

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

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



NEW QUESTION 1

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

Which Google Cloud product addresses HipLocal's business requirements for service level indicators and objectives?

- A. Cloud Profiler
- B. Cloud Monitoring
- C. Cloud Trace
- D. Cloud Logging

Answer: B

Explanation:

<https://cloud.google.com/stackdriver/docs/solutions/slo-monitoring#defn-sli>

NEW QUESTION 2

- (Exam Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 3

- (Exam Topic 1)

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 4

- (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 5

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 6

- (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 7

- (Exam Topic 2)

You have written a Cloud Function that accesses other Google Cloud resources. You want to secure the environment using the principle of least privilege. What should you do?

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 8

- (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 9

- (Exam Topic 2)

You have an application controlled by a managed instance group. When you deploy a new version of the application, costs should be minimized and the number of instances should not increase. You want to ensure that, when each new instance is created, the deployment only continues if the new instance is healthy. What should you do?

- A. Perform a rolling-action with maxSurge set to 1, maxUnavailable set to 0.
- B. Perform a rolling-action with maxSurge set to 0, maxUnavailable set to 1
- C. Perform a rolling-action with maxHealthy set to 1, maxUnhealthy set to 0.
- D. Perform a rolling-action with maxHealthy set to 0, maxUnhealthy set to 1.

Answer: A

Explanation:

Reference:

<https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups>

NEW QUESTION 10

- (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 10

- (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 15

- (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 16

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 20

- (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 25

- (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 cluster.
- H. Move your data to Cloud Storage and leverage the Cloud Dataproc connector to run jobs on that data.

Answer: D

NEW QUESTION 29

- (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 namespaces.
- 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-limits>

NEW QUESTION 33

- (Exam Topic 2)

Your team manages a Google Kubernetes Engine (GKE) cluster where an application is running. A different team is planning to integrate with this application. Before they start the integration, you need to ensure that the other team cannot make changes to your application, but they can deploy the integration on GKE. What should you do?

- A. Using Identity and Access Management (IAM), grant the Viewer IAM role on the cluster project to the other team.
- B. Create a new GKE cluster.
- C. Using Identity and Access Management (IAM), grant the Editor role on the cluster project to the other team.
- D. Create a new namespace in the existing cluster.
- E. Using Identity and Access Management (IAM), grant the Editor role on the cluster project to the other team.
- F. Create a new namespace in the existing cluster.
- G. Using Kubernetes role-based access control (RBAC), grant the Admin role on the new namespace to the other team.

Answer: D

NEW QUESTION 35

- (Exam Topic 2)

You are developing a new public-facing application that needs to retrieve specific properties in the metadata of users' objects in their respective Cloud Storage buckets. Due to privacy and data residency requirements, you must retrieve only the metadata and not the object data. You want to maximize the performance of the retrieval process. How should you retrieve the metadata?

- A. Use the patch method.
- B. Use the compose method.
- C. Use the copy method.
- D. Use the fields request parameter.

Answer: D

Explanation:

https://cloud.google.com/storage/docs/json_api/v1/objects/get

NEW QUESTION 37

- (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 API.
- 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 API.
- 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 API.
- 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 API.
- 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 42

- (Exam Topic 2)

You are creating an App Engine application that writes a file to any user's Google Drive. How should the application authenticate to the Google Drive API?

- A. With an OAuth Client ID that uses the <https://www.googleapis.com/auth/drive.file> scope to obtain an access token for each user.
- B. With an OAuth Client ID with delegated domain-wide authority.
- C. With the App Engine service account and <https://www.googleapis.com/auth/drive.file> scope that generates a signed JWT.
- D. With the App Engine service account with delegated domain-wide authority.

Answer: B

Explanation:

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

NEW QUESTION 45

- (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 50

- (Exam Topic 2)

You are porting an existing Apache/MySQL/PHP application stack from a single machine to Google Kubernetes Engine. You need to determine how to containerize the application. Your approach should follow Google-recommended best practices for availability. What should you do?

