

Google

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer



NEW QUESTION 1

- (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 2

- (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 3

- (Exam Topic 1)

In order for HipLocal to store application state and meet their stated business requirements, which database service should they migrate to?

- A. Cloud Spanner
- B. Cloud Datastore
- C. Cloud Memorystore as a cache
- D. Separate Cloud SQL clusters for each region

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

HipLocal has connected their Hadoop infrastructure to GCP using Cloud Interconnect in order to query data stored on persistent disks. Which IP strategy should they use?

- A. Create manual subnets.
- B. Create an auto mode subnet.
- C. Create multiple peered VPCs.
- D. Provision a single instance for NAT.

Answer: A

NEW QUESTION 5

- (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 6

- (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 7

- (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 8

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 9

- (Exam Topic 2)

You work for an organization that manages an ecommerce site. Your application is deployed behind a global HTTP(S) load balancer. You need to test a new product recommendation algorithm. You plan to use A/B testing to determine the new algorithm's effect on sales in a randomized way. How should you test this feature?

- A. Split traffic between versions using weights.
- B. Enable the new recommendation feature flag on a single instance.
- C. Mirror traffic to the new version of your application.
- D. Use HTTP header-based routing.

Answer: A

Explanation:

https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic_actions_weight-based_tra Deploying a new version of an existing production service generally incurs some risk. Even if your tests pass in staging, you probably don't want to subject 100% of your users to the new version immediately. With traffic management, you can define percentage-based traffic splits across multiple backend services.

For example, you can send 95% of the traffic to the previous version of your service and 5% to the new version of your service. After you've validated that the new production version works as expected, you can gradually shift the percentages until 100% of the traffic reaches the new version of your service. Traffic splitting is typically used for deploying new versions, A/B testing, service migration, and similar processes.

https://cloud.google.com/traffic-director/docs/advanced-traffic-management#weight-based_traffic_splitting_for_ https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#split_the_traffic_ https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic_actions_weight-based_tra

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:

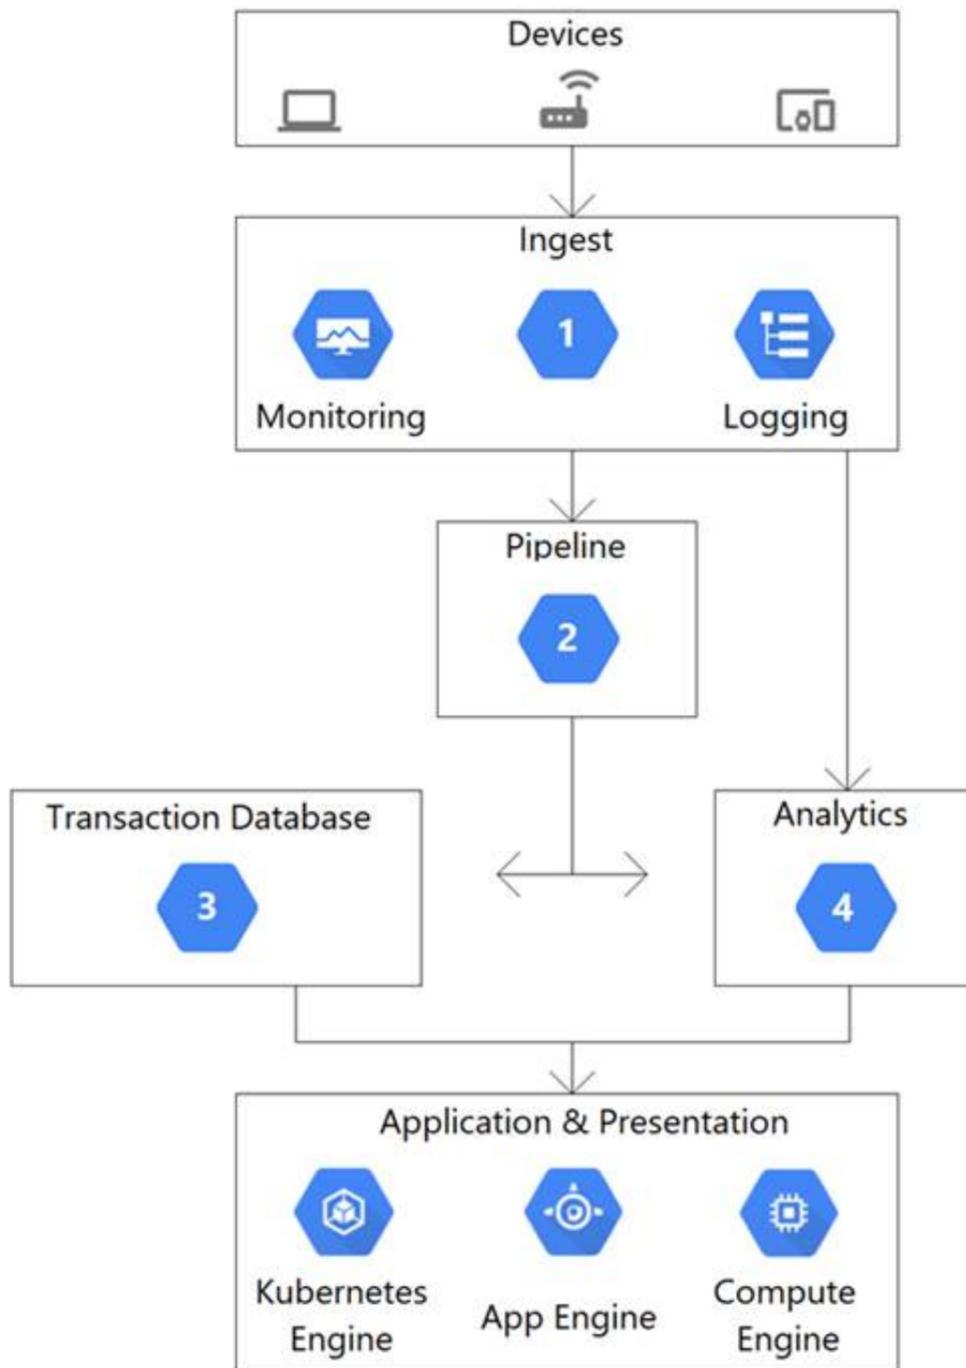
Reference:

