



Google

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

HipLocal's application uses Cloud Client Libraries to interact with Google Cloud. HipLocal needs to configure authentication and authorization in the Cloud Client Libraries to implement least privileged access for the application. What should they do?

- A. Create an API key
- B. Use the API key to interact with Google Cloud.
- C. Use the default compute service account to interact with Google Cloud.
- D. Create a service account for the application
- E. Export and deploy the private key for the application
- F. Use the service account to interact with Google Cloud.
- G. Create a service account for the application and for each Google Cloud API used by the application. Export and deploy the private keys used by the application
- H. Use the service account with one Google Cloud API to interact with Google Cloud.

Answer: A

NEW QUESTION 2

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

HipLocal is expanding into new locations. They must capture additional data each time the application is launched in a new European country. This is causing delays in the development process due to constant schema changes and a lack of environments for conducting testing on the application changes. How should they resolve the issue while meeting the business requirements?

- A. Create new Cloud SQL instances in Europe and North America for testing and deployment
- B. Provide developers with local MySQL instances to conduct testing on the application changes.
- C. Migrate data to Bigtable
- D. Instruct the development teams to use the Cloud SDK to emulate a local Bigtable development environment.
- E. Move from Cloud SQL to MySQL hosted on Compute Engine
- F. Replicate hosts across regions in the Americas and Europe
- G. Provide developers with local MySQL instances to conduct testing on the application changes.
- H. Migrate data to Firestore in Native mode and set up instances

Answer: B

NEW QUESTION 3

- (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 4

- (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 5

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

A recent security audit discovers that HipLocal's database credentials for their Compute Engine-hosted MySQL databases are stored in plain text on persistent disks. HipLocal needs to reduce the risk of these credentials being stolen. What should they do?

- A. Create a service account and download its key
- B. Use the key to authenticate to Cloud Key Management Service (KMS) to obtain the database credentials.
- C. Create a service account and download its key
- D. Use the key to authenticate to Cloud Key Management Service (KMS) to obtain a key used to decrypt the database credentials.
- E. Create a service account and grant it the roles/iam.serviceAccountUser role
- F. Impersonate as this account and authenticate using the Cloud SQL Proxy.
- G. Grant the roles/secretmanager.secretAccessor role to the Compute Engine service account
- H. Store and access the database credentials with the Secret Manager API.

Answer: D

Explanation:

<https://cloud.google.com/secret-manager/docs/overview>

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'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 9

- (Exam Topic 2)

You are developing an application that will allow clients to download a file from your website for a specific period of time. How should you design the application to complete this task while following Google-recommended best practices?

- A. Configure the application to send the file to the client as an email attachment.
- B. Generate and assign a Cloud Storage-signed URL for the fil
- C. Make the URL available for the client to download.
- D. Create a temporary Cloud Storage bucket with time expiration specified, and give download permissions to the bucke
- E. Copy the file, and send it to the client.
- F. Generate the HTTP cookies with time expiration specifie
- G. If the time is valid, copy the file from the Cloud Storage bucket, and make the file available for the client to download.

Answer: B

NEW QUESTION 10

- (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 10

- (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 15

- (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 19

- (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 20

- (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 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 database.
- H. Deploy the application to a Compute Engine managed instance group with autoscaling.

Answer: A

NEW QUESTION 25

- (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 28

- (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 31

- (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 34

- (Exam Topic 2)

You recently migrated a monolithic application to Google Cloud by breaking it down into microservices. One of the microservices is deployed using Cloud Functions. As you modernize the application, you make a change to the API of the service that is backward-incompatible. You need to support both existing callers who use the original API and new callers who use the new API. What should you do?

- A. Leave the original Cloud Function as-is and deploy a second Cloud Function with the new AP
- B. Use a load balancer to distribute calls between the versions.
- C. Leave the original Cloud Function as-is and deploy a second Cloud Function that includes only the changed AP
- D. Calls are automatically routed to the correct function.
- E. Leave the original Cloud Function as-is and deploy a second Cloud Function with the new AP
- F. Use Cloud Endpoints to provide an API gateway that exposes a versioned API.
- G. Re-deploy the Cloud Function after making code changes to support the new AP
- H. Requests for both versions of the API are fulfilled based on a version identifier included in the call.

