

Google.Professional Cloud Developer.vJun-2024.by.Lioane.156q

Number: Professional Cloud Developer  
Passing Score: 800  
Time Limit: 120  
File Version: 32.0

**Certification: Professional Cloud Developer**  
**Certification Full Name: Professional Cloud Developer**



## Exam A

### QUESTION 1

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.

**Correct Answer: C, E**

**Section:**

### QUESTION 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`

**Correct Answer: C**

**Section:**

### QUESTION 3

Your analytics system executes queries against a BigQuery dataset. The SQL query is executed in batch and passes the contents of a SQL file to the BigQuery CLI. Then it redirects the BigQuery CLI output to another process. However, you are getting a permission error from the BigQuery CLI when the queries are executed. You want to resolve the issue. What should you do?

- A. Grant the service account BigQuery Data Viewer and BigQuery Job User roles.
- B. Grant the service account BigQuery Data Editor and BigQuery Data Viewer roles.
- C. Create a view in BigQuery from the SQL query and SELECT\* from the view in the CLI.
- D. Create a new dataset in BigQuery, and copy the source table to the new dataset Query the new dataset and table from the CLI.

**Correct Answer: B**

**Section:**

### QUESTION 4



Your application is running on Compute Engine and is showing sustained failures for a small number of requests. You have narrowed the cause down to a single Compute Engine instance, but the instance is unresponsive to SSH. What should you do next?

- A. Reboot the machine.
- B. Enable and check the serial port output.
- C. Delete the machine and create a new one.
- D. Take a snapshot of the disk and attach it to a new machine.

**Correct Answer: A**

**Section:**

#### QUESTION 5

You configured your Compute Engine instance group to scale automatically according to overall CPU usage. However, your application's response latency increases sharply before the cluster has finished adding up instances. You want to provide a more consistent latency experience for your end users by changing the configuration of the instance group autoscaler. Which two configuration changes should you make? (Choose two.)

- A. Add the label "AUTOSCALE" to the instance group template.
- B. Decrease the cool-down period for instances added to the group.
- C. Increase the target CPU usage for the instance group autoscaler.
- D. Decrease the target CPU usage for the instance group autoscaler.
- E. Remove the health-check for individual VMs in the instance group.

**Correct Answer: A, C**

**Section:**

#### QUESTION 6

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.

**Correct Answer: A**

**Section:**

#### QUESTION 7

Your application requires service accounts to be authenticated to GCP products via credentials stored on its host Compute Engine virtual machine instances. You want to distribute these credentials to the host instances as securely as possible. What should you do?

- A. Use HTTP signed URLs to securely provide access to the required resources.
- B. Use the instance's service account Application Default Credentials to authenticate to the required resources.
- C. Generate a P12 file from the GCP Console after the instance is deployed, and copy the credentials to the host instance before starting the application.
- D. Commit the credential JSON file into your application's source repository, and have your CI/CD process package it with the software that is deployed to the instance.

**Correct Answer: B**

**Section:**



#### QUESTION 8

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

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

**Correct Answer: A**

**Section:**

#### QUESTION 9

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

What should you do?

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

**Correct Answer: A**

**Section:**



#### QUESTION 10

Your team develops services that run on Google Cloud. You need to build a data processing service and will use Cloud Functions. The data to be processed by the function is sensitive. You need to ensure that invocations can only happen from authorized services and follow Google-recommended best practices for securing functions. What should you do?

- A. Enable Identity-Aware Proxy in your project. Secure function access using its permissions.
- B. Create a service account with the Cloud Functions Viewer role. Use that service account to invoke the function.
- C. Create a service account with the Cloud Functions Invoker role. Use that service account to invoke the function.
- D. Create an OAuth 2.0 client ID for your calling service in the same project as the function you want to secure. Use those credentials to invoke the function.

**Correct Answer: C**

**Section:**

#### QUESTION 11

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.

**Correct Answer: A, D**

**Section:**

**Explanation:**

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

**QUESTION 12**

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. Change the web application DNS to point to the proxy Compute Engine instance. After authenticating, the Compute Engine instance forwards requests to and from the web application.
- D. Configure a Compute Engine instance that requests users to log in to their corporate account. Change the web application DNS to point to the proxy Compute Engine instance. After authenticating, the Compute Engine issues an HTTP redirect to a public IP address hosting the web application.

**Correct Answer: B**

**Section:**

**QUESTION 13**

You have an application that uses an HTTP Cloud Function to process user activity from both desktop browser and mobile application clients. This function will serve as the endpoint for all metric submissions using HTTP POST. Due to legacy restrictions, the function must be mapped to a domain that is separate from the domain requested by users on web or mobile sessions. The domain for the Cloud Function is <https://fn.example.com>. Desktop and mobile clients use the domain <https://www.example.com>. You need to add a header to the function's HTTP response so that only those browser and mobile sessions can submit metrics to the Cloud Function. Which response header should you add?

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



**Correct Answer: D**

**Section:**

**QUESTION 14**

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.

**Correct Answer: D**

**Section:**

**QUESTION 15**

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.

**Correct Answer: D**

**Section:**

#### QUESTION 16

You migrated some of your applications to Google Cloud. You are using a legacy monitoring platform deployed on-premises for both on-premises and cloud-deployed applications. You discover that your notification system is responding slowly to time-critical problems in the cloud applications. What should you do?

- A. Replace your monitoring platform with Cloud Monitoring.
- B. Install the Cloud Monitoring agent on your Compute Engine instances.
- C. Migrate some traffic back to your old platform. Perform A/B testing on the two platforms concurrently.
- D. Use Cloud Logging and Cloud Monitoring to capture logs, monitor, and send alerts. Send them to your existing platform.

**Correct Answer: D**

**Section:**

#### QUESTION 17

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

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

**Correct Answer: C**

**Section:**

#### QUESTION 18

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.

**Correct Answer: B**

**Section:**

#### QUESTION 19

You are planning to migrate a MySQL database to the managed Cloud SQL database for Google Cloud. You have Compute Engine virtual machine instances that will connect with this Cloud SQL instance. You do not want to whitelist IPs for the Compute Engine instances to be able to access Cloud SQL.

What should you do?

- A. Enable private IP for the Cloud SQL instance.
- B. Whitelist a project to access Cloud SQL, and add Compute Engine instances in the whitelisted project.
- C. Create a role in Cloud SQL that allows access to the database from external instances, and assign the Compute Engine instances to that role.
- D. Create a CloudSQL instance on one project. Create Compute engine instances in a different project. Create a VPN between these two projects to allow internal access to CloudSQL.

**Correct Answer: C**

**Section:**

#### QUESTION 20

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

**Correct Answer: C**

**Section:**

### QUESTION 21

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

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

**Correct Answer: A**

**Section:**

### QUESTION 22

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

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

**Correct Answer: C**

**Section:**

### QUESTION 23

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.

**Correct Answer: B**

**Section:**

### QUESTION 24

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.

**Correct Answer: A**

**Section:**

#### QUESTION 25

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

Which data structure should you use?

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

**Correct Answer: B**

**Section:**

#### QUESTION 26

Your company has a BigQuery dataset named 'Master' that keeps information about employee travel and expenses. This information is organized by employee department. That means employees should only be able to view information for their department. You want to apply a security framework to enforce this requirement with the minimum number of steps.

What should you do?

- A. Create a separate dataset for each department. Create a view with an appropriate WHERE clause to select records from a particular dataset for the specific department. Authorize this view to access records from your Master dataset. Give employees the permission to this department-specific dataset.
- B. Create a separate dataset for each department. Create a data pipeline for each department to copy appropriate information from the Master dataset to the specific dataset for the department. Give employees the permission to this department-specific dataset.
- C. Create a dataset named Master dataset. Create a separate view for each department in the Master dataset. Give employees access to the specific view for their department.
- D. Create a dataset named Master dataset. Create a separate table for each department in the Master dataset. Give employees access to the specific table for their department.

**Correct Answer: B**

**Section:**

#### QUESTION 27

You have an application in production. It is deployed on Compute Engine virtual machine instances controlled by a managed instance group. Traffic is routed to the instances via a HTTP(s) load balancer. Your users are unable to access your application. You want to implement a monitoring technique to alert you when the application is unavailable.