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

NEW QUESTION 10

- (Exam Topic 2)

This architectural diagram depicts a system that streams data from thousands of devices. You want to ingest data into a pipeline, store the data, and analyze the data using SQL statements. Which Google Cloud services should you use for steps 1, 2, 3, and 4?



- A. 1) App Engine 2) Pub/Sub 3) BigQuery 4) Firestore
- B. 1) Dataflow 2) Pub/Sub 3) Firestore 4) BigQuery
- C. 1) Pub/Sub 2) Dataflow 3) BigQuery 4) Firestore
- D. 1) Pub/Sub 2) Dataflow 3) Firestore 4) BigQuery

Answer: D

NEW QUESTION 12

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 13

- (Exam Topic 2)

Your existing application keeps user state information in a single MySQL database. This state information is very user-specific and depends heavily on how long a user has been using an application. The MySQL database is causing challenges to maintain and enhance the schema for various users.

Which storage option should you choose?

- A. Cloud SQL
- B. Cloud Storage
- C. Cloud Spanner
- D. Cloud Datastore/Firestore

Answer: A

Explanation:

Reference: <https://cloud.google.com/solutions/migrating-mysql-to-cloudsql-concept>

NEW QUESTION 14

- (Exam Topic 2)

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

- A. Deploy the website on App Engine and use traffic splitting.
- B. Deploy the website on App Engine as three separate services.
- C. Deploy the website on Cloud Functions and use traffic splitting.
- D. Deploy the website on Cloud Functions as three separate functions.

Answer: A

Explanation:

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

NEW QUESTION 17

- (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 21

- (Exam Topic 2)

You are writing a single-page web application with a user-interface that communicates with a third-party API for content using XMLHttpRequest. The data displayed on the UI by the API results is less critical than other data displayed on the same web page, so it is acceptable for some requests to not have the API data displayed in the UI. However, calls made to the API should not delay rendering of other parts of the user interface. You want your application to perform well when the API response is an error or a timeout. What should you do?

- A. Set the asynchronous option for your requests to the API to false and omit the widget displaying the API results when a timeout or error is encountered.
- B. Set the asynchronous option for your request to the API to true and omit the widget displaying the API results when a timeout or error is encountered.
- C. Catch timeout or error exceptions from the API call and keep trying with exponential backoff until the API response is successful.
- D. Catch timeout or error exceptions from the API call and display the error response in the UI widget.

Answer: A

NEW QUESTION 22

- (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 file.
- 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 file.
- 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 file.
- 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 file.
- 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 24

- (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 27

- (Exam Topic 2)

Your data is stored in Cloud Storage buckets. Fellow developers have reported that data downloaded from Cloud Storage is resulting in slow API performance. You want to research the issue to provide details to the GCP support team. Which command should you run?

