

# Google

## Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer



#### NEW QUESTION 1

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

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

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

**Answer: A**

#### NEW QUESTION 2

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

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

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

**Answer: B**

#### NEW QUESTION 3

- (Exam Topic 1)

HipLocal wants to improve the resilience of their MySQL deployment, while also meeting their business and technical requirements.

Which configuration should they choose?

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

**Answer: B**

#### NEW QUESTION 4

- (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 5

- (Exam Topic 1)

HipLocal has connected their Hadoop infrastructure to GCP using Cloud Interconnect in order to query data stored on persistent disks.

Which IP strategy should they use?

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

**Answer: A**

#### NEW QUESTION 6

- (Exam Topic 1)

Which service should HipLocal use to enable access to internal apps?

- A. Cloud VPN

- B. Cloud Armor
- C. Virtual Private Cloud
- D. Cloud Identity-Aware Proxy

**Answer:** D

**Explanation:**

Reference: <https://cloud.google.com/iap/docs/cloud-iap-for-on-prem-apps-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)

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

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

**Answer:** C

**NEW QUESTION 10**

- (Exam Topic 2)

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

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

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

**Answer:** A

**NEW QUESTION 10**

- (Exam Topic 2)

You 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 11**

- (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 14**

- (Exam Topic 2)

You are planning to add unit tests to your application. You need to be able to assert that published Pub/Sub messages are processed by your subscriber in order. You want the unit tests to be cost-effective and reliable. What should you do?

- A. Implement a mocking framework.
- B. Create a topic and subscription for each tester.
- C. Add a filter by tester to the subscription.
- D. Use the Pub/Sub emulator.

**Answer: D**

**Explanation:**

<https://cloud.google.com/pubsub/docs/emulator>, "Testing apps locally with the emulator".

**NEW QUESTION 16**

- (Exam Topic 2)

You want to create "fully baked" or "golden" Compute Engine images for your application. You need to bootstrap your application to connect to the appropriate database according to the environment the application is running on (test, staging, production). What should you do?

- A. Embed the appropriate database connection string in the image.
- B. Create a different image for each environment.
- C. When creating the Compute Engine instance, add a tag with the name of the database to be connected. In your application, query the Compute Engine API to pull the tags for the current instance, and use the tag to construct the appropriate database connection string.
- D. When creating the Compute Engine instance, create a metadata item with a key of "DATABASE" and a value for the appropriate database connection string.
- E. In your application, read the "DATABASE" environment variable, and use the value to connect to the appropriate database.
- F. When creating the Compute Engine instance, create a metadata item with a key of "DATABASE" and a value for the appropriate database connection string.
- G. In your application, query the metadata server for the "DATABASE" value, and use the value to connect to the appropriate database.

**Answer: C**

**NEW QUESTION 17**

- (Exam Topic 2)

You are developing a marquee stateless web application that will run on Google Cloud. The rate of the incoming user traffic is expected to be unpredictable, with no traffic on some days and large spikes on other days. You need the application to automatically scale up and down, and you need to minimize the cost associated with running the application. What should you do?

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

**Answer: A**

**NEW QUESTION 20**

- (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 22

- (Exam Topic 2)

You migrated your applications to Google Cloud Platform and kept your existing monitoring platform. You now find that your notification system is too slow for time critical problems. What should you do?

- A. Replace your entire monitoring platform with Stackdriver.
- B. Install the Stackdriver agents on your Compute Engine instances.
- C. Use Stackdriver to capture and alert on logs, then ship them to your existing platform.
- D. Migrate some traffic back to your old platform and perform AB testing on the two platforms concurrently.

**Answer:** B

**Explanation:**

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

#### NEW QUESTION 27

- (Exam Topic 2)

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

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

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 30

- (Exam Topic 2)

You work for a web development team at a small startup. Your team is developing a Node.js application using Google Cloud services, including Cloud Storage and Cloud Build. The team uses a Git repository for version control. Your manager calls you over the weekend and instructs you to make an emergency update to one of the company's websites, and you're the only developer available. You need to access Google Cloud to make the update, but you don't have your work laptop. You are not allowed to store source code locally on a non-corporate computer. How should you set up your developer environment?

- A. Use a text editor and the Git command line to send your source code updates as pull requests from a public computer.
- B. Use a text editor and the Git command line to send your source code updates as pull requests from a virtual machine running on a public computer.
- C. Use Cloud Shell and the built-in code editor for developmen
- D. Send your source code updates as pull requests.
- E. Use a Cloud Storage bucket to store the source code that you need to edi
- F. Mount the bucket to a public computer as a drive, and use a code editor to update the cod
- G. Turn on versioning for the bucket, and point it to the team's Git repository.