Which technique should you choose?

- A. Smoke tests
- B. Stackdriver uptime checks
- C. Cloud Load Balancing - health checks
- D. Managed instance group - health checks

**Correct Answer: B**

**Section:**

**Explanation:**  
476b8507f59c

**QUESTION 28**

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.

**Correct Answer: A**  
**Section:**

**QUESTION 29**

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.

**Correct Answer: C**  
**Section:**

**QUESTION 30**

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. Clients should use the name of the A record to connect to the service.
- C. 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/`.
- D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url `https://[API_NAME]/[API_VERSION]/`.

**Correct Answer: D**  
**Section:**

**QUESTION 31**

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

- A. Add a Stackdriver counter metric for path:`/api/alpha/`.

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

**Correct Answer: C**

**Section:**

#### QUESTION 32

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.

**Correct Answer: C**

**Section:**

#### QUESTION 33

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

**Correct Answer: A**

**Section:**

#### QUESTION 34

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

**Correct Answer: B**

**Section:**

#### QUESTION 35

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.

**Correct Answer: D**

**Section:**

**QUESTION 36**

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.

**Correct Answer: C**

**Section:**

**QUESTION 37**

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.

**Correct Answer: A**

**Section:**

**QUESTION 38**

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

**Correct Answer: D**

**Section:**

**QUESTION 39**

HipLocal wants to reduce the number of on-call engineers and eliminate manual scaling. Which two services should they choose? (Choose two.)



- A. Use Google App Engine services.
- B. Use serverless Google Cloud Functions.
- C. Use Knative to build and deploy serverless applications.
- D. Use Google Kubernetes Engine for automated deployments.
- E. Use a large Google Compute Engine cluster for deployments.

**Correct Answer: B, C**

**Section:**

#### QUESTION 40

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

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

**Correct Answer: B**

**Section:**

#### QUESTION 41

Which service should HipLocal use for their public APIs?

- A. Cloud Armor
- B. Cloud Functions
- C. Cloud Endpoints
- D. Shielded Virtual Machines

**Correct Answer: D**

**Section:**

#### QUESTION 42

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.

**Correct Answer: B**

**Section:**

#### QUESTION 43

Which database should HipLocal use for storing user activity?



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

**Correct Answer: A**

**Section:**

**QUESTION 44**

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.

**Correct Answer: D**

**Section:**

**QUESTION 45**

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.

**Correct Answer: B**

**Section:**

**Explanation:**

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

**QUESTION 46**

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

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

**Correct Answer: D**

**Section:**

**QUESTION 47**

For this question, refer to the HipLocal case study.



How should HipLocal redesign their architecture to ensure that the application scales to support a large increase in users?

- A. Use Google Kubernetes Engine (GKE) to run the application as a microservice. Run the MySQL database on a dedicated GKE node.
- B. Use multiple Compute Engine instances to run MySQL to store state information. Use a Google Cloud-managed load balancer to distribute the load between instances. Use managed instance groups for scaling.
- C. Use Memorystore to store session information and CloudSQL to store state information. Use a Google Cloud-managed load balancer to distribute the load between instances. Use managed instance groups for scaling.
- D. Use a Cloud Storage bucket to serve the application as a static website, and use another Cloud Storage bucket to store user state information.

**Correct Answer: D**

**Section:**

#### QUESTION 48

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 Splitting. Do not move QA traffic to the new versions if errors are found.

**Correct Answer: B**

**Section:**

#### QUESTION 49

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. Use the API key to interact with Google Cloud.
- B. Use the default compute service account to interact with Google Cloud.
- C. Create a service account for the application. Export and deploy the private key for the application. Use the service account to interact with Google Cloud.
- D. 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. Use the service account with one Google Cloud API to interact with Google Cloud.

**Correct Answer: A**

**Section:**

#### QUESTION 50

For this question refer to the HipLocal case study.

HipLocal wants to reduce the latency of their services for users in global locations. They have created read replicas of their database in locations where their users reside and configured their service to read traffic using those replicas. How should they further reduce latency for all database interactions with the least amount of effort?

- A. Migrate the database to Bigtable and use it to serve all global user traffic.
- B. Migrate the database to Cloud Spanner and use it to serve all global user traffic.
- C. Migrate the database to Firestore in Datastore mode and use it to serve all global user traffic.
- D. Migrate the services to Google Kubernetes Engine and use a load balancer service to better scale the application.

**Correct Answer: D**

**Section:**

**QUESTION 51**

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

**Correct Answer: B**

**Section:**

**Explanation:**

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

**QUESTION 52**

For this question, refer to the HipLocal case study.

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

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

**Correct Answer: D**

**Section:**

**Explanation:**

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

**QUESTION 53**

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

**Correct Answer: B**

**Section:**

**QUESTION 54**

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]/

**Correct Answer: A**

**Section:**

**Explanation:**

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

#### QUESTION 55

You are using Cloud Build to create a new Docker image on each source code commit to a Cloud Source Repositories repository. Your application is built on every commit to the master branch. You want to release specific commits made to the master branch in an automated method. What should you do?

- A. Manually trigger the build for new releases.
- B. Create a build trigger on a Git tag pattern. Use a Git tag convention for new releases.
- C. Create a build trigger on a Git branch name pattern. Use a Git branch naming convention for new releases.
- D. Commit your source code to a second Cloud Source Repositories repository with a second Cloud Build trigger. Use this repository for new releases only.

**Correct Answer: C**

**Section:**

#### QUESTION 56

You are designing a schema for a table that will be moved from MySQL to Cloud Bigtable. The MySQL table is as follows:

```
AccountActivity
(
  Account_id int,
  Event_timestamp datetime,
  Transaction_type string,
  Amount numeric(18, 4)
) primary key (Account_id, Event_timestamp)
```

How should you design a row key for Cloud Bigtable for this table?

- A. Set Account\_id as a key.
- B. Set Account\_id\_Event\_timestamp as a key.
- C. Set Event\_timestamp\_Account\_id as a key.
- D. Set Event\_timestamp as a key.

**Correct Answer: C**

**Section:**

#### QUESTION 57

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.

**Correct Answer: C**

**Section:**

**QUESTION 58**

You have an analytics application that runs hundreds of queries on BigQuery every few minutes using BigQuery API. You want to find out how much time these queries take to execute. What should you do?

- A. Use Stackdriver Monitoring to plot slot usage.
- B. Use Stackdriver Trace to plot API execution time.
- C. Use Stackdriver Trace to plot query execution time.
- D. Use Stackdriver Monitoring to plot query execution times.

**Correct Answer: D**

**Section:**

**QUESTION 59**

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

**Correct Answer: C**

**Section:**

**QUESTION 60**

You are deploying a single website on App Engine that needs to be accessible via the URL <http://www.altostrat.com/>. What should you do?

- A. Verify domain ownership with Webmaster Central. Create a DNS CNAME record to point to the App Engine canonical name ghs.googlehosted.com.
- B. Verify domain ownership with Webmaster Central. Define an A record pointing to the single global App Engine IP address.
- C. Define a mapping in dispatch.yaml to point the domain www.altostrat.com to your App Engine service. Create a DNS CNAME record to point to the App Engine canonical name ghs.googlehosted.com.
- D. Define a mapping in dispatch.yaml to point the domain www.altostrat.com to your App Engine service. Define an A record pointing to the single global App Engine IP address.

**Correct Answer: A**

**Section:**

**QUESTION 61**

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

**Correct Answer: C**

**Section:**

**QUESTION 62**

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

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

**Correct Answer: B**

**Section:**

**QUESTION 63**

Your application is built as a custom machine image. You have multiple unique deployments of the machine image. Each deployment is a separate managed instance group with its own template. Each deployment requires a unique set of configuration values. You want to provide these unique values to each deployment but use the same custom machine image in all deployments. You want to use out-of-the-box features of Compute Engine. What should you do?

- A. Place the unique configuration values in the persistent disk.
- B. Place the unique configuration values in a Cloud Bigtable table.
- C. Place the unique configuration values in the instance template startup script.
- D. Place the unique configuration values in the instance template instance metadata.