Answer: D

Explanation:

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

NEW QUESTION 37

- (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 length
- 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 38

- (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 40

- (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)

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 50

- (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 serves.
- 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 project.
- 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 serves.
- 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 53

- (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 57

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 58

- (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 63

- (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 67

- (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 69

- (Exam Topic 2)

Your team is developing an application in Google Cloud that executes with user identities maintained by Cloud Identity. Each of your application's users will have an associated Pub/Sub topic to which messages are published, and a Pub/Sub subscription where the same user will retrieve published messages. You need to

ensure that only authorized users can publish and subscribe to their own specific Pub/Sub topic and subscription. What should you do?



- A. Bind the user identity to the pubsub.publisher and pubsub.subscriber roles at the resource level.
- B. Grant the user identity the pubsub.publisher and pubsub.subscriber roles at the project level.
- C. Grant the user identity a custom role that contains the pubsub.topics.create and pubsub.subscriptions.create permissions.
- D. Configure the application to run as a service account that has the pubsub.publisher and pubsub.subscriber roles.

Answer: C

NEW QUESTION 70

- (Exam Topic 2)

Your team is writing a backend application to implement the business logic for an interactive voice response (IVR) system that will support a payroll application.

The IVR system has the following technical characteristics:

- Each customer phone call is associated with a unique IVR session.
- The IVR system creates a separate persistent gRPC connection to the backend for each session.
- If the connection is interrupted, the IVR system establishes a new connection, causing a slight latency for that call.

You need to determine which compute environment should be used to deploy the backend application. Using current call data, you determine that:

- Call duration ranges from 1 to 30 minutes.
- Calls are typically made during business hours.
- There are significant spikes of calls around certain known dates (e.g., pay days), or when large payroll changes occur.

You want to minimize cost, effort, and operational overhead. Where should you deploy the backend application?

- A. Compute Engine
- B. Google Kubernetes Engine cluster in Standard mode
- C. Cloud Functions
- D. Cloud Run

Answer: D

Explanation:

This page shows Cloud Run-specific details for developers who want to use gRPC to connect a Cloud Run service with other services, for example, to provide simple, high performance communication between internal microservices. You can use all gRPC types, streaming or unary, with Cloud Run.

Possible use cases include:

Communication between internal microservices.

High loads of data (gRPC uses protocol buffers, which are up to seven times faster than REST calls). Only a simple service definition is needed, you don't want to write a full client library.

Use streaming gRPCs in your gRPC server to build more responsive applications and APIs. <https://cloud.google.com/run/docs/tutorials/secure-services#:~:text=The%20backend%20service%20is%20priva>

NEW QUESTION 72

- (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 73

- (Exam Topic 2)

You are designing an application that consists of several microservices. Each microservice has its own RESTful API and will be deployed as a separate Kubernetes Service. You want to ensure that the consumers of these APIs aren't impacted when there is a change to your API, and also ensure that third-party systems aren't interrupted when new versions of the API are released. How should you configure the connection to the application following Google-recommended best practices?

- A. Use an Ingress that uses the API's URL to route requests to the appropriate backend.
- B. Leverage a Service Discovery system, and connect to the backend specified by the request.
- C. Use multiple clusters, and use DNS entries to route requests to separate versioned backends.
- D. Combine multiple versions in the same service, and then specify the API version in the POST request.

Answer: A

NEW QUESTION 78

- (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 platform.
- 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 80

- (Exam Topic 2)

You need to configure a Deployment on Google Kubernetes Engine (GKE). You want to include a check that verifies that the containers can connect to the database. If the Pod is failing to connect, you want a script on the container to run to complete a graceful shutdown. How should you configure the Deployment?

- A. Create two jobs: one that checks whether the container can connect to the database, and another that runs the shutdown script if the Pod is failing.
- B. Create the Deployment with a livenessProbe for the container that will fail if the container can't connect to the database.
- C. Configure a PreStop lifecycle handler that runs the shutdown script if the container is failing.
- D. Create the Deployment with a PostStart lifecycle handler that checks the service availability.
- E. Configure a PreStop lifecycle handler that runs the shutdown script if the container is failing.
- F. Create the Deployment with an initContainer that checks the service availability.
- G. Configure a PreStop lifecycle handler that runs the shutdown script if the Pod is failing.

Answer: B

Explanation:

<https://cloud.google.com/architecture/best-practices-for-running-cost-effective-kubernetes-applications-on-gke#>

NEW QUESTION 83

- (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 88

- (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 VM.
- B. Take debug snapshots of the functions that take the longest time.
- 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 agent.
- 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 identify time-intensive functions.