- A. `gsutil test -o output.json gs://my-bucket`
- B. `gsutil perfdiag -o output.json gs://my-bucket`
- C. `gcloud compute scp example-instance:~/test-data -o output.json gs://my-bucket`
- D. `gcloud services test -o output.json gs://my-bucket`

Answer: B

Explanation:

Reference: <https://groups.google.com/forum/#!topic/gce-discussion/xBI9Jq5HDsY>

NEW QUESTION 28

- (Exam Topic 2)

You developed a JavaScript web application that needs to access Google Drive's API and obtain permission from users to store files in their Google Drives. You need to select an authorization approach for your application. What should you do?

- A. Create an API key.
- B. Create a SAML token.
- C. Create a service account.
- D. Create an OAuth Client ID.

Answer: D

Explanation:

Reference: <https://developers.google.com/drive/api/v3/about-auth>

NEW QUESTION 33

- (Exam Topic 2)

You are designing a resource-sharing policy for applications used by different teams in a Google Kubernetes Engine cluster. You need to ensure that all applications can access the resources needed to run. What should you do? (Choose two.)

- A. Specify the resource limits and requests in the object specifications.
- B. Create a namespace for each team, and attach resource quotas to each namespace.
- C. Create a LimitRange to specify the default compute resource requirements for each namespace.
- D. Create a Kubernetes service account (KSA) for each application, and assign each KSA to the namespace.
- E. Use the Anthos Policy Controller to enforce label annotations on all namespace
- F. Use taints and tolerations to allow resource sharing for namespaces.

Answer: BC

Explanation:

<https://kubernetes.io/docs/concepts/policy/resource-quotas/> <https://kubernetes.io/docs/concepts/policy/limit-range/>
<https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-resource-requests-and-l>

NEW QUESTION 35

- (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 38

- (Exam Topic 2)

You are parsing a log file that contains three columns: a timestamp, an account number (a string), and a transaction amount (a number). You want to calculate the sum of all transaction amounts for each unique account number efficiently.

Which data structure should you use?

- A. A linked list
- B. A hash table
- C. A two-dimensional array
- D. A comma-delimited string

Answer: B

NEW QUESTION 43

- (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 44

- (Exam Topic 2)

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

- A. Write an exponential backoff process around the client library call.
- B. Write a one-second wait time backoff process around the client library call.
- C. Design a retry button in the application and ask users to click if the error occurs.
- D. Create a queue for the object and inform the users that the application will try again in 10 minutes.

Answer: A

NEW QUESTION 45

- (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 50

- (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 55

- (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 58

- (Exam Topic 2)

Your application is composed of a set of loosely coupled services orchestrated by code executed on Compute Engine. You want your application to easily bring up new Compute Engine instances that find and use a specific version of a service. How should this be configured?

- A. Define your service endpoint information as metadata that is retrieved at runtime and used to connect to the desired service.
- B. Define your service endpoint information as label data that is retrieved at runtime and used to connect to the desired service.
- C. Define your service endpoint information to be retrieved from an environment variable at runtime and used to connect to the desired service.
- D. Define your service to use a fixed hostname and port to connect to the desired servic
- E. Replace the service at the endpoint with your new version.

Answer: A

Explanation:

<https://cloud.google.com/service-infrastructure/docs/service-metadata/reference/rest#service-endpoint>

NEW QUESTION 63

- (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 66

- (Exam Topic 2)

You want to notify on-call engineers about a service degradation in production while minimizing development time. What should you do?

- A. Use Cloud Function to monitor resources and raise alerts.
- B. Use Cloud Pub/Sub to monitor resources and raise alerts.
- C. Use Stackdriver Error Reporting to capture errors and raise alerts.
- D. Use Stackdriver Monitoring to monitor resources and raise alerts.

Answer: A

NEW QUESTION 71

- (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 72

- (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 75

- (Exam Topic 2)

You recently developed a new service on Cloud Run. The new service authenticates using a custom service and then writes transactional information to a Cloud Spanner database. You need to verify that your application can support up to 5,000 read and 1,000 write transactions per second while identifying any bottlenecks that occur. Your test infrastructure must be able to autoscale. What should you do?

- A. Build a test harness to generate requests and deploy it to Cloud Ru
- B. Analyze the VPC Flow Logs using Cloud Logging.
- C. Create a Google Kubernetes Engine cluster running the Locust or JMeter images to dynamically generate load test

- D. Analyze the results using Cloud Trace.
- E. Create a Cloud Task to generate a test load.
- F. Use Cloud Scheduler to run 60,000 Cloud Task transactions per minute for 10 minutes.
- G. Analyze the results using Cloud Monitoring.
- H. Create a Compute Engine instance that uses a LAMP stack image from the Marketplace, and use Apache Bench to generate load tests against the service.
- I. Analyze the results using Cloud Trace.