**Correct Answer: A**

**Section:**

**QUESTION 64**

Your application performs well when tested locally, but it runs significantly slower when you deploy it to App Engine standard environment. You want to diagnose the problem. What should you do?

- A. File a ticket with Cloud Support indicating that the application performs faster locally.
- B. Use Stackdriver Debugger Snapshots to look at a point-in-time execution of the application.
- C. Use Stackdriver Trace to determine which functions within the application have higher latency.
- D. Add logging commands to the application and use Stackdriver Logging to check where the latency problem occurs.

**Correct Answer: D**

**Section:**

**QUESTION 65**

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.

**Correct Answer: C**

**Section:**

### QUESTION 66

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

**Correct Answer: D**

**Section:**

### QUESTION 67

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`

**Correct Answer: B**

**Section:**

### QUESTION 68

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.

**Correct Answer: D**

**Section:**

#### QUESTION 69

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. In the Cloud Function, publish a message to the Cloud Pub/Sub topic.
- D. 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.

**Correct Answer: C**

**Section:**

**Explanation:**

: <https://cloud.google.com/storage/docs/pubsub-notifications>

#### QUESTION 70

Your teammate has asked you to review the code below, which is adding a credit to an account balance in Cloud Datastore. Which improvement should you suggest your teammate make?

```
public Entity creditAccount(long accountId, long
creditAmount) {
    Entity account = datastore.get
(keyFactory.newKey(accountId));
    account = Entity.builder(account).set(
        "balance", account.getLong("balance")
+ creditAmount).build();
    datastore.put(account);
    return account;
}
```

- A. Get the entity with an ancestor query.
- B. Get and put the entity in a transaction.
- C. Use a strongly consistent transactional database.
- D. Don't return the account entity from the function.

**Correct Answer: A**

**Section:**

#### QUESTION 71

Your company stores their source code in a Cloud Source Repositories repository. Your company wants to build and test their code on each source code commit to the repository and requires a solution that is managed and has minimal operations overhead.

Which method should they use?

- A. Use Cloud Build with a trigger configured for each source code commit.
- B. Use Jenkins deployed via the Google Cloud Platform Marketplace, configured to watch for source code commits.
- C. Use a Compute Engine virtual machine instance with an open source continuous integration tool, configured to watch for source code commits.
- D. Use a source code commit trigger to push a message to a Cloud Pub/Sub topic that triggers an App Engine service to build the source code.

**Correct Answer: A**

**Section:**

**Explanation:**

<https://cloud.google.com/build/docs/automating-builds/create-manage-triggers#:~:text=A%20Cloud%20Build%20trigger%20automatically,changes%20that%20match%20certain%20criteria>.

#### QUESTION 72

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

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

**Correct Answer: B**

**Section:**

**Explanation:**

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

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

#### QUESTION 73

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

**Correct Answer: A**

**Section:**

**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%20enable%20you%20to%20allow%20or,Private%20Cloud%20\(VPC\)%20networks](https://cloud.google.com/armor/docs/security-policy-overview#:~:text=Google%20Cloud%20Armor%20security%20policies%20enable%20you%20to%20allow%20or,Private%20Cloud%20(VPC)%20networks)

'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'

#### QUESTION 74

Your API backend is running on multiple cloud providers. You want to generate reports for the network latency of your API. Which two steps should you take? (Choose two.)

- A. Use Zipkin collector to gather data.
- B. Use Fluentd agent to gather data.
- C. Use Stackdriver Trace to generate reports.
- D. Use Stackdriver Debugger to generate report.
- E. Use Stackdriver Profiler to generate report.

**Correct Answer: A, C**

**Section:**

**Explanation:**

<https://cloud.google.com/trace/docs/zipkin>

'receive traces from Zipkin clients and forward those traces to Cloud Trace for analysis.'

[https://cloud.google.com/trace/docs/quickstart#analysis\\_reports\\_window](https://cloud.google.com/trace/docs/quickstart#analysis_reports_window)

#### QUESTION 75

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

**Correct Answer: A**

**Section:**

#### QUESTION 76

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

**Correct Answer: A**

**Section:**

**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.'

#### QUESTION 77

You are developing an application that will be launched on Compute Engine instances into multiple distinct projects, each corresponding to the environments in your software development process (development, QA, staging, and production). The instances in each project have the same application code but a different configuration. During deployment, each instance should receive the application's configuration based on the environment it serves. You want to minimize the number of steps to configure this flow.



What should you do?

- A. When creating your instances, configure a startup script using the gcloud command to determine the project name that indicates the correct environment.
- B. In each project, configure a metadata key "environment" whose value is the environment it serves. Use your deployment tool to query the instance metadata and configure the application based on the "environment" value.
- C. Deploy your chosen deployment tool on an instance in each project. Use a deployment job to retrieve the appropriate configuration file from your version control system, and apply the configuration when deploying the application on each instance.
- D. During each instance launch, configure an instance custom-metadata key named "environment" whose value is the environment the instance serves. Use your deployment tool to query the instance metadata, and configure the application based on the "environment" value.

**Correct Answer: B**

**Section:**

#### QUESTION 78

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.

**Correct Answer: C**

**Section:**

#### QUESTION 79

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

**Correct Answer: D**

**Section:**

#### QUESTION 80

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 minutes. After the file is read, the application disconnects from the database.
- D. Write a shutdown script that publishes a message to the Pub/Sub topic announcing that a shutdown is in progress. After the application reads the message, it disconnects from the database.

**Correct Answer: D**





**Section:**

**QUESTION 81**

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 development. Send your source code updates as pull requests.
- D. Use a Cloud Storage bucket to store the source code that you need to edit. Mount the bucket to a public computer as a drive, and use a code editor to update the code. Turn on versioning for the bucket, and point it to the team's Git repository.

**Correct Answer: C**

**Section:**

**Explanation:**

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

**QUESTION 82**

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.

**Correct Answer: A, C**

**Section:**

**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)

**QUESTION 83**

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

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

**Correct Answer: A**

**Section:**

**QUESTION 84**

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.

**Correct Answer: C**

**Section:**

**Explanation:**

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

#### QUESTION 85

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. Create one pull subscription to allow the audit services to share the messages.
- B. Create one Pub/Sub topic. Create one pull subscription per audit service instance to allow the services to share the messages.
- C. Create one Pub/Sub topic. Create one push subscription with the endpoint pointing to a load balancer in front of the audit services.
- D. Create one Pub/Sub topic per authentication service. Create one pull subscription per topic to be used by one audit service.
- E. Create one Pub/Sub topic per authentication service. Create one push subscription per topic, with the endpoint pointing to one audit service.

**Correct Answer: A**

**Section:**

**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.'

#### QUESTION 86

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

**Correct Answer: A**

**Section:**

#### QUESTION 87

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 resources. The deployer is given permission to get the access token.
- B. Create a new service account that has a custom IAM role to access the resources. The deployer is given permission to get the access token.

- C. Create a new service account that has Editor authority to access the resources. The deployer is given permission to act as the new service account.
- D. Create a new service account that has a custom IAM role to access the resources. The deployer is given permission to act as the new service account.

**Correct Answer: D**

**Section:**

#### QUESTION 88

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.

**Correct Answer: B**

**Section:**

#### QUESTION 89

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

A)

Run `gcloud builds submit` in the directory that contains the application source code.

B)

Run `gcloud run deploy app-name --image gcr.io/$PROJECT_ID/app-name` in the directory that contains the application source code.

C)

Run `gcloud container images add-tag gcr.io/$PROJECT_ID/app-name gcr.io/$PROJECT_ID/app-name:latest` in the directory that contains the application source code.

D)

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

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

E)

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

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

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

**Correct Answer: A, D**

**Section:**

**Explanation:**

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

#### QUESTION 90

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

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

**Correct Answer: B**

**Section:**

**Explanation:**

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

#### QUESTION 91

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

**Correct Answer: C**

**Section:**

**Explanation:**

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



#### QUESTION 92

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

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

**Correct Answer: C**

**Section:**

**Explanation:**

[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)

#### QUESTION 93

You have a mixture of packaged and internally developed applications hosted on a Compute Engine instance that is running Linux. These applications write log records as text in local files. You want the logs to be written to Cloud Logging. What should you do?

