

Google

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer



NEW QUESTION 1

- (Exam Topic 1)

Which database should HipLocal use for storing user activity?

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

Answer: A

NEW QUESTION 2

- (Exam Topic 1)

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

Which configuration should they choose?

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

Answer: B

NEW QUESTION 3

- (Exam Topic 1)

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.

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

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

What should they do?

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

Answer: C

NEW QUESTION 5

- (Exam Topic 2)

You are developing a single-player mobile game backend that has unpredictable traffic patterns as users interact with the game throughout the day and night. You want to optimize costs by ensuring that you have enough resources to handle requests, but minimize over-provisioning. You also want the system to handle traffic spikes efficiently. Which compute platform should you use?

- A. Cloud Run
- B. Compute Engine with managed instance groups
- C. Compute Engine with unmanaged instance groups
- D. Google Kubernetes Engine using cluster autoscaling

Answer: A

NEW QUESTION 6

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 7

- (Exam Topic 2)

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 8

- (Exam Topic 2)

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 9

- (Exam Topic 2)

You made a typo in a low-level Linux configuration file that prevents your Compute Engine instance from booting to a normal run level. You just created the Compute Engine instance today and have done no other maintenance on it, other than tweaking files. How should you correct this error?

- A. Download the file using scp, change the file, and then upload the modified version
- B. Configure and log in to the Compute Engine instance through SSH, and change the file
- C. Configure and log in to the Compute Engine instance through the serial port, and change the file
- D. Configure and log in to the Compute Engine instance using a remote desktop client, and change the file

Answer: C

Explanation:

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

NEW QUESTION 10

- (Exam Topic 2)

You are using Cloud Run to host a web application. You need to securely obtain the application project ID and region where the application is running and display this information to users. You want to use the most performant approach. What should you do?

- A. Use HTTP requests to query the available metadata server at the <http://metadata.google.internal/endpoint> with the Metadata-Flavor: Google header.
- B. In the Google Cloud console, navigate to the Project Dashboard and gather configuration details. Navigate to the Cloud Run “Variables & Secrets” tab, and add the desired environment variables in Key:Value format.
- C. In the Google Cloud console, navigate to the Project Dashboard and gather configuration detail
- D. Write the application configuration information to Cloud Run's in-memory container filesystem.
- E. Make an API call to the Cloud Asset Inventory API from the application and format the request to include instance metadata.

Answer: B

NEW QUESTION 10

- (Exam Topic 2)

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

Answer: B

Explanation:

Reference: [https://medium.com/google-cloud/stackdriver-monitoring-automation-part-3-uptime-checks- 476b8507f59c](https://medium.com/google-cloud/stackdriver-monitoring-automation-part-3-uptime-checks-476b8507f59c)

NEW QUESTION 13

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 15

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 18

- (Exam Topic 2)

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.

Answer: B

Explanation:

Reference: <https://cloud.google.com/compute/docs/api/how-tos/authorization>

NEW QUESTION 23

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 27

- (Exam Topic 2)

You are using the Cloud Client Library to upload an image in your application to Cloud Storage. Users of the application report that occasionally the upload does not complete and the client library reports an HTTP 504 Gateway Timeout error. You want to make the application more resilient to errors. What changes to the application should you make?

- A. Write an exponential backoff process around the client library call.
- B. Write a one-second wait time backoff process around the client library call.
- C. Design a retry button in the application and ask users to click if the error occurs.
- D. Create a queue for the object and inform the users that the application will try again in 10 minutes.

Answer: A

NEW QUESTION 31

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 33

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 38

- (Exam Topic 2)

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

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

Answer: B

Explanation:

Reference: <https://cloud.google.com/compute/docs/metadata/overview>

NEW QUESTION 43

- (Exam Topic 2)

You want to notify on-call engineers about a service degradation in production while minimizing development time. What should you do?

- A. Use Cloud Function to monitor resources and raise alerts.
- B. Use Cloud Pub/Sub to monitor resources and raise alerts.
- C. Use Stackdriver Error Reporting to capture errors and raise alerts.
- D. Use Stackdriver Monitoring to monitor resources and raise alerts.

