Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

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

NEW QUESTION 1
- (Exam Topic 1)
For this question, refer to the HipLocal case study.
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 deploymen
B. Provide developers with local MySQL instances to conduct testing on the application changes.
C. Migrate data to Bigtabl
D. Instruct the development teams to use the Cloud SDK to emulate a local Bigtable development environment.
E. Move from Cloud SQL to MySQL hosted on Compute Engin
F. Replicate hosts across regions in the Americas and Europ
G. Provide developers with local MySQL instances to conduct testing on the application changes.
H. Migrate data to Firestore in Native mode and set up instan

Answer: B

NEW QUESTION 2
- (Exam Topic 1)
For this question, refer to the HipLocal case study.
A recent security audit discovers that HipLocal's database credentials for their Compute Engine-hosted MySQL databases are stored in plain text on persistent disks. HipLocal needs to reduce the risk of these credentials being stolen. What should they do?

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

Answer: D

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

NEW QUESTION 3
- (Exam Topic 1)
For this question, refer to the HipLocal case study.
How should HipLocal increase their API development speed while continuing to provide the QA team with a stable testing environment that meets feature requirements?

A. Include unit tests in their code, and prevent deployments to QA until all tests have a passing status.
B. Include performance tests in their code, and prevent deployments to QA until all tests have a passing status.
C. Create health checks for the QA environment, and redeploy the APIs at a later time if the environment is unhealthy.
D. Redeploy the APIs to App Engine using Traffic Splittin
E. Do not move QA traffic to the new versions if errors are found.

Answer: B

NEW QUESTION 4
- (Exam Topic 1)
Which service should HipLocal use for their public APIs?

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

Answer: D

NEW QUESTION 5
- (Exam Topic 2)
You 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 developing an application that will allow clients to download a file from your website for a specific period of time. How should you design the application to complete this task while following
Google-recommended best practices?

A. Configure the application to send the file to the client as an email attachment.
B. Generate and assign a Cloud Storage-signed URL for the fil
C. Make the URL available for the client to download.
D. Create a temporary Cloud Storage bucket with time expiration specified, and give download permissions to the bucke
E. Copy the file, and send it to the client.
F. Generate the HTTP cookies with time expiration specifie
G. If the time is valid, copy the file from the Cloud Storage bucket, and make the file available for the client to download.

Answer: B

NEW QUESTION 9
- (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 - heath checks
D. Managed instance group - heath checks

Answer: B

Explanation:
Reference: https://medium.com/google-cloud/stackdriver-monitoring-automation-part-3-uptime-checks- 476b8507f59c

NEW QUESTION 10
- (Exam Topic 2)
Your company has deployed a new API to App Engine Standard environment. During testing, the API is not behaving as expected. You want to monitor the application over time to diagnose the problem within the application code without redeploying the application.
Which tool should you use?

A. Stackdriver Trace
B. Stackdriver Monitoring
C. Stackdriver Debug Snapshots
D. Stackdriver Debug Logpoints

Answer: B

Explanation:

Reference: https://rominirani.com/gcp-stackdriver-tutorial-debug-snapshots-traces-logging-and-logpoints- 1ba49e4780e6

NEW QUESTION 10
- (Exam Topic 2)
You are planning to add unit tests to your application. You need to be able to assert that published Pub/Sub messages are processed by your subscriber in order. You want the unit tests to be cost-effective and reliable. What should you do?

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

Answer: D

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

NEW QUESTION 14
- (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 imag
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 strin
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 strin
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 18
- (Exam Topic 2)
Your company's corporate policy states that there must be a copyright comment at the very beginning of all source files. You want to write a custom step in Cloud Build that is triggered by each source commit. You need the trigger to validate that the source contains a copyright and add one for subsequent steps if not there. What should you do?

A. Build a new Docker container that examines the files in /workspace and then checks and adds a copyright for each source fil
B. Changed files are explicitly committed back to the source repository.
C. Build a new Docker container that examines the files in /workspace and then checks and adds a copyright for each source fil
D. Changed files do not need to be committed back to the source repository.
E. Build a new Docker container that examines the files in a Cloud Storage bucket and then checks and adds a copyright for each source fil
F. Changed files are written back to the Cloud Storage bucket.
G. Build a new Docker container that examines the files in a Cloud Storage bucket and then checks and adds a copyright for each source fil
H. Changed files are explicitly committed back to the source repository.