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 31

- (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 cluste
- 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 cluste
- 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 cluste
- G. Using Kubernetes role-based access control (RBAC), grant the Admin role on the new namespace to the other team.

**Answer:** D

#### NEW QUESTION 32

- (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 35**

- (Exam Topic 2)

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

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

**Answer:** C

**Explanation:**

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

The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload's CPU or memory consumption, or in response to custom metrics reported from within Kubernetes or external metrics from sources outside of your cluster.

**NEW QUESTION 37**

- (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 40**

- (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 42**

- (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 46**

- (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 databas
- 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 availabilit
- 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 availabilit
- 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 49**

- (Exam Topic 2)

Your team is building an application for a financial institution. The application's frontend runs on Compute Engine, and the data resides in Cloud SQL and one Cloud Storage bucket. The application will collect data containing PII, which will be stored in the Cloud SQL database and the Cloud Storage bucket. You need to secure the PII data. What should you do?

- A. 1) Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database2) Using IAM, allow only the frontend service account to access the Cloud Storage bucket
- B. 1) Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database2) Enable private access to allow the frontend to access the Cloud Storage bucket privately
- C. 1) Configure a private IP address for Cloud SQL2) Use VPC-SC to create a service perimeter3) Add the Cloud SQL database and the Cloud Storage bucket to the same service perimeter
- D. 1) Configure a private IP address for Cloud SQL2) Use VPC-SC to create a service perimeter3) Add the Cloud SQL database and the Cloud Storage bucket to different service perimeters

**Answer: C**

**NEW QUESTION 51**

- (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 acces
- 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](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 52**

- (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 bythe 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 57**

- (Exam Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 62**

- (Exam Topic 2)

Your development team is using Cloud Build to promote a Node.js application built on App Engine from your staging environment to production. The application relies on several directories of photos stored in a Cloud Storage bucket named webphotos-staging in the staging environment. After the promotion, these photos must be available in a Cloud Storage bucket named webphotos-prod in the production environment. You want to automate the process where possible. What should you do?

- A) Manually copy the photos to webphotos-prod.
- B) Add a startup script in the application's app.yaml file to move the photos from webphotos-staging to webphotos-prod.
- C) Add a build step in the cloudbuild.yaml file before the promotion step with the arguments:  

```
- name: gcr.io/cloud-builders/gsutil
  args: ['cp', '-r', 'gs://webphotos-staging',
        'gs://webphotos-prod']
  waitFor: ['-']
```
- D) Add a build step in the cloudbuild.yaml file before the promotion step with the arguments:

```
- name: gcr.io/cloud-builders/gcloud
  args: ['cp', '-A', 'gs://webphotos-staging',
        'gs://webphotos-prod']
  waitFor: ['-']
```

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 65**

- (Exam Topic 2)

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 70**

- (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 75

- (Exam Topic 2)

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

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

**Answer: A**

**Explanation:**

[https://cloud.google.com/iap/docs/signed-headers-howto#securing\\_iap\\_headers](https://cloud.google.com/iap/docs/signed-headers-howto#securing_iap_headers) (<https://cloud.google.com/endpoints/docs/openapi/authenticating-users-google-id>).  
<https://cloud.google.com/armor/docs/security-policy-overview#:~:text=Google%20Cloud%20Armor%20security%20policies,protect,your,application,by,providing,Layer,7,filtering,and,by,scrubbing,incoming,requests,for,common,web,attacks,or,other,Layer,7,attributes,to,potentially,block,traffic,before,it,reaches,your,load,balanced,backend,services,or,backend,buckets>" Google Cloud Armor security policies protect your application by providing Layer 7 filtering and by scrubbing incoming requests for common web attacks or other Layer 7 attributes to potentially block traffic before it reaches your load balanced backend services or backend buckets"

#### NEW QUESTION 78

- (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 80

- (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 83

- (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 85

- (Exam Topic 2)

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

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

**Answer:** A

#### Explanation:

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

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

### NEW QUESTION 86

- (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 88

- (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 92

- (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 acces
- 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/compute/docs/instances/connecting-advanced#cloud_iap) [https://cloud.google.com/solutions/connecting-securely#storing\\_host\\_keys\\_by\\_enabling\\_guest\\_attributes](https://cloud.google.com/solutions/connecting-securely#storing_host_keys_by_enabling_guest_attributes)