Answer: A

NEW QUESTION 48

- (Exam Topic 2)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 50

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 55

- (Exam Topic 2)

Before promoting your new application code to production, you want to conduct testing across a variety of different users. Although this plan is risky, you want to test the new version of the application with production users and you want to control which users are forwarded to the new version of the application based on their operating system. If bugs are discovered in the new version, you want to roll back the newly deployed version of the application as quickly as possible. What should you do?

- A. Deploy your application on Cloud Run
- B. Use traffic splitting to direct a subset of user traffic to the new version based on the revision tag.
- C. Deploy your application on Google Kubernetes Engine with Anthos Service Mesh
- D. Use traffic splitting to direct a subset of user traffic to the new version based on the user-agent header.
- E. Deploy your application on App Engine
- F. Use traffic splitting to direct a subset of user traffic to the new version based on the IP address.
- G. Deploy your application on Compute Engine
- H. Use Traffic Director to direct a subset of user traffic to the new version based on predefined weights.

Answer: B

NEW QUESTION 60

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 65

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 67

- (Exam Topic 2)

You are monitoring a web application that is written in Go and deployed in Google Kubernetes Engine. You notice an increase in CPU and memory utilization. You need to determine which source code is consuming the most CPU and memory resources. What should you do?

- A. Download, install, and start the Snapshot Debugger agent in your V
- B. Take debug snapshots of the functions that take the longest tim
- C. Review the call stack frame, and identify the local variables at that level in the stack.
- D. Import the Cloud Profiler package into your application, and initialize the Profiler agen
- E. Review the generated flame graph in the Google Cloud console to identify time-intensive functions.
- F. Import OpenTelemetry and Trace export packages into your application, and create the trace provider. Review the latency data for your application on the Trace overview page, and identify where bottlenecks are occurring.
- G. Create a Cloud Logging query that gathers the web application's log
- H. Write a Python script that calculates the difference between the timestamps from the beginning and the end of the application's longest functions to identity time-intensive functions.

Answer: B

NEW QUESTION 71

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 76

- (Exam Topic 2)

One of your deployed applications in Google Kubernetes Engine (GKE) is having intermittent performance issues. Your team uses a third-party logging solution. You want to install this solution on each node in your GKE cluster so you can view the logs. What should you do?

- A. Deploy the third-party solution as a DaemonSet
- B. Modify your container image to include the monitoring software
- C. Use SSH to connect to the GKE node, and install the software manually
- D. Deploy the third-party solution using Terraform and deploy the logging Pod as a Kubernetes Deployment

Answer: A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset#usage_patterns DaemonSets are useful for deploying ongoing background tasks that you need to run on all or certain nodes, and which do not require user intervention. Examples of such tasks include storage daemons like ceph, log collection daemons like fluent-bit, and node monitoring daemons like collectd.

NEW QUESTION 79

- (Exam Topic 2)

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.

Answer: AB

Explanation:

https://cloud.google.com/kms/docs/separation-of-duties#using_separate_project

NEW QUESTION 84

- (Exam Topic 2)

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

A) Manually copy the photos to webphotos-prod.

B) Add a startup script in the application's app.yaml file to move the photos from webphotos-staging to webphotos-prod.

C) Add a build step in the cloudbuild.yaml file before the promotion step with the arguments:

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

D) Add a build step in the cloudbuild.yaml file before the promotion step with the arguments:

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

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

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

NEW QUESTION 89

- (Exam Topic 2)

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

A. Reboot the machine.

B. Enable and check the serial port output.

C. Delete the machine and create a new one.

D. Take a snapshot of the disk and attach it to a new machine.

Answer: A

NEW QUESTION 90

- (Exam Topic 2)

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

A. Process the messages with a Dataproc job, and write the output to storage.

B. Process the messages with a Dataflow streaming pipeline using Apache Beam's PubSubIO package, and write the output to storage.

C. Process the messages with a Cloud Function, and write the results to a BigQuery location where you can run a job to deduplicate the data.