Answer: B

Explanation:

<https://cloud.google.com/architecture/distributed-load-testing-using-gke>

NEW QUESTION 80

- (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 85

- (Exam Topic 2)

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

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

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

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

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

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

Answer: AD

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/builds/submit> <https://cloud.google.com/artifact-registry/docs/configure-cloud-build>

NEW QUESTION 87

- (Exam Topic 2)

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: B

Explanation:

<https://kubernetes.io/docs/tasks/configure-pod-container/configure-liveness-readiness-startup-probes/#configure>

NEW QUESTION 91

- (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 Registry
- D. Review the vulnerability results before each deployment.
- E. Enable Container Analysis, and upload new container images to Artifact Registry
- 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 93

- (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 94

- (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 96

- (Exam Topic 2)

You have recently instrumented a new application with OpenTelemetry, and you want to check the latency of your application requests in Trace. You want to ensure that a specific request is always traced. What should you do?

- A. Wait 10 minutes, then verify that Trace captures those types of requests automatically.
- B. Write a custom script that sends this type of request repeatedly from your dev project.
- C. Use the Trace API to apply custom attributes to the trace.
- D. Add the X-Cloud-Trace-Context header to the request with the appropriate parameters.

Answer: D

Explanation:

<https://cloud.google.com/trace/docs/setup#force-trace>

Cloud Trace doesn't sample every request. To force a specific request to be traced, add an X-Cloud-Trace-Context header to the request.

NEW QUESTION 101

- (Exam Topic 2)

You are using Cloud Build for your CI/CD pipeline to complete several tasks, including copying certain files to Compute Engine virtual machines. Your pipeline requires a flat file that is generated in one builder in the pipeline to be accessible by subsequent builders in the same pipeline. How should you store the file so that all the builders in the pipeline can access it?

- A. Store and retrieve the file contents using Compute Engine instance metadata.
- B. Output the file contents to a file in /workspace
- C. Read from the same /workspace file in the subsequent build step.
- D. Use gsutil to output the file contents to a Cloud Storage object
- E. Read from the same object in the subsequent build step.
- F. Add a build argument that runs an HTTP POST via curl to a separate web server to persist the value in one build
- G. Use an HTTP GET via curl from the subsequent build step to read the value.

Answer: B

Explanation:

<https://cloud.google.com/build/docs/build-config-file-schema>

NEW QUESTION 102

- (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 105

- (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 106

- (Exam Topic 2)

You are writing a Compute Engine hosted application in project A that needs to securely authenticate to a Cloud Pub/Sub topic in project B. What should you do?

- A. Configure the instances with a service account owned by project
- B. Add the service account as a Cloud Pub/Sub publisher to project A.
- C. Configure the instances with a service account owned by project
- D. Add the service account as a publisher on the topic.
- E. Configure Application Default Credentials to use the private key of a service account owned by project
- F. Add the service account as a Cloud Pub/Sub publisher to project A.
- G. Configure Application Default Credentials to use the private key of a service account owned by project
- H. Add the service account as a publisher on the topic

Answer: B

Explanation:

<https://cloud.google.com/pubsub/docs/access-control>

"For example, suppose a service account in Cloud Project A wants to publish messages to a topic in Cloud Project B. You could accomplish this by granting the service account Edit permission in Cloud Project B"

NEW QUESTION 108

- (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 progress
- F. After the application reads the message, it disconnects from the database.

Answer: D

NEW QUESTION 109

- (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 110

- (Exam Topic 2)

Your application is deployed in a Google Kubernetes Engine (GKE) cluster. You want to expose this application publicly behind a Cloud Load Balancing HTTP(S) load balancer. What should you do?

- A. Configure a GKE Ingress resource.
- B. Configure a GKE Service resource.
- C. Configure a GKE Ingress resource with type: LoadBalancer.
- D. Configure a GKE Service resource with type: LoadBalancer.

Answer: A

Explanation:

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

NEW QUESTION 113

- (Exam Topic 2)

Your teammate has asked you to review the code below. Its purpose is to efficiently add a large number of small rows to a BigQuery table.

```
BigQuery service = BigQueryOptions.newBuilder().build().getService();

public void writeToBigQuery(Collection<Map<String, String>> rows){
    for(Map<String, String> row : rows) {
        InsertAllRequest insertRequest = InsertAllRequest.newBuilder(
            "datasetId", "tableId",
            InsertAllRequest.RowToInsert.of(row)).build();
        service.insertAll(insertRequest);
    }
}
```

Which improvement should you suggest your teammate make?