Answer: A

Explanation:
https://cloud.google.com/build/docs/configuring-builds/pass-data-between-steps#passing_data_using_workspace To pass data between build steps, store the assets produced by the build step in /workspace and these assets will be available to any subsequent build steps.

NEW QUESTION 23
- (Exam Topic 2)
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.

Answer: B

Explanation:
Reference: https://cloud.google.com/logging/docs/agent/logging/configuration

NEW QUESTION 25
- (Exam Topic 2)
You want to re-architect a monolithic application so that it follows a microservices model. You want to accomplish this efficiently while minimizing the impact of this change to the business.
Which approach should you take?

A. Deploy the application to Compute Engine and turn on autoscaling.
B. Replace the application's features with appropriate microservices in phases.
C. Refactor the monolithic application with appropriate microservices in a single effort and deploy it.

D. Build a new application with the appropriate microservices separate from the monolith and replace it when it is complete.

Answer: C

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

NEW QUESTION 26
- (Exam Topic 2)
You developed a JavaScript web application that needs to access Google Drive's API and obtain permission from users to store files in their Google Drives. You need to select an authorization approach for your application. What should you do?

A. Create an API key.
B. Create a SAML token.
C. Create a service account.
D. Create an OAuth Client ID.

Answer: D

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

NEW QUESTION 31
- (Exam Topic 2)
You are parsing a log file that contains three columns: a timestamp, an account number (a string), and a transaction amount (a number). You want to calculate the sum of all transaction amounts for each unique account number efficiently.
Which data structure should you use?

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

Answer: B

NEW QUESTION 35
- (Exam Topic 2)
You recently migrated a monolithic application to Google Cloud by breaking it down into microservices. One of the microservices is deployed using Cloud Functions. As you modernize the application, you make a change to the API of the service that is backward-incompatible. You need to support both existing callers who use the original API and new callers who use the new API. What should you do?

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

Answer: D

Explanation:
Reference: https://cloud.google.com/endpoints/docs/openapi/versioning-an-api

NEW QUESTION 40
- (Exam Topic 2)
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 ot 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.

Answer: AC

NEW QUESTION 41
- (Exam Topic 2)
You are porting an existing Apache/MySQL/PHP application stack from a single machine to Google Kubernetes Engine. You need to determine how to containerize the application. Your approach should follow Google-recommended best practices for availability. What should you do?

A. Package each component in a separate containe
B. Implement readiness and liveness probes.
C. Package the application in a single containe
D. Use a process management tool to manage each component.

E. Package each component in a separate containe
F. Use a script to orchestrate the launch of the components.
G. Package the application in a single containe
H. Use a bash script as an entrypoint to the container, and then spawn each component as a background job.

Answer: A

Explanation:
https://cloud.google.com/blog/products/containers-kubernetes/7-best-practices-for-building-containers https://cloud.google.com/architecture/best-practices-for-building-containers
"classic Apache/MySQL/PHP stack: you might be tempted to run all the components in a single container. However, the best practice is to use two or three different containers: one for Apache, one for MySQL, and potentially one for PHP if you are running PHP-FPM."

NEW QUESTION 45
- (Exam Topic 2)
Your company's product team has a new requirement based on customer demand to autoscale your stateless and distributed service running in a Google Kubernetes Engine (GKE) duster. You want to find a solution that minimizes changes because this feature will go live in two weeks. What should you do?

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

Answer: C

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler
The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload's CPU or memory consumption, or in response to custom metrics reported from within Kubernetes or external metrics from sources outside of your cluster.

NEW QUESTION 48
- (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 51
- (Exam Topic 2)
Your application is composed of a set of loosely coupled services orchestrated by code executed on Compute Engine. You want your application to easily bring up new Compute Engine instances that find and use a specific version of a service. How should this be configured?

A. Define your service endpoint information as metadata that is retrieved at runtime and used to connect to the desired service.
B. Define your service endpoint information as label data that is retrieved at runtime and used to connect to the desired service.
C. Define your service endpoint information to be retrieved from an environment variable at runtime and used to connect to the desired service.
D. Define your service to use a fixed hostname and port to connect to the desired servic
E. Replace the service at the endpoint with your new version.

Answer: A

Explanation:
https://cloud.google.com/service-infrastructure/docs/service-metadata/reference/rest#service-endpoint

NEW QUESTION 54
- (Exam Topic 2)
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.

