

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

<https://www.2passeasy.com/dumps/Professional-Cloud-Developer/>



NEW QUESTION 1

- (Exam Topic 1)

In order to meet their business requirements, how should HipLocal store their application state?

- A. Use local SSDs to store state.
- B. Put a memcache layer in front of MySQL.
- C. Move the state storage to Cloud Spanner.
- D. Replace the MySQL instance with Cloud SQL.

Answer: B

NEW QUESTION 2

- (Exam Topic 1)

Which database should HipLocal use for storing user activity?

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

Answer: A

NEW QUESTION 3

- (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 4

- (Exam Topic 1)

HipLocal's data science team wants to analyze user reviews. How should they prepare the data?

- A. Use the Cloud Data Loss Prevention API for redaction of the review dataset.
- B. Use the Cloud Data Loss Prevention API for de-identification of the review dataset.
- C. Use the Cloud Natural Language Processing API for redaction of the review dataset.
- D. Use the Cloud Natural Language Processing API for de-identification of the review dataset.

Answer: B

Explanation:

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

NEW QUESTION 5

- (Exam Topic 1)

HipLocal is configuring their access controls.

Which firewall configuration should they implement?

- A. Block all traffic on port 443.
- B. Allow all traffic into the network.
- C. Allow traffic on port 443 for a specific tag.
- D. Allow all traffic on port 443 into the network.

Answer: D

NEW QUESTION 6

- (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 rol
- F. Impersonate as this account and authenticate using the Cloud SQL Proxy.
- G. Grant the roles/secretmanager.secretAccessor role to the Compute Engine service accoun
- 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 7

- (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 8

- (Exam Topic 1)

HipLocal's .net-based auth service fails under intermittent load. What should they do?

- A. Use App Engine for autoscaling.
- B. Use Cloud Functions for autoscaling.
- C. Use a Compute Engine cluster for the service.
- D. Use a dedicated Compute Engine virtual machine instance for the service.

Answer: D

Explanation:

Reference: <https://www.qwiklabs.com/focuses/611?parent=catalog>

NEW QUESTION 9

- (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 10

- (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 10

- (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 11

- (Exam Topic 2)

Users are complaining that your Cloud Run-hosted website responds too slowly during traffic spikes. You want to provide a better user experience during traffic peaks. What should you do?

- A. Read application configuration and static data from the database on application startup.
- B. Package application configuration and static data into the application image during build time.
- C. Perform as much work as possible in the background after the response has been returned to the user.
- D. Ensure that timeout exceptions and errors cause the Cloud Run instance to exit quickly so a replacement instance can be started.

Answer: C

NEW QUESTION 15

- (Exam Topic 2)

You are developing an application hosted on Google Cloud that uses a MySQL relational database schema. The application will have a large volume of reads and writes to the database and will require backups and ongoing capacity planning. Your team does not have time to fully manage the database but can take on small administrative tasks. How should you host the database?

- A. Configure Cloud SQL to host the database, and import the schema into Cloud SQL.
- B. Deploy MySQL from the Google Cloud Marketplace to the database using a client, and import the schema.
- C. Configure Bigtable to host the database, and import the data into Bigtable.
- D. Configure Cloud Spanner to host the database, and import the schema into Cloud Spanner.
- E. Configure Firestore to host the database, and import the data into Firestore.

Answer: A

Explanation:

<https://cloud.google.com/spanner/docs/migrating-mysql-to-spanner#migration-process>

Cloud SQL: Cloud SQL is a web service that allows you to create, configure, and use relational databases that live in Google's cloud. It is a fully-managed service that maintains, manages, and administers your databases, allowing you to focus on your applications and services.

<https://cloud.google.com/sql/docs/mysql> Cloud SQL for MySQL is a fully-managed database service that helps you set up, maintain, manage, and administer your MySQL relational databases on Google Cloud Platform.

NEW QUESTION 16

- (Exam Topic 2)

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

- <https://yourcompany.com/students>
- <https://yourcompany.com/teachers>
- <https://yourcompany.com/classes>

You need to configure each API URL path to invoke a different function in your code. What should you do?

- A. Create one Cloud Function as a backend service exposed using an HTTPS load balancer.
- B. Create three Cloud Functions exposed directly.
- C. Create one Cloud Function exposed directly.
- D. Create three Cloud Functions as three backend services exposed using an HTTPS load balancer.

Answer: D

Explanation:

<https://cloud.google.com/load-balancing/docs/https/setup-global-ext-https-serverless>

NEW QUESTION 18

- (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 file.
- 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 bucket.
- E. Copy the file, and send it to the client.
- F. Generate the HTTP cookies with time expiration specified.
- 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 23

- (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 26

- (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 databas
- B. Deploy the application to Cloud Run.
- C. Build the application in C# with Firestore as the databas
- D. Deploy the application to App Engine flexible environment.
- E. Build the application in Python with CloudSQL as the databas
- F. Deploy the application to App Engine standard environment.
- G. Build the application in Python with Firestore as the databas
- H. Deploy the application to a Compute Engine managed instance group with autoscaling.