- A. Pipe the content of the files to the Linux Syslog daemon.
- B. Install a Google version of fluentd on the Compute Engine instance.
- C. Install a Google version of collectd on the Compute Engine instance.
- D. Using cron, schedule a job to copy the log files to Cloud Storage once a day.

**Correct Answer: B**

**Section:**



#### QUESTION 94

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. Create a different image for each environment.
- B. 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.
- C. When creating the Compute Engine instance, create a metadata item with a key of "DATABASE" and a value for the appropriate database connection string. In your application, read the "DATABASE" environment variable, and use the value to connect to the appropriate database.
- 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. In your application, query the metadata server for the "DATABASE" value, and use the value to connect to the appropriate database.

**Correct Answer: D**

**Section:**

#### QUESTION 95

You are developing a microservice-based application that will be deployed on a Google Kubernetes Engine cluster. The application needs to read and write to a Spanner database. You want to follow security best practices while minimizing code changes. How should you configure your application to retrieve Spanner credentials?

- A. Configure the appropriate service accounts, and use Workload Identity to run the pods.

- B. Store the application credentials as Kubernetes Secrets, and expose them as environment variables.
- C. Configure the appropriate routing rules, and use a VPC-native cluster to directly connect to the database.
- D. Store the application credentials using Cloud Key Management Service, and retrieve them whenever a database connection is made.

**Correct Answer: A**

**Section:**

**Explanation:**

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

#### QUESTION 96

You are deploying your application on a Compute Engine instance that communicates with Cloud SQL. You will use Cloud SQL Proxy to allow your application to communicate to the database using the service account associated with the application's instance. You want to follow the Google-recommended best practice of providing minimum access for the role assigned to the service account. What should you do?

- A. Assign the Project Editor role.
- B. Assign the Project Owner role.
- C. Assign the Cloud SQL Client role.
- D. Assign the Cloud SQL Editor role.

**Correct Answer: C**

**Section:**

#### QUESTION 97

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.

**Correct Answer: C**

**Section:**

**Explanation:**

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

#### QUESTION 98

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

**Correct Answer: D**

**Section:**

**QUESTION 99**

Your company just experienced a Google Kubernetes Engine (GKE) API outage due to a zone failure. You want to deploy a highly available GKE architecture that minimizes service interruption to users in the event of a future zone failure. What should you do?

- A. Deploy Zonal clusters
- B. Deploy Regional clusters
- C. Deploy Multi-Zone clusters
- D. Deploy GKE on-premises clusters

**Correct Answer: B**

**Section:**

**Explanation:**

[https://cloud.google.com/kubernetes-engine/docs/concepts/types-of-clusters#regional\\_clusters](https://cloud.google.com/kubernetes-engine/docs/concepts/types-of-clusters#regional_clusters)

A regional cluster has multiple replicas of the control plane, running in multiple zones within a given region. Nodes in a regional cluster can run in multiple zones or a single zone depending on the configured node locations. By default, GKE replicates each node pool across three zones of the control plane's region. When you create a cluster or when you add a new node pool, you can change the default configuration by specifying the zone(s) in which the cluster's nodes run. All zones must be within the same region as the control plane.

**QUESTION 100**

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.

**Correct Answer: B**

**Section:**

**Explanation:**

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

**QUESTION 101**

You are developing a microservice-based application that will run on Google Kubernetes Engine (GKE). Some of the services need to access different Google Cloud APIs. How should you set up authentication of these services in the cluster following Google-recommended best practices? (Choose two.)

- A. Use the service account attached to the GKE node.
- B. Enable Workload Identity in the cluster via the gcloud command-line tool.
- C. Access the Google service account keys from a secret management service.
- D. Store the Google service account keys in a central secret management service.
- E. Use gcloud to bind the Kubernetes service account and the Google service account using roles/iam.workloadIdentity.

**Correct Answer: B, E**

**Section:**

**Explanation:**

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

**QUESTION 102**

Your development team has been tasked with maintaining a .NET legacy application. The application incurs occasional changes and was recently updated. Your goal is to ensure that the application provides consistent results

while moving through the CI/CD pipeline from environment to environment. You want to minimize the cost of deployment while making sure that external factors and dependencies between hosting environments are not problematic. Containers are not yet approved in your organization. What should you do?

- A. Rewrite the application using .NET Core, and deploy to Cloud Run. Use revisions to separate the environments.
- B. Use Cloud Build to deploy the application as a new Compute Engine image for each build. Use this image in each environment.
- C. Deploy the application using MS Web Deploy, and make sure to always use the latest, patched MS Windows Server base image in Compute Engine.
- D. Use Cloud Build to package the application, and deploy to a Google Kubernetes Engine cluster. Use namespaces to separate the environments.

**Correct Answer: B**

**Section:**

**Explanation:**

[https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud#phase\\_1\\_rehost\\_in\\_the\\_cloud](https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud#phase_1_rehost_in_the_cloud)

<https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud>

#### QUESTION 103

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

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

**Correct Answer: C**

**Section:**

**Explanation:**

[https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#perform\\_a\\_bluegreen\\_deployment](https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#perform_a_bluegreen_deployment)

#### QUESTION 104

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

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

**Correct Answer: C**

**Section:**

#### QUESTION 105

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. Create a subscription for each downstream development team to subscribe to their step's topic.
- C. Create a Pub/Sub topic for timesheet submissions. Create a subscription for each downstream development team to subscribe to the topic.
- D. Create a timesheet microservice deployed to Google Kubernetes Engine. The microservice calls each downstream step and waits for a successful response before calling the next step.

**Correct Answer: C**

**Section:**

#### QUESTION 106

You are designing an application that uses a microservices architecture. You are planning to deploy the application in the cloud and on-premises. You want to make sure the application can scale up on demand and also use managed services as much as possible. What should you do?

- A. Deploy open source Istio in a multi-cluster deployment on multiple Google Kubernetes Engine (GKE) clusters managed by Anthos.
- B. Create a GKE cluster in each environment with Anthos, and use Cloud Run for Anthos to deploy your application to each cluster.
- C. Install a GKE cluster in each environment with Anthos, and use Cloud Build to create a Deployment for your application in each cluster.
- D. Create a GKE cluster in the cloud and install open-source Kubernetes on-premises. Use an external load balancer service to distribute traffic across the two environments.

**Correct Answer: B**

**Section:**

**Explanation:**

<https://cloud.google.com/anthos/run>

Integrated with Anthos, Cloud Run for Anthos provides a flexible serverless development platform for hybrid and multicloud environments. Cloud Run for Anthos is Google's managed and fully supported Knative offering, an open source project that enables serverless workloads on Kubernetes.

#### QUESTION 107

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

**Correct Answer: A**

**Section:**

**Explanation:**

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

#### QUESTION 108

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

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

**Correct Answer: D**

**Section:**

**QUESTION 109**

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.

**Correct Answer: D**

**Section:**

**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.

**QUESTION 110**

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 images. Review the vulnerability results before each deployment.
- B. Enable Container Analysis, and upload new container images to Artifact Registry. Review the vulnerability results before each deployment.
- C. Enable Container Analysis, and upload new container images to Artifact Registry. Review the critical vulnerability results before each deployment.
- D. Use the Container Analysis REST API to call Container Analysis to scan new container images. Review the vulnerability results before each deployment.

**Correct Answer: D**

**Section:**

**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).

**QUESTION 111**

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

**Correct Answer: D**

**Section:**

**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.

#### QUESTION 112

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

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

**Correct Answer: A**

**Section:**

**Explanation:**

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

#### QUESTION 113

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

**Correct Answer: D**

**Section:**

#### QUESTION 114

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

**Correct Answer: B**

**Section:**

#### QUESTION 115

You are developing a Java Web Server that needs to interact with Google Cloud services via the Google Cloud API on the user's behalf. Users should be able to authenticate to the Google Cloud API using their Google Cloud identities. Which workflow should you implement in your web application?