Answer: A

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

NEW QUESTION 57
- (Exam Topic 2)
Your team is responsible for maintaining an application that aggregates news articles from many different sources. Your monitoring dashboard contains publicly accessible real-time reports and runs on a Compute Engine instance as a web application. External stakeholders and analysts need to access these reports via a secure channel without authentication. How should you configure this secure channel?

A. Add a public IP address to the instanc
B. Use the service account key of the instance to encrypt the traffic.
C. Use Cloud Scheduler to trigger Cloud Build every hour to create an export from the report
D. Store the reports in a public Cloud Storage bucket.
E. Add an HTTP(S) load balancer in front of the monitoring dashboar
F. Configure Identity-Aware Proxy to secure the communication channel.
G. Add an HTTP(S) load balancer in front of the monitoring dashboar
H. Set up a Google-managed SSL certificate on the load balancer for traffic encryption.

Answer: D

Explanation:
https://cloud.google.com/load-balancing/docs/ssl-certificates/google-managed-certs

NEW QUESTION 58
- (Exam Topic 2)
Your team develops stateless services that run on Google Kubernetes Engine (GKE). You need to deploy a new service that will only be accessed by other services running in the GKE cluster. The service will need to scale as quickly as possible to respond to changing load. What should you do?

A. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
B. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a NodePort Service.
C. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
D. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

Answer: C

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

NEW QUESTION 59
- (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 63
- (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 67
- (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 anincremental 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 69
- (Exam Topic 2)
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.

Answer: AC

Explanation:
https://cloud.google.com/run/docs/deploying#images https://cloud.google.com/blog/products/containers-kubernetes/google-cloud-now-supports-buildpacks

NEW QUESTION 74
- (Exam Topic 2)
You have a container deployed on Google Kubernetes Engine. The container can sometimes be slow to launch, so you have implemented a liveness probe. You notice that the liveness probe occasionally fails on
launch. What should you do?

A. Add a startup probe.
B. Increase the initial delay for the liveness probe.
C. Increase the CPU limit for the container.
D. Add a readiness probe.

Answer: B

Explanation:
https://kubernetes.io/docs/tasks/configure-pod-container/configure-liveness-readiness-startup-probes/#configure

NEW QUESTION 75
- (Exam Topic 2)
You have two tables in an ANSI-SQL compliant database with identical columns that you need to quickly combine into a single table, removing duplicate rows from the result set.
What should you do?

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

Answer: C

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

NEW QUESTION 77
- (Exam Topic 2)
You need to deploy resources from your laptop to Google Cloud using Terraform. Resources in your Google Cloud environment must be created using a service account. Your Cloud Identity has the roles/iam.serviceAccountTokenCreator Identity and Access Management (IAM) role and the necessary permissions to deploy the resources using Terraform. You want to set up your development environment to deploy the desired resources following Google-recommended best practices. What should you do?

A. 1) Download the service account's key file in JSON format, and store it locally on your laptop.2) Set the GOOGLE_APPLICATION_CREDENTIALS environment variable to the path of your downloaded key file.
B. 1) Run the following command from a command line: gcloud config set auth/impersonate_service_account service-account-name@project.iam.gserviceacccount.com.2) Set the GOOGLE_OAUTH_ACCESS_TOKEN environment variable to the value that is returned bythe gcloud auth print-access-token command.
C. 1) Run the following command from a command line: gcloud auth application-default login.2) In the browser window that opens, authenticate using your personal credentials.
D. 1) Store the service account's key file in JSON format in Hashicorp Vault.2) Integrate Terraform with Vault to retrieve the key file dynamically, and authenticate to Vault using a short-lived access token.

Answer: D

Explanation:
https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#file-system Whenever possible, avoid storing service account keys on a file system. If you can't avoid storing keys on
disk, make sure to restrict access to the key file, configure file access auditing, and encrypt the underlying disk.
https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#software-keystore In situations where using a hardware-based key store isn't

viable, use a software-based key store to manage
service account keys. Similar to hardware-based options, a software-based key store lets users or applications
use service account keys without revealing the private key. Software-based key store solutions can help you control key access in a fine-grained manner and can
also ensure that each key access is logged.