D. Retrieve the messages with a Dataflow streaming pipeline, store them in Cloud Bigtable, and use another Dataflow streaming pipeline to deduplicate messages.

Answer: B

Explanation:

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

NEW QUESTION 95

- (Exam Topic 2)

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

A. Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department.

B. Verify the provided JSON Web Token within the application.

C. Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department.

D. Issue client-side certificates to everybody in the finance team and verify the certificates in the application.

E. Configure Cloud Armor Security Policies to restrict access to only corporate IP address range.

F. Verify the provided JSON Web Token within the application.

G. Configure Cloud Armor Security Policies to restrict access to only corporate IP address range.

H. Issue client side certificates to everybody in the finance team and verify the certificates in the application.

Answer: A

Explanation:

https://cloud.google.com/iap/docs/signed-headers-howto#securing_iap_headers (<https://cloud.google.com/endpoints/docs/openapi/authenticating-users-google-id>).

<https://cloud.google.com/armor/docs/security-policy-overview#:~:text=Google%20Cloud%20Armor%20security> "Google Cloud Armor security policies protect your application by providing Layer 7 filtering and by

scrubbing incoming requests for common web attacks or other Layer 7 attributes to potentially block traffic before it reaches your load balanced backend services

or backend buckets"

NEW QUESTION 97

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 102

- (Exam Topic 2)

Your company's development teams want to use Cloud Build in their projects to build and push Docker images to Container Registry. The operations team requires all Docker images to be published to a centralized, securely managed Docker registry that the operations team manages. What should you do?

- A. Use Container Registry to create a registry in each development team's projec
- B. Configure the Cloud Build build to push the Docker image to the project's registr
- C. Grant the operations team access to each development team's registry.
- D. Create a separate project for the operations team that has Container Registry configure
- E. Assign appropriate permissions to the Cloud Build service account in each developer team's project to allow access to the operation team's registry.
- F. Create a separate project for the operations team that has Container Registry configure
- G. Create a Service Account for each development team and assign the appropriate permissions to allow it access to the operations team's registr
- H. Store the service account key file in the source code repository and use it to authenticate against the operations team's registry.
- I. Create a separate project for the operations team that has the open source Docker Registry deployed on a Compute Engine virtual machine instanc
- J. Create a username and password for each development tea
- K. Store the username and password in the source code repository and use it to authenticate against the operations team's Docker registry.

Answer: A

Explanation:

Reference: <https://cloud.google.com/container-registry/>

NEW QUESTION 107

- (Exam Topic 2)

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

- A. Deploy the Pub/Sub and Cloud Run emulators on your local machin
- B. Deploy the application locally, and change the logging level in the application to DEBUG or INF
- C. Write mock messages to topic A, and then analyze the logs.
- D. Use the gcloud CLI to write mock messages to topic
- E. Change the logging level in the application to DEBUG or INFO, and then analyze the logs.
- F. Deploy the Pub/Sub emulator on your local machin
- G. Point the production application to your local Pub/Sub topic
- H. Write mock messages to topic A, and then analyze the logs.
- I. Use the Google Cloud console to write mock messages to topic
- J. Change the logging level in the application to DEBUG or INFO, and then analyze the logs.

Answer: A

NEW QUESTION 109

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 110

- (Exam Topic 2)

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 Ru
- B. Use revisions to separate the environments.

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

Answer: B

Explanation:

https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud#phase_1_rehost_in_
<https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud>

NEW QUESTION 112

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 117

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 120

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 125

- (Exam Topic 2)

Your company has a BigQuery data mart that provides analytics information to hundreds of employees. One user wants to run jobs without interrupting important workloads. This user isn't concerned about the time it takes to run these jobs. You want to fulfill this request while minimizing cost to the company and the effort required on your part. What should you do?

- A. Ask the user to run the jobs as batch jobs.
- B. Create a separate project for the user to run jobs.
- C. Add the user as a job.user role in the existing project.
- D. Allow the user to run jobs when important workloads are not running.