Answer: A

NEW QUESTION 31

- (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 35

- (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 theAPI 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 40

- (Exam Topic 2)

You want to re-architect a monolithic application so that it follows a microservices model. You want to accomplish this efficiently while minimizing the impact of this change to the business.

Which approach should you take?

- A. Deploy the application to Compute Engine and turn on autoscaling.
- B. Replace the application's features with appropriate microservices in phases.
- C. Refactor the monolithic application with appropriate microservices in a single effort and deploy it.
- D. Build a new application with the appropriate microservices separate from the monolith and replace it when it is complete.

Answer: C

Explanation:

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

NEW QUESTION 44

- (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 48

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 50

- (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 55

- (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 59

- (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 60

- (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 63

- (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 67

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 69

- (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 service.
- 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 73

- (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 77

- (Exam Topic 2)

You are building a mobile application that will store hierarchical data structures in a database. The application will enable users working offline to sync changes when they are back online. A backend service will enrich the data in the database using a service account. The application is expected to be very popular and needs to scale seamlessly and securely. Which database and IAM role should you use?

- A. Use Cloud SQL, and assign the roles/cloudsql.editor role to the service account.
- B. Use Bigtable, and assign the roles/bigtable.viewer role to the service account.
- C. Use Firestore in Native mode and assign the roles/datastore.user role to the service account.

D. Use Firestore in Datastore mode and assign the roles/datastore.viewer role to the service account.

Answer: C

Explanation:

<https://firebase.google.com/docs/firestore/manage-data/enable-offline>

Cloud Firestore supports offline data persistence. This feature caches a copy of the Cloud Firestore data that your app is actively using, so your app can access the data when the device is offline. You can write, read, listen to, and query the cached data. When the device comes back online, Cloud Firestore synchronizes any local changes made by your app to the Cloud Firestore backend.

NEW QUESTION 78

- (Exam Topic 2)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 82

- (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 loa
- F. Use Cloud Scheduler to run 60,000 Cloud Task transactions per minute for 10 minute
- 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 servic
- I. Analyze the results using Cloud Trace.

Answer: B

Explanation:

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

NEW QUESTION 86

- (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 88

- (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 89

- (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 90

- (Exam Topic 2)

You have an application that uses an HTTP Cloud Function to process user activity from both desktop browser and mobile application clients. This function will serve as the endpoint for all metric submissions using HTTP POST.

Due to legacy restrictions, the function must be mapped to a domain that is separate from the domain requested by users on web or mobile sessions. The domain for the Cloud Function is <https://fn.example.com>. Desktop and mobile clients use the domain <https://www.example.com>. You need to add a header to the function's HTTP response so that only those browser and mobile sessions can submit metrics to the Cloud Function. Which response header should you add?

- A. Access-Control-Allow-Origin: *
- B. Access-Control-Allow-Origin: https://*.example.com
- C. Access-Control-Allow-Origin: <https://fn.example.com>
- D. Access-Control-Allow-origin: <https://www.example.com>

Answer: D

NEW QUESTION 95

- (Exam Topic 2)

You have two tables in an ANSI-SQL compliant database with identical columns that you need to quickly combine into a single table, removing duplicate rows from the result set.

What should you do?

- A. Use the JOIN operator in SQL to combine the tables.
- B. Use nested WITH statements to combine the tables.
- C. Use the UNION operator in SQL to combine the tables.
- D. Use the UNION ALL operator in SQL to combine the tables.

Answer: C

Explanation:

Reference: https://www.techonthenet.com/sql/union_all.php

NEW QUESTION 97

- (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 98

- (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 99

- (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 /workspac
- C. Read from the same /workspace file in the subsequent build step.
- D. Use gsutil to output the file contents to a Cloud Storage objec
- 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 builde
- 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 103

- (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 105

- (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 107

- (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 111

- (Exam Topic 2)

You are designing an application that will subscribe to and receive messages from a single Pub/Sub topic and insert corresponding rows into a database. Your application runs on Linux and leverages preemptible virtual machines to reduce costs. You need to create a shutdown script that will initiate a graceful shutdown. What should you do?

- A. Write a shutdown script that uses inter-process signals to notify the application process to disconnect from the database.
- B. Write a shutdown script that broadcasts a message to all signed-in users that the Compute Engine instance is going down and instructs them to save current work and sign out.
- C. Write a shutdown script that writes a file in a location that is being polled by the application once every five minute
- D. After the file is read, the application disconnects from the database.
- E. Write a shutdown script that publishes a message to the Pub/Sub topic announcing that a shutdown is in progres
- F. After the application reads the message, it disconnects from the database.