**NEW QUESTION 96**

- (Exam Topic 2)

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

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

**Answer:** D

**Explanation:**

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

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

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

**NEW QUESTION 98**

- (Exam Topic 2)

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

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

**Answer:** D

**NEW QUESTION 102**

- (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 is 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 103**

- (Exam Topic 2)

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

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

**Answer: D**

#### NEW QUESTION 105

- (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 108

- (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 112

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

#### NEW QUESTION 114

- (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 119

- (Exam Topic 2)

You manage a microservices application on Google Kubernetes Engine (GKE) using Istio. You secure the communication channels between your microservices by implementing an Istio AuthorizationPolicy, a Kubernetes NetworkPolicy, and mTLS on your GKE cluster. You discover that HTTP requests between two Pods to specific URLs fail, while other requests to other URLs succeed. What is the cause of the connection issue?

- A. A Kubernetes NetworkPolicy resource is blocking HTTP traffic between the Pods.
- B. The Pod initiating the HTTP requests is attempting to connect to the target Pod via an incorrect TCP port.
- C. The Authorization Policy of your cluster is blocking HTTP requests for specific paths within your application.
- D. The cluster has mTLS configured in permissive mode, but the Pod's sidecar proxy is sending unencrypted traffic in plain text.

**Answer: C**

#### NEW QUESTION 120

- (Exam Topic 2)

You are planning to deploy hundreds of microservices in your Google Kubernetes Engine (GKE) cluster. How should you secure communication between the microservices on GKE using a managed service?

- A. Use global HTTP(S) Load Balancing with managed SSL certificates to protect your services
- B. Deploy open source Istio in your GKE cluster, and enable mTLS in your Service Mesh
- C. Install cert-manager on GKE to automatically renew the SSL certificates.
- D. Install Anthos Service Mesh, and enable mTLS in your Service Mesh.

**Answer: D**

#### Explanation:

[https://cloud.google.com/service-mesh/docs/overview#security\\_benefits](https://cloud.google.com/service-mesh/docs/overview#security_benefits)

- Ensures encryption in transit. Using mTLS for authentication also ensures that all TCP communications are encrypted in transit.

#### NEW QUESTION 123

- (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 124

- (Exam Topic 2)

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

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

**Answer: C**

#### NEW QUESTION 127

- (Exam Topic 2)

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

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

**Answer:** B

**Explanation:**

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

#### NEW QUESTION 129

- (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 133

- (Exam Topic 2)

The development teams in your company want to manage resources from their local environments. You have been asked to enable developer access to each team's Google Cloud projects. You want to maximize efficiency while following Google-recommended best practices. What should you do?

- A. Add the users to their projects, assign the relevant roles to the users, and then provide the users with each relevant Project ID.
- B. Add the users to their projects, assign the relevant roles to the users, and then provide the users with each relevant Project Number.
- C. Create groups, add the users to their groups, assign the relevant roles to the groups, and then provide the users with each relevant Project ID.
- D. Create groups, add the users to their groups, assign the relevant roles to the groups, and then provide the users with each relevant Project Number.

**Answer:** C

#### NEW QUESTION 135

- (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 139

- (Exam Topic 2)

Your application is controlled by a managed instance group. You want to share a large read-only data set between all the instances in the managed instance group. You want to ensure that each instance can start quickly and can access the data set via its filesystem with very low latency. You also want to minimize the Total cost of the solution. What should you do?

- A. Move the data to a Cloud Storage bucket, and mount the bucket on the filesystem using Cloud Storage FUSE.
- B. Move the data to a Cloud Storage bucket, and copy the data to the boot disk of the instance via a startup script.
- C. Move the data to a Compute Engine persistent disk, and attach the disk in read-only mode to multiple Compute Engine virtual machine instances.
- D. Move the data to a Compute Engine persistent disk, take a snapshot, create multiple disks from the snapshot, and attach each disk to its own instance.

**Answer:** C

#### NEW QUESTION 144

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

**NEW QUESTION 145**

- (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 148**

- (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 153**

- (Exam Topic 2)

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

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

**Answer: D**

**Explanation:**

Reference:

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

**NEW QUESTION 158**

- (Exam Topic 2)

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

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 160**

- (Exam Topic 2)

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

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

**Answer:** D

**Explanation:**

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

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

**NEW QUESTION 165**

- (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 170**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

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

2nd - Questions and Answers in PDF Format

### **Professional-Cloud-Developer Practice Exam Features:**

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

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