- A. 1) When a user arrives at your application, prompt them for their Google username and password. 2) Store an SHA password hash in your application's database along with the user's username. 3) The application

authenticates to the Google Cloud API using HTTPs requests with the user's username and password hash in the Authorization request header.

- B. 1) When a user arrives at your application, prompt them for their Google username and password. 2) Forward the user's username and password in an HTTPS request to the Google Cloud authorization server, and request an access token. 3) The Google server validates the user's credentials and returns an access token to the application. 4) The application uses the access token to call the Google Cloud API.
- C. 1) When a user arrives at your application, route them to a Google Cloud consent screen with a list of requested permissions that prompts the user to sign in with SSO to their Google Account. 2) After the user signs in and provides consent, your application receives an authorization code from a Google server. 3) The Google server returns the authorization code to the user, which is stored in the browser's cookies. 4) The user authenticates to the Google Cloud API using the authorization code in the cookie.
- D. 1) When a user arrives at your application, route them to a Google Cloud consent screen with a list of requested permissions that prompts the user to sign in with SSO to their Google Account. 2) After the user signs in and provides consent, your application receives an authorization code from a Google server. 3) The application requests a Google Server to exchange the authorization code with an access token. 4) The Google server responds with the access token that is used by the application to call the Google Cloud API.

**Correct Answer: D**

**Section:**

**Explanation:**

<https://developers.google.com/identity/protocols/oauth2#webserver>

The Google OAuth 2.0 endpoint supports web server applications that use languages and frameworks such as PHP, Java, Python, Ruby, and ASP.NET. The authorization sequence begins when your application redirects a browser to a Google URL; the URL includes query parameters that indicate the type of access being requested. Google handles the user authentication, session selection, and user consent. The result is an authorization code, which the application can exchange for an access token and a refresh token.

#### QUESTION 116

You recently developed a new application. You want to deploy the application on Cloud Run without a Dockerfile. Your organization requires that all container images are pushed to a centrally managed container repository. How should you build your container using Google Cloud services? (Choose two.)

- A. Push your source code to Artifact Registry.
- B. Submit a Cloud Build job to push the image.
- C. Use the pack build command with pack CLI.
- D. Include the --source flag with the gcloud run deploy CLI command.
- E. Include the --platform=kubernetes flag with the gcloud run deploy CLI command.



**Correct Answer: A, C**

**Section:**

**Explanation:**

<https://cloud.google.com/run/docs/deploying#images>

<https://cloud.google.com/blog/products/containers-kubernetes/google-cloud-now-supports-buildpacks>

#### QUESTION 117

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. After 1 year, move the logs to another bucket with a 2-year retention period.

**Correct Answer: C**

**Section:**

**Explanation:**

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

**QUESTION 118**

Your company has a new security initiative that requires all data stored in Google Cloud to be encrypted by customer-managed encryption keys. You plan to use Cloud Key Management Service (KMS) to configure access to the keys. You need to follow the 'separation of duties' principle and Google-recommended best practices. What should you do? (Choose two.)

- A. Provision Cloud KMS in its own project.
- B. Do not assign an owner to the Cloud KMS project.
- C. Provision Cloud KMS in the project where the keys are being used.
- D. Grant the roles/cloudkms.admin role to the owner of the project where the keys from Cloud KMS are being used.
- E. Grant an owner role for the Cloud KMS project to a different user than the owner of the project where the keys from Cloud KMS are being used.

**Correct Answer: A, B**

**Section:**

**Explanation:**

[https://cloud.google.com/kms/docs/separation-of-duties#using\\_separate\\_project](https://cloud.google.com/kms/docs/separation-of-duties#using_separate_project)

**QUESTION 119**

You need to migrate a standalone Java application running in an on-premises Linux virtual machine (VM) to Google Cloud in a cost-effective manner. You decide not to take the lift-and-shift approach, and instead you plan to modernize the application by converting it to a container. How should you accomplish this task?

- A. Use Migrate for Anthos to migrate the VM to your Google Kubernetes Engine (GKE) cluster as a container.
- B. Export the VM as a raw disk and import it as an image. Create a Compute Engine instance from the Imported image.
- C. Use Migrate for Compute Engine to migrate the VM to a Compute Engine instance, and use Cloud Build to convert it to a container.
- D. Use Jib to build a Docker image from your source code, and upload it to Artifact Registry. Deploy the application in a GKE cluster, and test the application.

**Correct Answer: D**

**Section:**

**Explanation:**

<https://cloud.google.com/blog/products/application-development/introducing-jib-build-java-docker-images-better>

**QUESTION 120**

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.

**Correct Answer: D**

**Section:**

**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.

**QUESTION 121**

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.

**Correct Answer: D**

**Section:**

**Explanation:**

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

#### QUESTION 122

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.

**Correct Answer: C**

**Section:**

#### QUESTION 123

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

**Correct Answer: C**

**Section:**

**Explanation:**

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/choosing-networks-routing#route-alignment>

#### QUESTION 124

Your company has deployed a new API to a Compute Engine instance. During testing, the API is not behaving as expected. You want to monitor the application over 12 hours to diagnose the problem within the application code without redeploying the application. Which tool should you use?

- A. Cloud Trace
- B. Cloud Monitoring
- C. Cloud Debugger logpoints



D. Cloud Debugger snapshots

**Correct Answer: C**

**Section:**

**Explanation:**

<https://cloud.google.com/debugger/docs/using/logpoints>

Logpoints allow you to inject logging into running services without restarting or interfering with the normal function of the service

#### QUESTION 125

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

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

**Correct Answer: D**

**Section:**

#### QUESTION 126

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 database 2) 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 database 2) Enable private access to allow the frontend to access the Cloud Storage bucket privately
- C. 1) Configure a private IP address for Cloud SQL 2) Use VPC-SC to create a service perimeter 3) Add the Cloud SQL database and the Cloud Storage bucket to the same service perimeter
- D. 1) Configure a private IP address for Cloud SQL 2) Use VPC-SC to create a service perimeter 3) Add the Cloud SQL database and the Cloud Storage bucket to different service perimeters

**Correct Answer: C**

**Section:**

#### QUESTION 127

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 rooms. For each room, create a document that lists the contents of the messages
- B. Create a collection for the rooms. For each room, create a collection that contains a document for each message
- C. Create a collection for the rooms. For each room, create a document that contains a collection for documents, each of which contains a message.
- D. Create a collection for the rooms, and create a document for each room. Create a separate collection for messages, with one document per message. Each room's document contains a list of references to the messages.

**Correct Answer: C**

**Section:**

**Explanation:**

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

#### QUESTION 128

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

**Correct Answer: C**

**Section:**

#### QUESTION 129

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

**Correct Answer: C**

**Section:**

**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.

#### QUESTION 130

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

- A. 1) Create a managed instance group. 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.
- B. 1) Create an unmanaged instance group. Replicate the static content across the VMs. 2) Create an external HTTP(S) load balancer 3) Enable Cloud CDN, and send traffic to the unmanaged instance group.
- C. 1) Create a Standard storage class, regional Cloud Storage bucket. 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
- D. 1) Create a Standard storage class, multi-regional Cloud Storage bucket. 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.

**Correct Answer: D**

**Section:**

#### QUESTION 131

You are writing from a Go application to a Cloud Spanner database. You want to optimize your application's performance using Google-recommended best practices. What should you do?

- A. Write to Cloud Spanner using Cloud Client Libraries.
- B. Write to Cloud Spanner using Google API Client Libraries
- C. Write to Cloud Spanner using a custom gRPC client library.
- D. Write to Cloud Spanner using a third-party HTTP client library.



**Correct Answer: A**

**Section:**

**Explanation:**

<https://cloud.google.com/apis/docs/cloud-client-libraries>

"Cloud Client Libraries are the recommended option for accessing Cloud APIs programmatically, where available. Cloud Client Libraries use the latest client library models"

<https://cloud.google.com/apis/docs/client-libraries-explained>

<https://cloud.google.com/go/docs/reference>

**QUESTION 132**

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.

**Correct Answer: A**

**Section:**

**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.

**QUESTION 133**

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.

**Correct Answer: D**

**Section:**

**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.

**QUESTION 134**

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.