- A. Package each component in a separate container.
- B. Implement readiness and liveness probes.
- C. Package the application in a single container.
- D. Use a process management tool to manage each component.
- E. Package each component in a separate container.
- F. Use a script to orchestrate the launch of the components.
- G. Package the application in a single container.
- H. Use a bash script as an entrypoint to the container, and then spawn each component as a background job.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/containers-kubernetes/7-best-practices-for-building-containers> <https://cloud.google.com/architecture/best-practices-for-building-containers>

"classic Apache/MySQL/PHP stack: you might be tempted to run all the components in a single container. However, the best practice is to use two or three different containers: one for Apache, one for MySQL, and potentially one for PHP if you are running PHP-FPM."

NEW QUESTION 55

- (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 58

- (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 60

- (Exam Topic 2)

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

- A. Install the Stackdriver Logging Agent and configure it to send the application logs.
- B. Use a Stackdriver Logging Library to log directly from the application to Stackdriver Logging.

- C. Provide the log file folder path in the metadata of the instance to configure it to send the application logs.
- D. Change the application to log to /var/log so that its logs are automatically sent to Stackdriver Logging.

Answer: A

NEW QUESTION 63

- (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 68

- (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 72

- (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 Run
- B. Analyze the VPC Flow Logs using Cloud Logging.
- C. Create a Google Kubernetes Engine cluster running the Locust or JMeter images to dynamically generate load test
- D. Analyze the results using Cloud Trace.
- E. Create a Cloud Task to generate a test load
- F. Use Cloud Scheduler to run 60,000 Cloud Task transactions per minute for 10 minutes
- G. Analyze the results using Cloud Monitoring.
- H. Create a Compute Engine instance that uses a LAMP stack image from the Marketplace, and use Apache Bench to generate load tests against the service
- I. Analyze the results using Cloud Trace.

Answer: B

Explanation:

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

NEW QUESTION 74

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 78

- (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 82

- (Exam Topic 2)

Your web application is deployed to the corporate intranet. You need to migrate the web application to Google Cloud. The web application must be available only to company employees and accessible to employees as they travel. You need to ensure the security and accessibility of the web application while minimizing application changes. What should you do?

- A. Configure the application to check authentication credentials for each HTTP(S) request to the application.
- B. Configure Identity-Aware Proxy to allow employees to access the application through its public IP address.
- C. Configure a Compute Engine instance that requests users to log in to their corporate account.
- D. Change the web application DNS to point to the proxy Compute Engine instance.
- E. After authenticating, the Compute Engine instance forwards requests to and from the web application.
- F. Configure a Compute Engine instance that requests users to log in to their corporate account.
- G. Change the web application DNS to point to the proxy Compute Engine instance.
- H. After authenticating, the Compute Engine issues an HTTP redirect to a public IP address hosting the web application.

Answer: B

NEW QUESTION 86

- (Exam Topic 2)

You are deploying a microservices application to Google Kubernetes Engine (GKE). The application will receive daily updates. You expect to deploy a large number of distinct containers that will run on the Linux operating system (OS). You want to be alerted to any known OS vulnerabilities in the new containers. You want to follow Google-recommended best practices. What should you do?

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

Answer: B

Explanation:

<https://cloud.google.com/container-analysis/docs/automated-scanning-howto> <https://cloud.google.com/container-analysis/docs/os-overview> says: The Container Scanning API allows you to automate OS vulnerability detection, scanning each time you push an image to Container Registry or Artifact Registry. Enabling this API also triggers language package scans for Go and Java vulnerabilities (Preview).

NEW QUESTION 89

- (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 92

- (Exam Topic 2)

You are a developer at a large organization. You have an application written in Go running in a production Google Kubernetes Engine (GKE) cluster. You need to add a new feature that requires access to BigQuery. You want to grant BigQuery access to your GKE cluster following Google-recommended best practices. What should you do?

- A. Create a Google service account with BigQuery access.
- B. Add the JSON key to Secret Manager, and use the Go client library to access the JSON key.

- C. Create a Google service account with BigQuery acces
- D. Add the Google service account JSON key as a Kubernetes secret, and configure the application to use this secret.
- E. Create a Google service account with BigQuery acces
- F. Add the Google service account JSON key to Secret Manager, and use an init container to access the secret for the application to use.
- G. Create a Google service account and a Kubernetes service accoun
- H. Configure Workload Identity on the GKE cluster, and reference the Kubernetes service account on the application Deployment.

