Cloud Shell to the cluster's list of authorized networks.

NEW QUESTION 258

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Professional-Cloud-Developer Practice Exam Features:

- * Professional-Cloud-Developer Questions and Answers Updated Frequently
- * Professional-Cloud-Developer Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-Developer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Professional-Cloud-Developer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-Developer Practice Test Here](#)