NEW QUESTION 80
- (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 81
- (Exam Topic 2)
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 repositor
B. Create a Pub/Sub trigger that runs the build file when an event is published to the topic.
C. Create a build trigger that runs the build file in response to a repository code being pushed to the development branch.
D. Create a webhook build trigger that runs the build file in response to HTTP POST calls to the webhook URL.
E. Create a Cron job that runs the following command every 24 hours: gcloud builds submit.

Answer: B

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.

NEW QUESTION 82
- (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 83
- (Exam Topic 2)
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.

Answer: BE

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

NEW QUESTION 88

- (Exam Topic 2)
You are running a web application on Google Kubernetes Engine that you inherited. You want to determine whether the application is using libraries with known vulnerabilities or is vulnerable to XSS attacks. Which service should you use?

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

Answer: C

Explanation:
https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview
Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to exercise as many user inputs and event handlers as possible.

NEW QUESTION 89
- (Exam Topic 2)
You plan to make a simple HTML application available on the internet. This site keeps information about FAQs for your application. The application is static and contains images, HTML, CSS, and Javascript. You want to make this application available on the internet with as few steps as possible.
What should you do?

A. Upload your application to Cloud Storage.
B. Upload your application to an App Engine environment.
C. Create a Compute Engine instance with Apache web server installe
D. Configure Apache web server to host the application.
E. Containerize your application firs
F. Deploy this container to Google Kubernetes Engine (GKE) and assignan external IP address to the GKE pod hosting the application.

Answer: A

Explanation:
Reference: https://cloud.google.com/storage/docs/hosting-static-website

NEW QUESTION 91
- (Exam Topic 2)
You manage an ecommerce application that processes purchases from customers who can subsequently cancel or change those purchases. You discover that order volumes are highly variable and the backend
order-processing system can only process one request at a time. You want to ensure seamless performance for customers regardless of usage volume. It is crucial that customers' order update requests are performed in the sequence in which they were generated. What should you do?

A. Send the purchase and change requests over WebSockets to the backend.
B. Send the purchase and change requests as REST requests to the backend.
C. Use a Pub/Sub subscriber in pull mode and use a data store to manage ordering.
D. Use a Pub/Sub subscriber in push mode and use a data store to manage ordering.

Answer: C

Explanation:
https://cloud.google.com/pubsub/docs/pull

NEW QUESTION 92
- (Exam Topic 2)
You are working on a social media application. You plan to add a feature that allows users to upload images. These images will be 2 MB – 1 GB in size. You want to minimize their infrastructure operations overhead for this feature. What should you do?

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

Answer: B

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

NEW QUESTION 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 98
- (Exam Topic 2)
You are developing an application that will store and access sensitive unstructured data objects in a Cloud Storage bucket. To comply with regulatory requirements, you need to ensure that all data objects are available for at least 7 years after their initial creation. Objects created more than 3 years ago are accessed very infrequently (less than once a year). You need to configure object storage while ensuring that storage cost is optimized. What should you do? (Choose two.)

A. Set a retention policy on the bucket with a period of 7 years.
B. Use IAM Conditions to provide access to objects 7 years after the object creation date.
C. Enable Object Versioning to prevent objects from being accidentally deleted for 7 years after object creation.
D. Create an object lifecycle policy on the bucket that moves objects from Standard Storage to Archive Storage after 3 years.
E. Implement a Cloud Function that checks the age of each object in the bucket and moves the objects older than 3 years to a second bucket with the Archive Storage clas
F. Use Cloud Scheduler to trigger the Cloud Function on a daily schedule.

Answer: AD

Explanation:
https://cloud.google.com/storage/docs/bucket-lock
This page discusses the Bucket Lock feature, which allows you to configure a data retention policy for a Cloud Storage bucket that governs how long objects in the bucket must be retained. The feature also allows you to lock the data retention policy, permanently preventing the policy from being reduced or removed.
https://cloud.google.com/storage/docs/storage-classes#archive
Archive storage is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Unlike the "coldest" storage services offered by other Cloud providers, your data is available within milliseconds, not hours or days.
Archive storage is the best choice for data that you plan to access less than once a year.

NEW QUESTION 102
- (Exam Topic 2)
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 Centra
B. Create a DNS CNAME record to point to the App Engine canonical name ghs.googlehosted.com.
C. Verify domain ownership with Webmaster Centra
D. Define an A record pointing to the single global App Engine IP address.
E. 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.
F. 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.