Answer: B

NEW QUESTION 91

- (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 92

- (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 94

- (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 99

- (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 100

- (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 105

- (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 110

- (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 114

- (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 116

- (Exam Topic 2)

Your team detected a spike of errors in an application running on Cloud Run in your production project. The application is configured to read messages from Pub/Sub topic A, process the messages, and write the messages to topic B. You want to conduct tests to identify the cause of the errors. You can use a set of mock messages for testing. What should you do?

- A. Deploy the Pub/Sub and Cloud Run emulators on your local machine
- B. Deploy the application locally, and change the logging level in the application to DEBUG or INFO
- C. Write mock messages to topic A, and then analyze the logs.
- D. Use the gcloud CLI to write mock messages to topic
- E. Change the logging level in the application to DEBUG or INFO, and then analyze the logs.

- F. Deploy the Pub/Sub emulator on your local machine
- G. Point the production application to your local Pub/Sub topic
- H. Write mock messages to topic A, and then analyze the logs.
- I. Use the Google Cloud console to write mock messages to topic
- J. Change the logging level in the application to DEBUG or INFO, and then analyze the logs.

Answer: A

NEW QUESTION 118

- (Exam Topic 2)

You are load testing your server application. During the first 30 seconds, you observe that a previously inactive Cloud Storage bucket is now servicing 2000 write requests per second and 7500 read requests per second. Your application is now receiving intermittent 5xx and 429 HTTP responses from the Cloud Storage JSON API as the demand escalates. You want to decrease the failed responses from the Cloud Storage API. What should you do?

- A. Distribute the uploads across a large number of individual storage buckets.
- B. Use the XML API instead of the JSON API for interfacing with Cloud Storage.
- C. Pass the HTTP response codes back to clients that are invoking the uploads from your application.
- D. Limit the upload rate from your application clients so that the dormant bucket's peak request rate is reached more gradually.

Answer: A

Explanation:

Reference: <https://cloud.google.com/storage/docs/request-rate>

NEW QUESTION 120

- (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 123

- (Exam Topic 2)

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

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

Answer: C

Explanation:

Reference: https://cloud.google.com/compute/docs/instances/connecting-advanced#cloud_iap https://cloud.google.com/solutions/connecting-securely#storing_host_keys_by_enabling_guest_attributes

NEW QUESTION 128

- (Exam Topic 2)

You need to copy directory local-scripts and all of its contents from your local workstation to a Compute Engine virtual machine instance. Which command should you use?

- A. `gsutil cp --project "my-gcp-project" -r ~/local-scripts/ gcp-instance-name:~/server-scripts/ --zone "us-east1-b"`
- B. `gsutil cp --project "my-gcp-project" -R ~/local-scripts/ gcp-instance-name:~/server-scripts/ --zone "us-east1-b"`
- C. `gcloud compute scp --project "my-gcp-project" --recurse ~/local-scripts/ gcpinstance-name:~/server-scripts/ --zone "us-east1-b"`
- D. `gcloud compute mv --project "my-gcp-project" --recurse ~/local-scripts/ gcpinstance-name:~/server-scripts/ --zone "us-east1-b"`

Answer: C

Explanation:

Reference: <https://cloud.google.com/sdk/gcloud/reference/compute/copy-files>

NEW QUESTION 129

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 130

- (Exam Topic 2)

You are building a highly available and globally accessible application that will serve static content to users. You need to configure the storage and serving components. You want to minimize management overhead and latency while maximizing reliability for users. What should you do?

- A. 1) Create a managed instance group
- B. Replicate the static content across the virtual machines (VMs)2) Create an external HTTP(S) load balancer.3) Enable Cloud CDN, and send traffic to the managed instance group.
- C. 1) Create an unmanaged instance group
- D. Replicate the static content across the VMs.2) Create an external HTTP(S) load balancer3) Enable Cloud CDN, and send traffic to the unmanaged instance group.
- E. 1) Create a Standard storage class, regional Cloud Storage bucket
- F. Put the static content in the bucket2) Reserve an external IP address, and create an external HTTP(S) load balancer3) Enable Cloud CDN, and send traffic to your backend bucket
- G. 1) Create a Standard storage class, multi-regional Cloud Storage bucket
- H. Put the static content in the bucket.2) Reserve an external IP address, and create an external HTTP(S) load balancer.3) Enable Cloud CDN, and send traffic to your backend bucket.