Answer: D

NEW QUESTION 113

- (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 118

- (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 122

- (Exam Topic 2)

You are working on a social media application. You plan to add a feature that allows users to upload images. These images will be 2 MB – 1 GB in size. You want to minimize their infrastructure operations overhead for this feature. What should you do?

- A. Change the application to accept images directly and store them in the database that stores other user information.
- B. Change the application to create signed URLs for Cloud Storag
- C. Transfer these signed URLs to the client application to upload images to Cloud Storage.
- D. Set up a web server on GCP to accept user images and create a file store to keep uploaded file
- E. Change the application to retrieve images from the file store.
- F. Create a separate bucket for each user in Cloud Storag
- G. Assign a separate service account to allow write access on each bucke
- H. Transfer service account credentials to the client application based on user informatio
- I. The application uses this service account to upload images to Cloud Storage.

Answer: B

Explanation:

Reference:

<https://cloud.google.com/blog/products/storage-data-transfer/uploading-images-directly-to-cloud-storage-by-usi>

NEW QUESTION 127

- (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 128

- (Exam Topic 2)

You want to migrate an on-premises container running in Knative to Google Cloud. You need to make sure that the migration doesn't affect your application's deployment strategy, and you want to use a fully managed service. Which Google Cloud service should you use to deploy your container?

- A. Cloud Run
- B. Compute Engine
- C. Google Kubernetes Engine
- D. App Engine flexible environment

Answer: A

Explanation:

<https://cloud.google.com/blog/products/serverless/knative-based-cloud-run-services-are-ga>

NEW QUESTION 129

- (Exam Topic 2)

You are in the final stage of migrating an on-premises data center to Google Cloud. You are quickly approaching your deadline, and discover that a web API is running on a server slated for decommissioning. You need to recommend a solution to modernize this API while migrating to Google Cloud. The modernized web API must meet the following requirements:

- Autoscales during high traffic periods at the end of each month
- Written in Python 3.x
- Developers must be able to rapidly deploy new versions in response to frequent code changes

You want to minimize cost, effort, and operational overhead of this migration. What should you do?

- A. Modernize and deploy the code on App Engine flexible environment.
- B. Modernize and deploy the code on App Engine standard environment.
- C. Deploy the modernized application to an n1-standard-1 Compute Engine instance.
- D. Ask the development team to re-write the application to run as a Docker container on Google Kubernetes Engine.

Answer: B

Explanation:

<https://cloud.google.com/appengine/docs/standard>

NEW QUESTION 132

- (Exam Topic 2)

You are developing an internal application that will allow employees to organize community events within your company. You deployed your application on a single Compute Engine instance. Your company uses Google Workspace (formerly G Suite), and you need to ensure that the company employees can authenticate to the application from anywhere. What should you do?

- A. Add a public IP address to your instance, and restrict access to the instance using firewall rule
- B. Allow your company's proxy as the only source IP address.
- C. Add an HTTP(S) load balancer in front of the instance, and set up Identity-Aware Proxy (IAP). Configure the IAP settings to allow your company domain to access the website.
- D. Set up a VPN tunnel between your company network and your instance's VPC location on Google Clou
- E. Configure the required firewall rules and routing information to both the on-premises and Google Cloud networks.
- F. Add a public IP address to your instance, and allow traffic from the interne
- G. Generate a random hash, and create a subdomain that includes this hash and points to your instanc
- H. Distribute this DNS address to your company's employees.

Answer: B

Explanation:

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

NEW QUESTION 135

- (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 139

- (Exam Topic 2)

You are a developer working on an internal application for payroll processing. You are building a component of the application that allows an employee to submit a timesheet, which then initiates several steps:

- An email is sent to the employee and manager, notifying them that the timesheet was submitted.
- A timesheet is sent to payroll processing for the vendor's API.
- A timesheet is sent to the data warehouse for headcount planning.

These steps are not dependent on each other and can be completed in any order. New steps are being considered and will be implemented by different development teams. Each development team will implement the error handling specific to their step. What should you do?

- A. Deploy a Cloud Function for each step that calls the corresponding downstream system to complete the required action.
- B. Create a Pub/Sub topic for each step
- C. Create a subscription for each downstream development team to subscribe to their step's topic.
- D. Create a Pub/Sub topic for timesheet submission
- E. Create a subscription for each downstream development team to subscribe to the topic.
- F. Create a timesheet microservice deployed to Google Kubernetes Engine
- G. The microservice calls each downstream step and waits for a successful response before calling the next step.

Answer: C

NEW QUESTION 140

- (Exam Topic 2)

You are developing an application that will handle requests from end users. You need to secure a Cloud Function called by the application to allow authorized end users to authenticate to the function via the application while restricting access to unauthorized users. You will integrate Google Sign-In as part of the solution and want to follow Google-recommended best practices. What should you do?