Answer: D

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity#what_is

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

NEW QUESTION 94

- (Exam Topic 2)

You are developing an ecommerce application that stores customer, order, and inventory data as relational tables inside Cloud Spanner. During a recent load test, you discover that Spanner performance is not scaling linearly as expected. Which of the following is the cause?

- A. The use of 64-bit numeric types for 32-bit numbers.
- B. The use of the STRING data type for arbitrary-precision values.
- C. The use of Version 1 UUIDs as primary keys that increase monotonically.
- D. The use of LIKE instead of STARTS_WITH keyword for parameterized SQL queries.

Answer: C

NEW QUESTION 97

- (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 99

- (Exam Topic 2)

You are running a web application on Google Kubernetes Engine that you inherited. You want to determine whether the application is using libraries with known vulnerabilities or is vulnerable to XSS attacks. Which service should you use?

- A. Google Cloud Armor
- B. Debugger
- C. Web Security Scanner
- D. Error Reporting

Answer: C

Explanation:

<https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview>

Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to exercise as many user inputs and event handlers as possible.

NEW QUESTION 104

- (Exam Topic 2)

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

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

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

Which improvement should you suggest your teammate make?

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

Answer: B

NEW QUESTION 105

- (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 110

- (Exam Topic 2)

You are using Cloud Build build to promote a Docker image to Development, Test, and Production environments. You need to ensure that the same Docker image is deployed to each of these environments. How should you identify the Docker image in your build?

- A. Use the latest Docker image tag.
- B. Use a unique Docker image name.
- C. Use the digest of the Docker image.
- D. Use a semantic version Docker image tag.

Answer: D

NEW QUESTION 114

- (Exam Topic 2)

You are running a containerized application on Google Kubernetes Engine. Your container images are stored in Container Registry. Your team uses CI/CD practices. You need to prevent the deployment of containers with known critical vulnerabilities. What should you do?

- A. • Use Web Security Scanner to automatically crawl your application• Review your application logs for scan results, and provide an attestation that the container is free of known critical vulnerabilities• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed
- B. • Use Web Security Scanner to automatically crawl your application• Review the scan results in the scan details page in the Cloud Console, and provide an attestation that the container is free of known critical vulnerabilities• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed
- C. • Enable the Container Scanning API to perform vulnerability scanning• Review vulnerability reporting in Container Registry in the Cloud Console, and provide an attestation that the container is free of known critical vulnerabilities• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed
- D. • Enable the Container Scanning API to perform vulnerability scanning• Programmatically review vulnerability reporting through the Container Scanning API, and provide an attestation that the container is free of known critical vulnerabilities• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

Answer: D

Explanation:

<https://cloud.google.com/binary-authorization/docs/creating-attestations-kritis>

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

NEW QUESTION 118

- (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 123

- (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 126

- (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 129

- (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 130

- (Exam Topic 2)

You are developing an application that reads credit card data from a Pub/Sub subscription. You have written code and completed unit testing. You need to test the Pub/Sub integration before deploying to Google Cloud. What should you do?

- A. Create a service to publish messages, and deploy the Pub/Sub emulator
- B. Generate random content in the publishing service, and publish to the emulator.
- C. Create a service to publish messages to your application
- D. Collect the messages from Pub/Sub in production, and replay them through the publishing service.
- E. Create a service to publish messages, and deploy the Pub/Sub emulator
- F. Collect the messages from Pub/Sub in production, and publish them to the emulator.
- G. Create a service to publish messages, and deploy the Pub/Sub emulator
- H. Publish a standard set of testing messages from the publishing service to the emulator.

Answer: D

NEW QUESTION 134

- (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 136

- (Exam Topic 2)

Your team develops services that run on Google Kubernetes Engine. Your team's code is stored in Cloud Source Repositories. You need to quickly identify bugs in the code before it is deployed to production. You want to invest in automation to improve developer feedback and make the process as efficient as possible. What should you do?