- A. Include multiple rows with each request.
- B. Perform the inserts in parallel by creating multiple threads.
- C. Write each row to a Cloud Storage object, then load into BigQuery.
- D. Write each row to a Cloud Storage object in parallel, then load into BigQuery.

Answer: B

NEW QUESTION 116

- (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 121

- (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 123

- (Exam Topic 2)

Your application is logging to Stackdriver. You want to get the count of all requests on all `/api/alpha/*` endpoints. What should you do?

- A. Add a Stackdriver counter metric for path:`/api/alpha/`.
- B. Add a Stackdriver counter metric for endpoint:`/api/alpha/*`.
- C. Export the logs to Cloud Storage and count lines matching `/api/alpha`.
- D. Export the logs to Cloud Pub/Sub and count lines matching `/api/alpha`.

Answer: C

NEW QUESTION 125

- (Exam Topic 2)

You manage an application that runs in a Compute Engine instance. You also have multiple backend services executing in stand-alone Docker containers running in Compute Engine instances. The Compute Engine instances supporting the backend services are scaled by managed instance groups in multiple regions. You want your calling application to be loosely coupled. You need to be able to invoke distinct service implementations that are chosen based on the value of an HTTP header found in the request. Which Google Cloud feature should you use to invoke the backend services?

- A. Traffic Director
- B. Service Directory
- C. Anthos Service Mesh
- D. Internal HTTP(S) Load Balancing

Answer: D

NEW QUESTION 127

- (Exam Topic 2)

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

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

Answer: A

Explanation:

Reference: <https://cloud.google.com/architecture/application-deployment-and-testing-strategies>

NEW QUESTION 130

- (Exam Topic 2)

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

- A. Assign the Google Cloud service account to your GKE Pod using Workload Identity.
- B. Export the Google Cloud service account, and share it with the Pod as a Kubernetes Secret.
- C. Export the Google Cloud service account, and embed it in the source code of the application.
- D. Export the Google Cloud service account, and upload it to HashiCorp Vault to generate a dynamic service account for your application.

Answer: A

Explanation:

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

Applications running on GKE might need access to Google Cloud APIs such as Compute Engine API, BigQuery Storage API, or Machine Learning APIs. Workload Identity allows a Kubernetes service account in your GKE cluster to act as an IAM service account. Pods that use the configured Kubernetes service account automatically authenticate as the IAM service account when accessing Google Cloud APIs. Using Workload Identity allows you to assign distinct, fine-grained identities and authorization for each application in your cluster.

NEW QUESTION 131

- (Exam Topic 2)

You want to view the memory usage of your application deployed on Compute Engine. What should you do?

- A. Install the Stackdriver Client Library.
- B. Install the Stackdriver Monitoring Agent.
- C. Use the Stackdriver Metrics Explorer.
- D. Use the Google Cloud Platform Console.

Answer: C

Explanation:

Reference:

<https://stackoverflow.com/questions/43991246/google-cloud-platform-how-to-monitor-memory-usage-of-vm-in>

NEW QUESTION 136

- (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 140

- (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 142

- (Exam Topic 2)

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

What should you do?

- A. Ask the user to run the jobs as batch jobs.
- B. Create a separate project for the user to run jobs.
- C. Add the user as a job.user role in the existing project.
- D. Allow the user to run jobs when important workloads are not running.

Answer: B

NEW QUESTION 144

- (Exam Topic 2)

You are configuring a continuous integration pipeline using Cloud Build to automate the deployment of new container images to Google Kubernetes Engine (GKE). The pipeline builds the application from its source code, runs unit and integration tests in separate steps, and pushes the container to Container Registry. The application runs on a Python web server.

The Dockerfile is as follows: FROM python:3.7-alpine - COPY ./app WORKDIR /app

RUN pip install -r requirements.txt CMD ["gunicorn", "-w 4", "main:app"]

You notice that Cloud Build runs are taking longer than expected to complete. You want to decrease the build time. What should you do? (Choose two.)

- A. Select a virtual machine (VM) size with higher CPU for Cloud Build runs.
- B. Deploy a Container Registry on a Compute Engine VM in a VPC, and use it to store the final images.
- C. Cache the Docker image for subsequent builds using the --cache-from argument in your build config file.
- D. Change the base image in the Dockerfile to ubuntu:latest, and install Python 3.7 using a package manager utility.
- E. Store application source code on Cloud Storage, and configure the pipeline to use gsutil to download the source code.