- A. Deploy from a source code repository and grant users the roles/cloudfunctions.viewer role.
- B. Deploy from a source code repository and grant users the roles/cloudfunctions.invoker role
- C. Deploy from your local machine using gcloud and grant users the roles/cloudfunctions.admin role
- D. Deploy from your local machine using gcloud and grant users the roles/cloudfunctions.developer role

Answer: C

NEW QUESTION 144

- (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 149

- (Exam Topic 2)

Your company has a data warehouse that keeps your application information in BigQuery. The BigQuery data warehouse keeps 2 PBs of user data. Recently, your company expanded your user base to include EU users and needs to comply with these requirements:

Your company must be able to delete all user account information upon user request. All EU user data must be stored in a single region specifically for EU users. Which two actions should you take? (Choose two.)

- A. Use BigQuery federated queries to query data from Cloud Storage.
- B. Create a dataset in the EU region that will keep information about EU users only.
- C. Create a Cloud Storage bucket in the EU region to store information for EU users only.
- D. Re-upload your data using to a Cloud Dataflow pipeline by filtering your user records out.
- E. Use DML statements in BigQuery to update/delete user records based on their requests.

Answer: CE

Explanation:

Reference: <https://cloud.google.com/solutions/bigquery-data-warehouse>

NEW QUESTION 151

- (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 155

- (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 balance
- 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 160

- (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 163

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 167

- (Exam Topic 2)

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

- A. Blue/green deployment
- B. Canary deployment
- C. Rolling deployment
- D. Recreate deployment

Answer: A

NEW QUESTION 168

- (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 171

- (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 176

- (Exam Topic 2)

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

- A. Use Cloud Logging to check for error logs, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- B. Use Cloud Trace to monitor the requests per sec of incoming requests to Spanner, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- C. Use Cloud Monitoring to monitor the CPU utilization, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- D. Use Snapshot Debugger to check for application errors, and reduce Spanner processing units by small increments until you find the minimum capacity required.

Answer: C

Explanation:

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

NEW QUESTION 179

- (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 184

- (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 188

- (Exam Topic 2)

Your team develops services that run on Google Kubernetes Engine. You need to standardize their log data using Google-recommended practices and make the data more useful in the fewest number of steps. What should you do? (Choose two.)

- A. Create aggregated exports on application logs to BigQuery to facilitate log analytics.
- B. Create aggregated exports on application logs to Cloud Storage to facilitate log analytics.
- C. Write log output to standard output (stdout) as single-line JSON to be ingested into Cloud Logging as structured logs.
- D. Mandate the use of the Logging API in the application code to write structured logs to Cloud Logging.
- E. Mandate the use of the Pub/Sub API to write structured data to Pub/Sub and create a Dataflow streaming pipeline to normalize logs and write them to BigQuery for analytics.

Answer: AC

Explanation:

https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs#best_practices

NEW QUESTION 191

- (Exam Topic 2)

You are a SaaS provider deploying dedicated blogging software to customers in your Google Kubernetes Engine (GKE) cluster. You want to configure a secure multi-tenant platform to ensure that each customer has access to only their own blog and can't affect the workloads of other customers. What should you do?

- A. Enable Application-layer Secrets on the GKE cluster to protect the cluster.
- B. Deploy a namespace per tenant and use Network Policies in each blog deployment.
- C. Use GKE Audit Logging to identify malicious containers and delete them on discovery.
- D. Build a custom image of the blogging software and use Binary Authorization to prevent untrusted image deployments.

Answer: B

Explanation:

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

NEW QUESTION 196

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 197

- (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 200

- (Exam Topic 2)

You have containerized a legacy application that stores its configuration on an NFS share. You need to deploy this application to Google Kubernetes Engine (GKE) and do not want the application serving traffic until after the configuration has been retrieved. What should you do?

- A. Use the gsutil utility to copy files from within the Docker container at startup, and start the service using an ENTRYPOINT script.
- B. Create a PersistentVolumeClaim on the GKE cluster.

- C. Access the configuration files from the volume, and start the service using an ENTRYPOINT script.
- D. Use the COPY statement in the Dockerfile to load the configuration into the container image
- E. Verify that the configuration is available, and start the service using an ENTRYPOINT script.
- F. Add a startup script to the GKE instance group to mount the NFS share at node startup
- G. Copy the configuration files into the container, and start the service using an ENTRYPOINT script.

Answer: D

Explanation:

Reference: <https://cloud.google.com/compute/docs/instances/startup-scripts/linux>

NEW QUESTION 201

- (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 202

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Professional-Cloud-Developer Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Professional-Cloud-Developer Product From:

<https://www.2passeasy.com/dumps/Professional-Cloud-Developer/>

Money Back Guarantee

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