- A. Use Spinnaker to automate building container images from code based on Git tags.
- B. Use Cloud Build to automate building container images from code based on Git tags.
- C. Use Spinnaker to automate deploying container images to the production environment.
- D. Use Cloud Build to automate building container images from code based on forked versions.

Answer: A

Explanation:

Reference: <https://spinnaker.io/docs/guides/tutorials/codelabs/kubernetes-v2-source-to-prod/>

NEW QUESTION 140

- (Exam Topic 2)

You are a developer working with the CI/CD team to troubleshoot a new feature that your team introduced. The CI/CD team used HashiCorp Packer to create a new Compute Engine image from your development branch. The image was successfully built, but is not booting up. You need to investigate the issue with the CI/CD team. What should you do?

- A. Create a new feature branch, and ask the build team to rebuild the image.
- B. Shut down the deployed virtual machine, export the disk, and then mount the disk locally to access the boot logs.
- C. Install Packer locally, build the Compute Engine image locally, and then run it in your personal Google Cloud project.
- D. Check Compute Engine OS logs using the serial port, and check the Cloud Logging logs to confirm access to the serial port.

Answer: D

Explanation:

<https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-using-serial-console>

NEW QUESTION 145

- (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 148

- (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 153

- (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 158

- (Exam Topic 2)

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

instance_class: B1

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

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

Answer: C

NEW QUESTION 163

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 168

- (Exam Topic 2)

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

Which GKE object should you use?

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

Answer: B

Explanation:

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

NEW QUESTION 170

- (Exam Topic 2)

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

```

export NAME=load-balancer

# create network
gcloud compute networks create ${NAME}

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

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

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

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

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

```

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

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

Answer: C

Explanation:

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

NEW QUESTION 171

- (Exam Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 174

- (Exam Topic 2)

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

How should you handle this error?

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

Answer: C

Explanation:

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

NEW QUESTION 179

- (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 183

- (Exam Topic 2)

Your development team has built several Cloud Functions using Java along with corresponding integration and service tests. You are building and deploying the functions and launching the tests using Cloud Build. Your Cloud Build job is reporting deployment failures immediately after successfully validating the code. What should you do?

- A. Check the maximum number of Cloud Function instances.
- B. Verify that your Cloud Build trigger has the correct build parameters.
- C. Retry the tests using the truncated exponential backoff polling strategy.
- D. Verify that the Cloud Build service account is assigned the Cloud Functions Developer role.

Answer: D

Explanation:

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

NEW QUESTION 186

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 190

- (Exam Topic 2)

You have an application running in App Engine. Your application is instrumented with Stackdriver Trace. The `/product-details` request reports details about four known unique products at `/sku-details` as shown below. You want to reduce the time it takes for the request to complete. What should you do?

Timeline



- A. Increase the size of the instance class.
- B. Change the Persistent Disk type to SSD.
- C. Change `/product-details` to perform the requests in parallel.
- D. Store the `/sku-details` information in a database, and replace the webservice call with a database query.

Answer: C

NEW QUESTION 195

- (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 199

- (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 202

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 207

- (Exam Topic 2)

You have an on-premises application that authenticates to the Cloud Storage API using a user-managed service account with a user-managed key. The application connects to Cloud Storage using Private Google Access over a Dedicated Interconnect link. You discover that requests from the application to access objects in the Cloud Storage bucket are failing with a 403 Permission Denied error code. What is the likely cause of this issue?

- A. The folder structure inside the bucket and object paths have changed.
- B. The permissions of the service account's predefined role have changed.
- C. The service account key has been rotated but not updated on the application server.
- D. The Interconnect link from the on-premises data center to Google Cloud is experiencing a temporary outage.