Answer: B

NEW QUESTION 126

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 129

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 134

- (Exam Topic 2)

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

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

The command fails with the error:

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

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

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

Answer: A

NEW QUESTION 137

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 140

- (Exam Topic 2)

You work at a rapidly growing financial technology startup. You manage the payment processing application written in Go and hosted on Cloud Run in the Singapore region (asia-southeast1). The payment processing application processes data stored in a Cloud Storage bucket that is also located in the Singapore region.

The startup plans to expand further into the Asia Pacific region. You plan to deploy the Payment Gateway in Jakarta, Hong Kong, and Taiwan over the next six months. Each location has data residency requirements that require customer data to reside in the country where the transaction was made. You want to minimize the cost of these deployments. What should you do?

- A. Create a Cloud Storage bucket in each region, and create a Cloud Run service of the payment processing application in each region.
- B. Create a Cloud Storage bucket in each region, and create three Cloud Run services of the payment processing application in the Singapore region.
- C. Create three Cloud Storage buckets in the Asia multi-region, and create three Cloud Run services of the payment processing application in the Singapore region.
- D. Create three Cloud Storage buckets in the Asia multi-region, and create three Cloud Run revisions of the payment processing application in the Singapore region.

Answer: A

NEW QUESTION 142

- (Exam Topic 2)

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

What should you do to correct the problem?

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

Answer: B

NEW QUESTION 144

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 147

- (Exam Topic 2)

You are developing an HTTP API hosted on a Compute Engine virtual machine instance that needs to be invoked by multiple clients within the same Virtual Private Cloud (VPC). You want clients to be able to get the IP address of the service. What should you do?

- A. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Clients should use this IP address to connect to the service.
- B. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Then, define an A record in Cloud DNS
- C. Clients should use the name of the A record to connect to the service.
- D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[INSTANCE_NAME\].\[ZONE\].c.\[PROJECT_ID\].internal/](https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/).
- E. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[API_NAME\]/\[API_VERSION\]/](https://[API_NAME]/[API_VERSION]/).

Answer: D

NEW QUESTION 152

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 153

- (Exam Topic 2)

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

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

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

Answer: C

NEW QUESTION 158

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 163

- (Exam Topic 2)

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.

Answer: A

Explanation:

Reference: <https://cloud.google.com/compute/docs/instance-groups>

NEW QUESTION 164

- (Exam Topic 2)

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

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

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

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

Answer: AC

NEW QUESTION 165

- (Exam Topic 2)

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

What should you do?

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

Answer: C

NEW QUESTION 167

- (Exam Topic 2)

You are creating and running containers across different projects in Google Cloud. The application you are developing needs to access Google Cloud services from within Google Kubernetes Engine (GKE).

What should you do?

- A. Assign a Google service account to the GKE nodes.
- B. Use a Google service account to run the Pod with Workload Identity.

- C. Store the Google service account credentials as a Kubernetes Secret.
- D. Use a Google service account with GKE role-based access control (RBAC).

Answer: B

Explanation:

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

NEW QUESTION 171

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 173

- (Exam Topic 2)

You are trying to connect to your Google Kubernetes Engine (GKE) cluster using `kubectl` from Cloud Shell. You have deployed your GKE cluster with a public endpoint. From Cloud Shell, you run the following command:

```
gcloud container clusters get-credentials <cluster-name> \
  --zone <zone> --project <project-name> \
```

You notice that the `kubectl` commands time out without returning an error message. What is the most likely cause of this issue?

- A. Your user account does not have privileges to interact with the cluster using `kubectl`.
- B. Your Cloud Shell external IP address is not part of the authorized networks of the cluster.
- C. The Cloud Shell is not part of the same VPC as the GKE cluster.
- D. A VPC firewall is blocking access to the cluster's endpoint.

Answer: B

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/private-clusters#cloud_shell

If you want to use Cloud Shell to access the cluster, you must add the public IP address of your Cloud Shell to the cluster's list of authorized networks.

NEW QUESTION 178

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Professional-Cloud-Developer Practice Exam Features:

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

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