**Correct Answer: A**

**Section:**

**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.

#### QUESTION 135

You are developing a new web application using Cloud Run and committing code to Cloud Source Repositories. You want to deploy new code in the most efficient way possible. You have already created a Cloud Build YAML file that builds a container and runs the following command: `gcloud run deploy`. What should you do next?

- A. Create a Pub/Sub topic to be notified when code is pushed to the repository. Create a Pub/Sub trigger that runs the build file when an event is published to the topic.
- B. Create a build trigger that runs the build file in response to a repository code being pushed to the development branch.
- C. Create a webhook build trigger that runs the build file in response to HTTP POST calls to the webhook URL.
- D. Create a Cron job that runs the following command every 24 hours: `gcloud builds submit`.

**Correct Answer: B**

**Section:**

**Explanation:**

<https://cloud.google.com/build/docs/triggers>

Cloud Build uses build triggers to enable CI/CD automation. You can configure triggers to listen for incoming events, such as when a new commit is pushed to a repository or when a pull request is initiated, and then automatically execute a build when new events come in. You can also configure triggers to build code on any changes to your source repository or only on changes that match certain criteria.

#### QUESTION 136

Your team has created an application that is hosted on a Google Kubernetes Engine (GKE) cluster. You need to connect the application to a legacy REST service that is deployed in two GKE clusters in two different regions. You want to connect your application to the legacy service in a way that is resilient and requires the fewest number of steps. You also want to be able to run probe-based health checks on the legacy service on a separate port. How should you set up the connection?

- A. Use Traffic Director with a sidecar proxy to connect the application to the service.
- B. Use a proxyless Traffic Director configuration to connect the application to the service.
- C. Configure the legacy service's firewall to allow health checks originating from the proxy.
- D. Configure the legacy service's firewall to allow health checks originating from the application.
- E. Configure the legacy service's firewall to allow health checks originating from the Traffic Director control plane.

**Correct Answer: A, C**

**Section:**

**Explanation:**

<https://cloud.google.com/traffic-director/docs/advanced-setup#routing-rule-maps> <https://cloud.google.com/traffic-director/docs/advanced-setup>

A) Using Traffic Director with a sidecar proxy can provide resilience for your application by allowing for failover to the secondary region in the event of an outage. The sidecar proxy can route traffic to the legacy service in either of the two GKE clusters, ensuring high availability. C. Configuring the legacy service's firewall to allow health checks originating from the proxy allows the proxy to periodically check the health of the legacy service and ensure that it is functioning properly. This helps to ensure that traffic is only routed to healthy instances of the legacy service, further improving the resilience of the setup.

#### QUESTION 137

You work for a financial services company that has a container-first approach. Your team develops microservices applications. You have a Cloud Build pipeline that creates a container image, runs regression tests, and publishes the image to Artifact Registry. You need to ensure that only containers that have passed the regression tests are deployed to Google Kubernetes Engine (GKE) clusters. You have already enabled Binary Authorization on the GKE

clusters What should you do next?

- A. Deploy Voucher Server and Voucher Client Components. After a container image has passed the regression tests, run Voucher Client as a step in the Cloud Build pipeline.
- B. Set the Pod Security Standard level to Restricted for the relevant namespaces Digitally sign the container images that have passed the regression tests as a step in the Cloud Build pipeline.
- C. Create an attestor and a policy. Create an attestation for the container images that have passed the regression tests as a step in the Cloud Build pipeline.
- D. Create an attestor and a policy Run a vulnerability scan to create an attestation for the container image as a step in the Cloud Build pipeline.

**Correct Answer: C**

**Section:**

#### QUESTION 138

You have an ecommerce application hosted in Google Kubernetes Engine (GKE) that receives external requests and forwards them to third-party APIs external to Google Cloud. The third-party APIs are responsible for credit card processing, shipping, and inventory management using the process shown in the diagram.

Your customers are reporting that the ecommerce application is running slowly at unpredictable times. The application doesn't report any metrics You need to determine the cause of the inconsistent performance What should you do?



- A. Install the Ops Agent inside your container and configure it to gather application metrics.
- B. Install the OpenTelemetry library for your respective language, and instrument your application.
- C. Modify your application to read and forward the x-Cloud-Trace-context header when it calls the downstream services
- D. Enable Managed Service for Prometheus on the GKE cluster to gather application metrics.

**Correct Answer: B**

**Section:**

#### QUESTION 139

You are deploying a Python application to Cloud Run using Cloud Build. The Cloud Build pipeline is shown below:

```

steps:
- name: python
  entrypoint: pip
  args: ["install", "-r", "requirements.txt", "--user"]

- name: 'gcr.io/cloud-builders/docker'
  args: ['build', '-t',
        'us-central1-docker.pkg.dev/${PROJECT_ID}/${_REPO_NAME}/myimage:${SHORT_SHA}',
        '.']

- name: 'gcr.io/cloud-builders/docker'
  args: ['push', 'us-central1-
docker.pkg.dev/${PROJECT_ID}/${_REPO_NAME}/myimage:${SHORT_SHA}']

- name: google/cloud-sdk
  args: ['gcloud', 'run', 'deploy', 'helloworld-${SHORT_SHA}',
        '--image=us-central1-
docker.pkg.dev/${PROJECT_ID}/${_REPO_NAME}/myimage:${SHORT_SHA}',
        '--region', 'us-central1', '--platform', 'managed',
        '--allow-unauthenticated']

```

You want to optimize deployment times and avoid unnecessary steps. What should you do?

- A. Remove the step that pushes the container to Artifact Registry.
- B. Add the `---cache-from` argument to the Docker build step in your build config file.
- C. Store image artifacts in a Cloud Storage bucket in the same region as the Cloud Run instance.
- D. Deploy a new Docker registry in a VPC and use Cloud Build worker pools inside the VPC to run the build pipeline.

**Correct Answer: B**

**Section:**

**Explanation:**

[https://cloud.google.com/build/docs/optimize-builds/speeding-up-builds#using\\_a\\_cached\\_docker\\_image](https://cloud.google.com/build/docs/optimize-builds/speeding-up-builds#using_a_cached_docker_image)

#### QUESTION 140

You are a developer at a large organization. You are deploying a web application to Google Kubernetes Engine (GKE). The DevOps team has built a CI/CD pipeline that uses Cloud Deploy to deploy the application to Dev Test,

and Prod clusters in GKE. After Cloud Deploy successfully deploys the application to the Dev cluster you want to automatically promote it to the Test Cluster. How should you configure this process following Google-recommended best practices?

- A. 1 Create a Cloud Build trigger that listens for SUCCEEDED Pub/Sub messages from the clouddeploy-operations topic. 2 Configure Cloud Build to include a step that promotes the application to the Test cluster
- B. 1 Create a Cloud Function that calls the Google Cloud Deploy API to promote the application to the Test cluster 2 Configure this function to be triggered by SUCCEEDED Pub/Sub messages from the cloud-builds topic
- C. 1 Create a Cloud Function that calls the Google Cloud Deploy API to promote the application to the Test cluster 2 Configure this function to be triggered by SUCCEEDED Pub/Sub messages from the clouddeploy operations topic
- D. 1 Create a Cloud Build pipeline that uses the gke-deploy builder 2 Create a Cloud Build trigger that listens to SUCCEEDED Pub/Sub messages from the cloud-builds topic 3 Configure this pipeline to run a deployment step to the Test cluster

**Correct Answer: B**

**Section:**

#### QUESTION 141

You need to containerize a web application that will be hosted on Google Cloud behind a global load balancer with SSL certificates. You don't have the time to develop authentication at the application level, and you want to offload SSL encryption and management from your application. You want to configure the architecture using managed services where possible What should you do?

- A. Host the application on Compute Engine, and configure Cloud Endpoints for your application.
- B. Host the application on Google Kubernetes Engine and use Identity-Aware Proxy (IAP) with Cloud Load Balancing and Google-managed certificates.
- C. Host the application on Google Kubernetes Engine, and deploy an NGINX Ingress Controller to handle authentication.
- D. Host the application on Google Kubernetes Engine, and deploy cert-manager to manage SSL certificates.