Answer: AC

Explanation:

<https://cloud.google.com/build/docs/optimize-builds/increase-vcpu-for-builds>

By default, Cloud Build runs your builds on a standard virtual machine (VM). In addition to the standard VM, Cloud Build provides several high-CPU VM types to run builds. To increase the speed of your build, select a machine with a higher vCPU to run builds. Keep in mind that although selecting a high vCPU machine increases your build speed, it may also increase the startup time of your build as Cloud Build only starts non-standard machines on demand.

https://cloud.google.com/build/docs/optimize-builds/speeding-up-builds#using_a_cached_docker_image

The easiest way to increase the speed of your Docker image build is by specifying a cached image that can be used for subsequent builds. You can specify the cached image by adding the --cache-from argument in your build config file, which will instruct Docker to build using that image as a cache source.

NEW QUESTION 148

- (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 150

- (Exam Topic 2)

You are developing an HTTP API hosted on a Compute Engine virtual machine instance that needs to be invoked by multiple clients within the same Virtual Private Cloud (VPC). You want clients to be able to get the IP address of the service. What should you do?

- A. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Clients should use this IP address to connect to the service.
- B. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Then, define an A record in Cloud DNS.
- C. Clients should use the name of the A record to connect to the service.
- D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[INSTANCE_NAME\].\[ZONE\].c.\[PROJECT_ID\].internal/](https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/).
- E. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[API_NAME\]/\[API_VERSION\]/](https://[API_NAME]/[API_VERSION]/).

Answer: D

NEW QUESTION 155

- (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 158

- (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 159

- (Exam Topic 2)

You plan to deploy a new application revision with a Deployment resource to Google Kubernetes Engine (GKE) in production. The container might not work correctly. You want to minimize risk in case there are issues after deploying the revision. You want to follow Google-recommended best practices. What should

you do?

- A. Perform a rolling update with a PodDisruptionBudget of 80%.
- B. Perform a rolling update with a HorizontalPodAutoscaler scale-down policy value of 0.
- C. Convert the Deployment to a StatefulSet, and perform a rolling update with a PodDisruptionBudget of 80%.
- D. Convert the Deployment to a StatefulSet, and perform a rolling update with a HorizontalPodAutoscaler scale-down policy value of 0.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/containers-kubernetes/ensuring-reliability-and-uptime-for-your-gke-clus> Setting PodDisruptionBudget ensures that your workloads have a sufficient number of replicas, even during maintenance. Using the PDB, you can define a number (or percentage) of pods that can be terminated, even if terminating them brings the current replica count below the desired value. With PDB configured, Kubernetes will drain a node following the configured disruption schedule. New pods will be deployed on other available nodes. This approach ensures Kubernetes schedules workloads in an optimal way while controlling the disruption based on the PDB configuration.

<https://blog.knoldus.com/how-to-avoid-outages-in-your-kubernetes-cluster-using-pdb/>

NEW QUESTION 163

- (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 166

- (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 169

- (Exam Topic 2)

You recently deployed your application in Google Kubernetes Engine, and now need to release a new version of your application. You need the ability to instantly roll back to the previous version in case there are issues with the new version. Which deployment model should you use?

- A. Perform a rolling deployment, and test your new application after the deployment is complete.
- B. Perform A/B testing, and test your application periodically after the new tests are implemented.
- C. Perform a blue/green deployment, and test your new application after the deployment is complete.
- D. Perform a canary deployment, and test your new application periodically after the new version is deployed.

Answer: C

NEW QUESTION 173

- (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 174

- (Exam Topic 2)

You are designing a schema for a Cloud Spanner customer database. You want to store a phone number array field in a customer table. You also want to allow users to search customers by phone number. How should you design this schema?

- A. Create a table named Customer
- B. Add an Array field in a table that will hold phone numbers for the customer.
- C. Create a table named Customer
- D. Create a table named Phone
- E. Add a CustomerId field in the Phones table to find the CustomerId from a phone number.
- F. Create a table named Customer
- G. Add an Array field in a table that will hold phone numbers for the customer.
- H. Create a secondary index on the Array field.
- I. Create a table named Customers as a parent table
- J. Create a table named Phones, and interleave this table into the Customer table
- K. Create an index on the phone number field in the Phones table.

Answer: C

NEW QUESTION 176

- (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 181

- (Exam Topic 2)

You are developing an ecommerce web application that uses App Engine standard environment and Memorystore for Redis. When a user logs into the app, the application caches the user's information (e.g., session, name, address, preferences), which is stored for quick retrieval during checkout. While testing your application in a browser, you get a 502 Bad Gateway error. You have determined that the application is not connecting to Memorystore. What is the reason for this error?