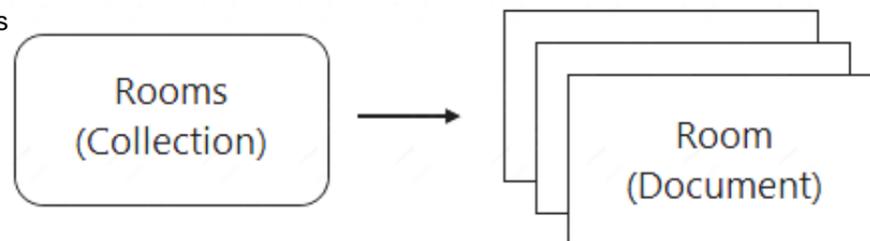
Answer: C

NEW QUESTION 212

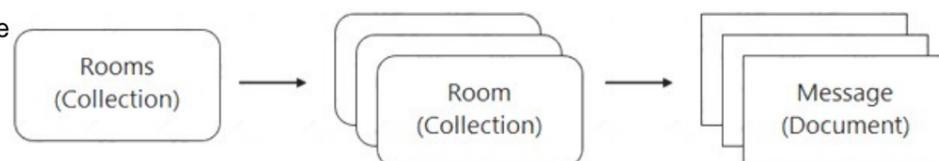
- (Exam Topic 2)

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?

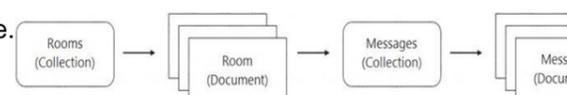
- A. Create a collection for the room
- B. For each room, create a document that lists the contents of the messages



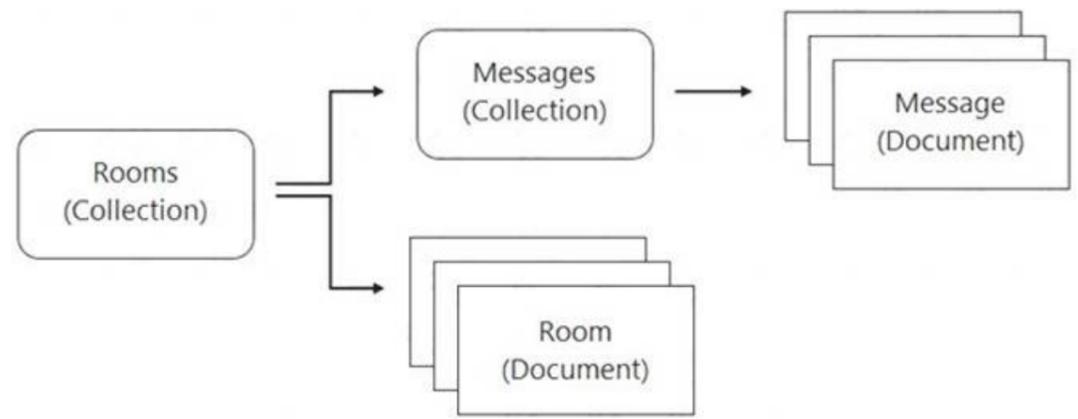
- C. Create a collection for the room
- D. For each room, create a collection that contains a document for each message



- 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 217

- (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 219

- (Exam Topic 2)

Your team is developing unit tests for Cloud Function code. The code is stored in a Cloud Source Repositories repository. You are responsible for implementing the tests. Only a specific service account has the necessary permissions to deploy the code to Cloud Functions. You want to ensure that the code cannot be deployed without first passing the tests. How should you configure the unit testing process?

- A. Configure Cloud Build to deploy the Cloud Function
- B. If the code passes the tests, a deployment approval is sent to you.
- C. Configure Cloud Build to deploy the Cloud Function, using the specific service account as the build agent
- D. Run the unit tests after successful deployment.
- E. Configure Cloud Build to run the unit test
- F. If the code passes the tests, the developer deploys the Cloud Function.
- G. Configure Cloud Build to run the unit tests, using the specific service account as the build agent
- H. If the code passes the tests, Cloud Build deploys the Cloud Function.

Answer: D

NEW QUESTION 223

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 228

- (Exam Topic 2)

You want to upload files from an on-premises virtual machine to Google Cloud Storage as part of a data migration. These files will be consumed by Cloud DataProc Hadoop cluster in a GCP environment. Which command should you use?

- A. `gsutil cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/`
- B. `gcloud cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/`
- C. `hadoop fs cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/`

D. gcloud dataproc cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/

Answer: A

Explanation:

The gsutil cp command allows you to copy data between your local file. storage. boto files generated by running "gsutil config"

NEW QUESTION 230

.....

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