**Correct Answer: B**

**Section:**

#### QUESTION 142

You are a developer at a financial institution You use Cloud Shell to interact with Google Cloud services. User data is currently stored on an ephemeral disk however a recently passed regulation mandates that you can no longer store sensitive information on an ephemeral disk. You need to implement a new storage solution for your user data You want to minimize code changes Where should you store your user data'?

- A. Store user data on a Cloud Shell home disk and log in at least every 120 days to prevent its deletion
- B. Store user data on a persistent disk in a Compute Engine instance
- C. Store user data in BigQuery tables
- D. Store user data in a Cloud Storage bucket

**Correct Answer: B**

**Section:**

#### QUESTION 143

You need to load-test a set of REST API endpoints that are deployed to Cloud Run. The API responds to HTTP POST requests Your load tests must meet the following requirements:

- \* Load is initiated from multiple parallel threads
- \* User traffic to the API originates from multiple source IP addresses.
- \* Load can be scaled up using additional test instances

You want to follow Google-recommended best practices How should you configure the load testing'?

- A. Create an image that has cURL installed and configure cURL to run a test plan Deploy the image in a managed instance group, and run one instance of the image for each VM.
- B. Create an image that has cURL installed and configure cURL to run a test plan Deploy the image in an unmanaged instance group, and run one instance of the image for each VM.
- C. Deploy a distributed load testing framework on a private Google Kubernetes Engine Cluster Deploy additional Pods as needed to initiate more traffic and support the number of concurrent users.



D. Download the container image of a distributed load testing framework on Cloud Shell Sequentially start several instances of the container on Cloud Shell to increase the load on the API.

**Correct Answer: C**

**Section:**

#### QUESTION 144

You are a lead developer working on a new retail system that runs on Cloud Run and Firestore. A web UI requirement is for the user to be able to browse through alt products. A few months after go-live, you notice that Cloud Run instances are terminated with HTTP 500: Container instances are exceeding memory limits errors during busy times

This error coincides with spikes in the number of Firestore queries

You need to prevent Cloud Run from crashing and decrease the number of Firestore queries. You want to use a solution that optimizes system performance What should you do?

- A. Create a custom index over the products
- B. Modify the query that returns the product list using cursors with limits
- C. Modify the Cloud Run configuration to increase the memory limits
- D. Modify the query that returns the product list using integer offsets

**Correct Answer: B**

**Section:**

#### QUESTION 145

You are using Cloud Run to host a global ecommerce web application. Your company's design team is creating a new color scheme for the web app. You have been tasked with determining whether the new color scheme will increase sales You want to conduct testing on live production traffic How should you design the study?

- A. Use an external HTTP(S) load balancer to route a predetermined percentage of traffic to two different color schemes of your application Analyze the results to determine whether there is a statistically significant difference in sales.
- B. Use an external HTTP(S) load balancer to route traffic to the original color scheme while the new deployment is created and tested After testing is complete reroute all traffic to the new color scheme Analyze the results to determine whether there is a statistically significant difference in sales.
- C. Enable a feature flag that displays the new color scheme to half of all users. Monitor sales to see whether they increase for this group of users.
- D. Use an external HTTP(S) load balancer to mirror traffic to the new version of your application Analyze the results to determine whether there is a statistically significant difference in sales.

**Correct Answer: A**

**Section:**

#### QUESTION 146

Your team is setting up a build pipeline for an application that will run in Google Kubernetes Engine (GKE). For security reasons, you only want images produced by the pipeline to be deployed to your GKE cluster. Which combination of Google Cloud services should you use?

- A. Google Cloud Deploy, Artifact Registry, and Google Cloud Armor
- B. Google Cloud Deploy, Cloud Storage and Google Cloud Armor
- C. Cloud Build, Cloud Storage, and Binary Authorization
- D. Cloud Build, Artifact Registry and Binary Authorization

**Correct Answer: A**

**Section:**

#### QUESTION 147

You are developing an online gaming platform as a microservices application on Google Kubernetes Engine (GKE). Users on social media are complaining about long loading times for certain URL requests to the application.

You need to investigate performance bottlenecks in the application and identify. which HTTP requests have a significantly high latency span in user requests What should you do?

- A. Instrument your microservices by installing the OpenTelemetry tracing package Update your application code to send traces to Trace for inspection and analysis Create an analysis report on Trace to analyze user requests
- B. Configure GKE workload metrics using kubectl Select all Pods to send their metrics to Cloud Monitoring. Create a custom dashboard of application metrics in Cloud Monitoring to determine performance bottlenecks of your GKE cluster
- C. Install tcpdump on your GKE nodes. Run tcpdump to capture network traffic over an extended period of time to collect data Analyze the data files using Wireshark to determine the cause of high latency
- D. Update your microservices to log HTTP request methods and URL paths to STDOUT Use the logs router to send container logs to Cloud Logging Create filters in Cloud Logging to evaluate the latency of user requests across different methods and URL paths.

**Correct Answer: B**

**Section:**

#### QUESTION 148

You are working on a new application that is deployed on Cloud Run and uses Cloud Functions Each time new features are added, new Cloud Functions and Cloud Run services are deployed You use ENV variables to keep track of the services and enable interservice communication but the maintenance of the ENV variables has become difficult. You want to implement dynamic discovery in a scalable way. What should you do?

- A. Create a Service Directory Namespace Use API calls to register the services during deployment, and query during runtime.
- B. Configure your microservices to use the Cloud Run Admin and Cloud Functions APIs to query for deployed Cloud Run services and Cloud Functions in the Google Cloud project.
- C. Deploy Hashicorp Consul on a single Compute Engine Instance Register the services with Consul during deployment and query during runtime
- D. Rename the Cloud Functions and Cloud Run services endpoints using a well-documented naming convention

**Correct Answer: B**

**Section:**

#### QUESTION 149

You have deployed a Java application to Cloud Run. Your application requires access to a database hosted on Cloud SQL Due to regulatory requirements: your connection to the Cloud SQL instance must use its internal IP address. How should you configure the connectivity while following Google-recommended best practices'?

- A. Configure your Cloud Run service with a Cloud SQL connection.
- B. Configure your Cloud Run service to use a Serverless VPC Access connector
- C. Configure your application to use the Cloud SQL Java connector
- D. Configure your application to connect to an instance of the Cloud SQL Auth proxy

**Correct Answer: A**

**Section:**

#### QUESTION 150

Your team is developing a Cloud Function triggered by Cloud Storage Events. You want to accelerate testing and development of your Cloud Function while following Google-recommended best practices. What should you do?

- A. Install the Functions Frameworks library, and configure the Cloud Function on localhost. Make a copy of the function, and make edits to the new version Test the new version using curl.
- B. Make a copy of the Cloud Function, and rewrite the code to be HTTP-triggered Edit and test the new version by triggering the HTTP endpoint. Send mock requests to the new function to evaluate the functionality.
- C. Make a copy of the Cloud Function in the Google Cloud Console Use the Cloud console's in-line editor to make source code changes to the new function Modify your web application to call the new function and test the new version in production.
- D. Create a new Cloud Function that is triggered when Cloud Audit Logs detects the cloudfunctions. functions. sourceCodeSet operation in the original Cloud Function Send mock requests to the new function to evaluate the functionality.

**Correct Answer: A**

**Section:**

**QUESTION 151**

Your team recently deployed an application on Google Kubernetes Engine (GKE). You are monitoring your application and want to be alerted when the average memory consumption of your containers is under 20% or above 80% How should you configure the alerts?

- A. Create a Cloud Function that consumes the Monitoring API. Create a schedule to trigger the Cloud Function hourly and alert you if the average memory consumption is outside the defined range
- B. In Cloud Monitoring, create an alerting policy to notify you if the average memory consumption is outside the defined range
- C. Create a Cloud Function that runs on a schedule, executes kubectl top on all the workloads on the cluster, and sends an email alert if the average memory consumption is outside the defined range
- D. Write a script that pulls the memory consumption of the instance at the OS level and sends an email alert if the average memory consumption is outside the defined range

**Correct Answer: B**

**Section:**

**QUESTION 152**

You are developing a new application. You want the application to be triggered only when a given file is updated in your Cloud Storage bucket. Your trigger might change, so your process must support different types of triggers. You want the configuration to be simple so that multiple team members can update the triggers in the future. What should you do?