- A. Your Memorystore for Redis instance was deployed without a public IP address.
- B. You configured your Serverless VPC Access connector in a different region than your App Engine instance.
- C. The firewall rule allowing a connection between App Engine and Memorystore was removed during an infrastructure update by the DevOps team.
- D. You configured your application to use a Serverless VPC Access connector on a different subnet in a different availability zone than your App Engine instance.

Answer: A

Explanation:

Reference: <https://cloud.google.com/endpoints/docs/openapi/troubleshoot-response-errors>

NEW QUESTION 185

- (Exam Topic 2)

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. Your Cloud Build job is reporting deployment failures immediately after successfully validating the code. What should you do?

- 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: D

Explanation:

<https://cloud.google.com/build/docs/securing-builds/configure-access-for-cloud-build-service-account>

NEW QUESTION 186

- (Exam Topic 2)

You have been tasked with planning the migration of your company's application from on-premises to Google Cloud. Your company's monolithic application is an ecommerce website. The application will be migrated to microservices deployed on Google Cloud in stages. The majority of your company's revenue is generated through online sales, so it is important to minimize risk during the migration. You need to prioritize features and select the first functionality to migrate. What should you do?

- A. Migrate the Product catalog, which has integrations to the frontend and product database.
- B. Migrate Payment processing, which has integrations to the frontend, order database, and third-party payment vendor.
- C. Migrate Order fulfillment, which has integrations to the order database, inventory system, and third-party shipping vendor.
- D. Migrate the Shopping cart, which has integrations to the frontend, cart database, inventory system, and payment processing system.

Answer: A

NEW QUESTION 191

- (Exam Topic 2)

You deployed a new application to Google Kubernetes Engine and are experiencing some performance degradation. Your logs are being written to Cloud Logging, and you are using a Prometheus sidecar model for capturing metrics. You need to correlate the metrics and data from the logs to troubleshoot the performance issue and send real-time alerts while minimizing costs. What should you do?

- A. Create custom metrics from the Cloud Logging logs, and use Prometheus to import the results using the Cloud Monitoring REST API.
- B. Export the Cloud Logging logs and the Prometheus metrics to Cloud Bigtable.
- C. Run a query to join the results, and analyze in Google Data Studio.
- D. Export the Cloud Logging logs and stream the Prometheus metrics to BigQuery.
- E. Run a recurring query to join the results, and send notifications using Cloud Tasks.
- F. Export the Prometheus metrics and use Cloud Monitoring to view them as external metric.
- G. Configure Cloud Monitoring to create log-based metrics from the logs, and correlate them with the Prometheus data.

Answer: D

Explanation:

Reference:

<https://cloud.google.com/blog/products/operations/troubleshoot-gke-faster-with-monitoring-data-in-your-logs>

NEW QUESTION 193

- (Exam Topic 2)

You are building a CI/CD pipeline that consists of a version control system, Cloud Build, and Container Registry. Each time a new tag is pushed to the repository, a Cloud Build job is triggered, which runs unit tests on the new code, builds a new Docker container image, and pushes it into Container Registry. The last step of your pipeline should deploy the new container to your production Google Kubernetes Engine (GKE) cluster. You need to select a tool and deployment strategy that meets the following requirements:

- Zero downtime is incurred
 - Testing is fully automated
 - Allows for testing before being rolled out to users
 - Can quickly rollback if needed
- What should you do?

- A. Trigger a Spinnaker pipeline configured as an A/B test of your new code and, if it is successful, deploy the container to production.
- B. Trigger a Spinnaker pipeline configured as a canary test of your new code and, if it is successful, deploy the container to production.
- C. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a canary test.
- D. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a shadow test.

Answer: D

Explanation:

https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#perform_a_shad With a shadow test, you test the new version of your application by mirroring user traffic from the current application version without impacting the user requests.

NEW QUESTION 197

- (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 200

- (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 205

- (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 projec
- 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:

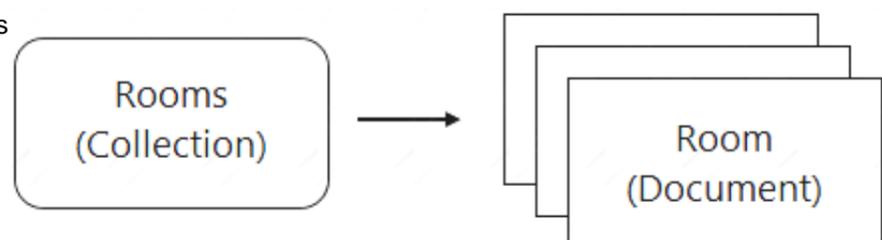
Reference: <https://cloud.google.com/sql/docs/mysql/connect-external-app>