Answer: D

NEW QUESTION 131

- (Exam Topic 2)

The new version of your containerized application has been tested and is ready to deploy to production on Google Kubernetes Engine. You were not able to fully load-test the new version in pre-production environments, and you need to make sure that it does not have performance problems once deployed. Your deployment must be automated. What should you do?

- A. Use Cloud Load Balancing to slowly ramp up traffic between versions
- B. Use Cloud Monitoring to look for performance issues.
- C. Deploy the application via a continuous delivery pipeline using canary deployment
- D. Use Cloud Monitoring to look for performance issues
- E. and ramp up traffic as the metrics support it.
- F. Deploy the application via a continuous delivery pipeline using blue/green deployment
- G. Use Cloud Monitoring to look for performance issues, and launch fully when the metrics support it.
- H. Deploy the application using kubectl and set the spec.updateStrategy.type to RollingUpdate
- I. Use Cloud Monitoring to look for performance issues, and run the kubectl rollback command if there are any issues.

Answer: C

Explanation:

https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#perform_a_blueg

NEW QUESTION 134

- (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 ["unicorn", "-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 138

- (Exam Topic 2)

You are creating a Google Kubernetes Engine (GKE) cluster and run this command:

```
> gcloud container clusters create large-cluster --num-nodes 200
```

The command fails with the error:

```
insufficient regional quota to satisfy request: resource "CPUS": request requires '200.0' and is short '176.0'. project has a quota of '24.0' with '24.0' available
```

You want to resolve the issue. What should you do?

- A. Request additional GKE quota in the GCP Console.
- B. Request additional Compute Engine quota in the GCP Console.
- C. Open a support case to request additional GKE quota.
- D. Decouple services in the cluster, and rewrite new clusters to function with fewer cores.

Answer: A

NEW QUESTION 139

- (Exam Topic 2)

You have an HTTP Cloud Function that is called via POST. Each submission's request body has a flat, unnested JSON structure containing numeric and text data. After the Cloud Function completes, the collected data should be immediately available for ongoing and complex analytics by many users in parallel. How should you persist the submissions?

- A. Directly persist each POST request's JSON data into Datastore.
- B. Transform the POST request's JSON data, and stream it into BigQuery.
- C. Transform the POST request's JSON data, and store it in a regional Cloud SQL cluster.
- D. Persist each POST request's JSON data as an individual file within Cloud Storage, with the file name containing the request identifier.

Answer: D

NEW QUESTION 140

- (Exam Topic 2)

Your code is running on Cloud Functions in project A. It is supposed to write an object in a Cloud Storage bucket owned by project B. However, the write call is failing with the error "403 Forbidden". What should you do to correct the problem?

- A. Grant your user account the roles/storage.objectCreator role for the Cloud Storage bucket.
- B. Grant your user account the roles/iam.serviceAccountUser role for the service-PROJECTA@gcf-adminrobot.iam.gserviceaccount.com service account.
- C. Grant the service-PROJECTA@gcf-admin-robot.iam.gserviceaccount.com service account the roles/storage.objectCreator role for the Cloud Storage bucket.
- D. Enable the Cloud Storage API in project B.

Answer: B

NEW QUESTION 144

- (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/`.
- E. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url `https://[API_NAME]/[API_VERSION]/`.

Answer: D

NEW QUESTION 146

- (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 147

- (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 148

- (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 blue/green deployment, and test your new application after the deployment is complete.
- E. Perform a canary deployment, and test your new application periodically after the new version is deployed.

Answer: C

NEW QUESTION 153

- (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 155

- (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 156

- (Exam Topic 2)

You are using Cloud Build to build and test application source code stored in Cloud Source Repositories. The build process requires a build tool not available in the Cloud Build environment.

What should you do?

- A. Download the binary from the internet during the build process.
- B. Build a custom cloud builder image and reference the image in your build steps.
- C. Include the binary in your Cloud Source Repositories repository and reference it in your build scripts.
- D. Ask to have the binary added to the Cloud Build environment by filing a feature request against the Cloud Build public Issue Tracker.

Answer: B

NEW QUESTION 159

- (Exam Topic 2)

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

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

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

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

Answer: AC

NEW QUESTION 164

- (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 168

- (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 169


- (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 170

- (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 172

- (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 174

- (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 175

.....

Relate Links

100% Pass Your Professional-Cloud-Developer Exam with ExamBible Prep Materials

<https://www.exambible.com/Professional-Cloud-Developer-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>