- A. Create an Eventarc trigger that monitors your Cloud Storage bucket for a specific filename, and set the target as Cloud Run.
- B. Configure Cloud Storage events to be sent to Pub/Sub, and use Pub/Sub events to trigger a Cloud Build job that executes your application.
- C. Configure a Firebase function that executes your application and is triggered when an object is updated in Cloud Storage.
- D. Configure a Cloud Function that executes your application and is triggered when an object is updated in Cloud Storage.

**Correct Answer: C**

**Section:**

**QUESTION 153**

You are developing a flower ordering application. Currently you have three microservices.

- \* Order Service (receives the orders).
- \* Order Fulfillment Service (processes the orders).
- \* Notification Service (notifies the customer when the order is filled).

You need to determine how the services will communicate with each other. You want incoming orders to be processed quickly and you need to collect order information for fulfillment. You also want to make sure orders are not lost between your services and are able to communicate asynchronously. How should the requests be processed?

A.



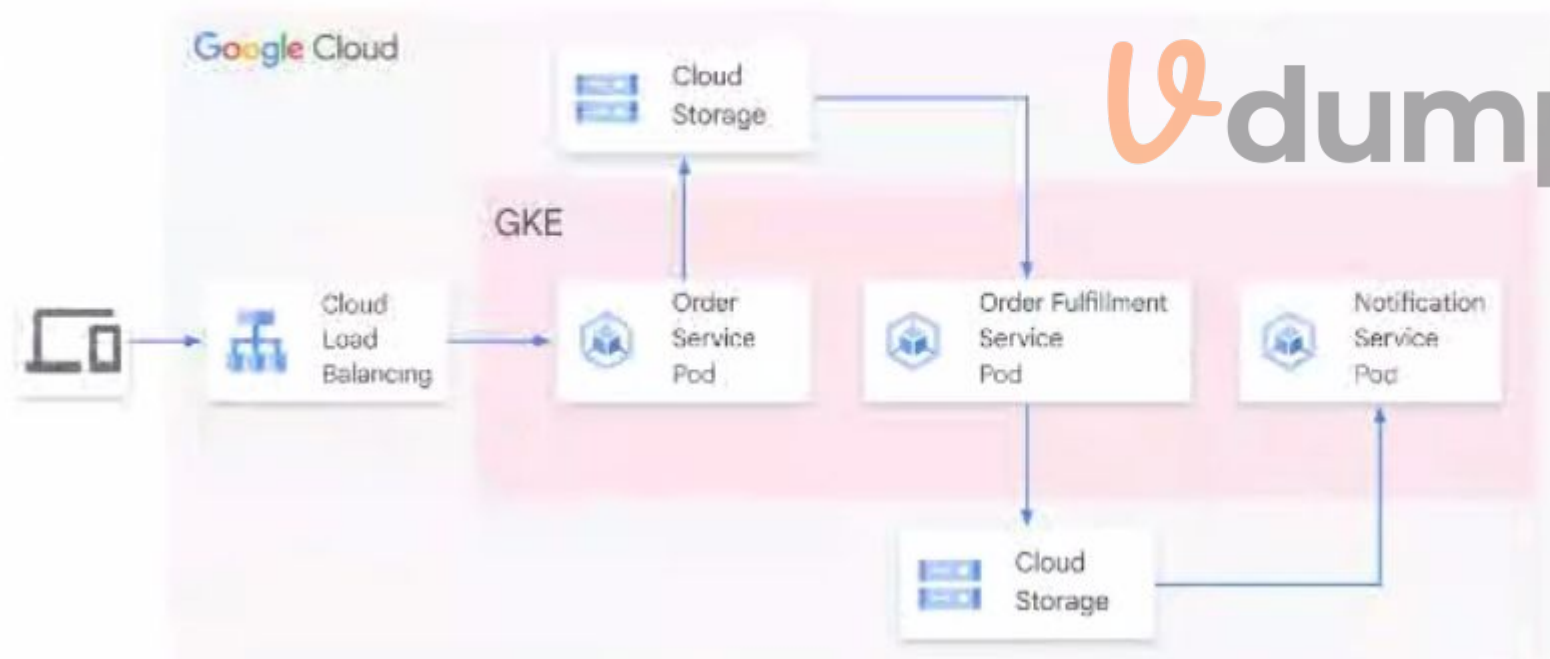


Order request → Order Service → Order Fulfillment Service → Notification Service



B.

Order request → Order Service → Cloud Storage bucket → Order Fulfillment Service → Cloud Storage bucket → Notification Service



C.

Order request → Order Service → Firestore database → Order Fulfillment Service → Firestore database → Notification Service



D.

Order request → Order Service → Pub/Sub queue → Order Fulfillment Service → Firestore database → Pub/Sub queue → Notification Service



Correct Answer: A  
Section:

QUESTION 154

You need to deploy an internet-facing microservices application to Google Kubernetes Engine (GKE). You want to validate new features using the A/B testing method. You have the following requirements for deploying new

container image releases

\* There is no downtime when new container images are deployed.

\* New production releases are tested and verified using a subset of production users.

What should you do?

- A. 1 Configure your CI/CD pipeline to update the Deployment manifest file by replacing the container version with the latest version. 2 Recreate the Pods in your cluster by applying the Deployment manifest file. 3 Validate the application's performance by comparing its functionality with the previous release version and roll back if an issue arises.
- B. 1 install the Anthos Service Mesh on your GKE cluster. 2 Create two Deployments on the GKE cluster and label them with different version names. 3 Create a VirtualService with a routing rule to send a small percentage of traffic to the Deployment that references the new version of the application.
- C. 1 Create a second namespace on GKE for the new release version. 2 Create a Deployment configuration for the second namespace with the desired number of Pods. 3 Deploy new container versions in the second namespace. 4 Update the ingress configuration to route traffic to the namespace with the new container versions.
- D. 1. Implement a rolling update pattern by replacing the Pods gradually with the new release version. 2 Validate the application's performance for the new subset of users during the rollout and roll back if an issue arises.

**Correct Answer: C**

**Section:**

#### QUESTION 155

You are developing an event-driven application. You have created a topic to receive messages sent to Pub/Sub. You want those messages to be processed in real time. You need the application to be independent from any other system and only incur compute costs when new messages arrive. You want to configure the simplest and most efficient architecture. What should you do?

- A. Deploy your code on Cloud Functions. Use a Pub/Sub trigger to invoke the Cloud Function. Use the Pub/Sub API to create a pull subscription to the Pub/Sub topic and read messages from it.
- B. Deploy your code on Cloud Functions. Use a Pub/Sub trigger to handle new messages in the topic.
- C. Deploy the application on Google Kubernetes Engine. Use the Pub/Sub API to create a pull subscription to the Pub/Sub topic and read messages from it.
- D. Deploy the application on Compute Engine. Use a Pub/Sub push subscription to process new messages in the topic.

**Correct Answer: B**

**Section:**

**Explanation:**

<https://cloud.google.com/functions/docs/calling/pubsub>

#### QUESTION 156

You are a developer at a large organization. Your team uses Git for source code management (SCM). You want to ensure that your team follows Google-recommended best practices to manage code to drive higher rates of software delivery. Which SCM process should your team use?

- A. Each developer commits their code to the main branch before each product release, conducts testing, and rolls back if integration issues are detected.
- B. Each group of developers copies the repository, commits their changes to their repository, and merges their code into the main repository before each product release.
- C. Each developer creates a branch for their own work, commits their changes to their branch, and merges their code into the main branch daily.
- D. Each group of developers creates a feature branch from the main branch for their work, commits their changes to their branch, and merges their code into the main branch after the change advisory board approves it.

**Correct Answer: D**

**Section:**

**Explanation:**

Use a centralized repository. A centralized repository is a single location where all of your team's code is stored. This makes it easy for everyone to access the latest code, and it also helps to prevent conflicts. Use branches. Branches are a way to create a separate version of the code for development purposes. This allows developers to work on new features or bug fixes without affecting the main branch of the code.