NEW QUESTION 210

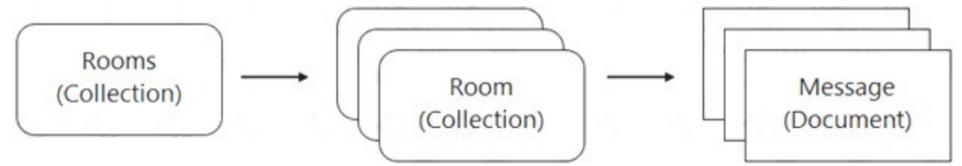
- (Exam Topic 2)

You are designing a chat room application that will host multiple rooms and retain the message history for each room. You have selected Firestore as your database. How should you represent the data in Firestore?

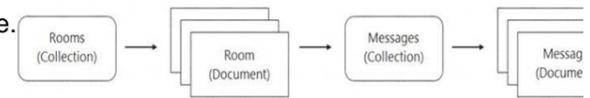
- A. Create a collection for the room
- B. For each room, create a document that lists the contents of the messages



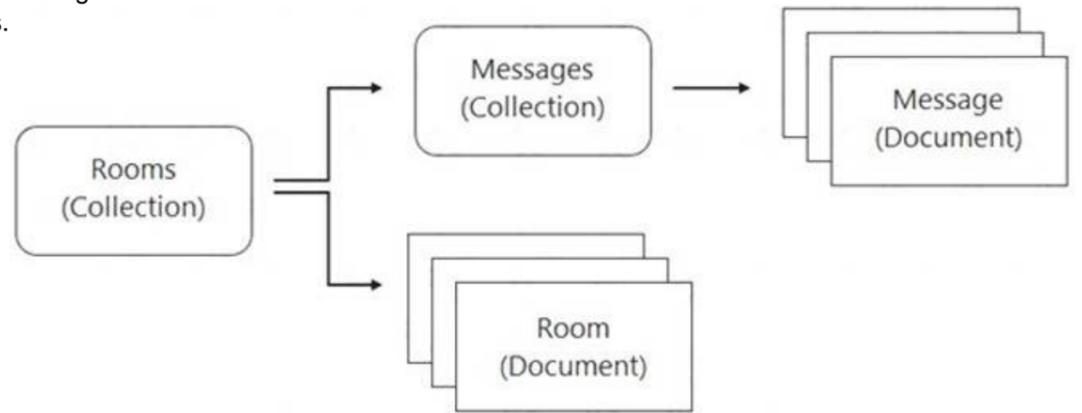
- C. Create a collection for the room
- D. For each room, create a collection that contains a document for each message



- E. Create a collection for the room
- F. For each room, create a document that contains a collection for documents, each of which contains a message.



- G. Create a collection for the rooms, and create a document for each room
- H. Create a separate collection for messages, with one document per message
- I. Each room's document contains a list of references to the messages.



Answer: C

Explanation:

<https://firebase.google.com/docs/firestore/data-model#hierarchical-data>

NEW QUESTION 214

- (Exam Topic 2)

You are developing a JPEG image-resizing API hosted on Google Kubernetes Engine (GKE). Callers of the service will exist within the same GKE cluster. You want clients to be able to get the IP address of the service.

What should you do?

- A. Define a GKE Service
- B. Clients should use the name of the A record in Cloud DNS to find the service's cluster IP address.
- C. Define a GKE Service
- D. Clients should use the service name in the URL to connect to the service.
- E. Define a GKE Endpoint
- F. Clients should get the endpoint name from the appropriate environment variable in the client container.
- G. Define a GKE Endpoint
- H. Clients should get the endpoint name from Cloud DNS.

Answer: C

NEW QUESTION 219

- (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 222

- (Exam Topic 2)

Your company's development teams want to use various open source operating systems in their Docker builds. When images are created in published containers in your company's environment, you need to scan them for Common Vulnerabilities and Exposures (CVEs). The scanning process must not impact software development agility. You want to use managed services where possible. What should you do?

- A. Enable the Vulnerability scanning setting in the Container Registry.
- B. Create a Cloud Function that is triggered on a code check-in and scan the code for CVEs.
- C. Disallow the use of non-commercially supported base images in your development environment.
- D. Use Cloud Monitoring to review the output of Cloud Build to determine whether a vulnerable version has been used.

Answer: A

Explanation:

<https://cloud.google.com/container-analysis/docs/os-overview>

NEW QUESTION 227

.....

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](#)