Answer: A

Explanation:
Reference: https://cloud.google.com/appengine/docs/flexible/dotnet/mapping-custom-domains?hl=fa

NEW QUESTION 104
- (Exam Topic 2)
You are developing an internal application that will allow employees to organize community events within your company. You deployed your application on a single Compute Engine instance. Your company uses Google Workspace (formerly G Suite), and you need to ensure that the company employees can authenticate to the application from anywhere. What should you do?

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

Answer: B

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

NEW QUESTION 107
- (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)
You are load testing your server application. During the first 30 seconds, you observe that a previously inactive Cloud Storage bucket is now servicing 2000 write requests per second and 7500 read requests per second. Your application is now receiving intermittent 5xx and 429 HTTP responses from the Cloud Storage JSON API as the demand escalates. You want to decrease the failed responses from the Cloud Storage API. What should you do?

A. Distribute the uploads across a large number of individual storage buckets.
B. Use the XML API instead of the JSON API for interfacing with Cloud Storage.
C. Pass the HTTP response codes back to clients that are invoking the uploads from your application.
D. Limit the upload rate from your application clients so that the dormant bucket's peak request rate is reached more gradually.

Answer: A

Explanation:
Reference: https://cloud.google.com/storage/docs/request-rate

NEW QUESTION 112
- (Exam Topic 2)
You are developing a web application that will be accessible over both HTTP and HTTPS and will run on Compute Engine instances. On occasion, you will need to SSH from your remote laptop into one of the Compute Engine instances to conduct maintenance on the app. How should you configure the instances while following Google-recommended best practices?

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

Answer: C

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

NEW QUESTION 115
- (Exam Topic 2)
Your team develops services that run on Google Kubernetes Engine. Your team's code is stored in Cloud Source Repositories. You need to quickly identify bugs in the code before it is deployed to production. You want to invest in automation to improve developer feedback and make the process as efficient as possible. What should you do?

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

Answer: A

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

NEW QUESTION 117
- (Exam Topic 2)
You are configuring a continuous integration pipeline using Cloud Build to automate the deployment of new container images to Google Kubernetes Engine (GKE). The pipeline builds the application from its source code, runs unit and integration tests in separate steps, and pushes the container to Container Registry. The application runs on a Python web server.
The Dockerfile is as follows: FROM python:3.7-alpine - COPY . /app WORKDIR /app
RUN pip install -r requirements.txt CMD ["gunicorn", "-w 4", "main:app"]
You notice that Cloud Build runs are taking longer than expected to complete. You want to decrease the build time. What should you do? (Choose two.)

A. Select a virtual machine (VM) size with higher CPU for Cloud Build runs.
B. Deploy a Container Registry on a Compute Engine VM in a VPC, and use it to store the final images.
C. Cache the Docker image for subsequent builds using the -- cache-from argument in your build config file.
D. Change the base image in the Dockerfile to ubuntu:latest, and install Python 3.7 using a package manager utility.
E. Store application source code on Cloud Storage, and configure the pipeline to use gsutil to download the source code.

Answer: AC

Explanation:
https://cloud.google.com/build/docs/optimize-builds/increase-vcpu-for-builds
By default, Cloud Build runs your builds on a standard virtual machine (VM). In addition to the standard VM, Cloud Build provides several high-CPU VM types to run builds. To increase the speed of your build, select a machine with a higher vCPU to run builds. Keep in mind that although selecting a high vCPU machine increases your build speed, it may also increase the startup time of your build as Cloud Build only starts non-standard machines on demand.

https://cloud.google.com/build/docs/optimize-builds/speeding-up-builds#using_a_cached_docker_image
The easiest way to increase the speed of your Docker image build is by specifying a cached image that can be used for subsequent builds. You can specify the cached image by adding the --cache-from argument in your build config file, which will instruct Docker to build using that image as a cache source.

NEW QUESTION 119
- (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 122
- (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 123
- (Exam Topic 2)
Your App Engine standard configuration is as follows: service: production
instance_class: B1
You want to limit the application to 5 instances. Which code snippet should you include in your configuration?

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

Answer: C

NEW QUESTION 124
- (Exam Topic 2)
You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. Your application can scale horizontally, and each instance of your application needs to have a stable network identity and its own persistent disk.
Which GKE object should you use?

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

Answer: B

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

NEW QUESTION 127
- (Exam Topic 2)
You recently joined a new team that has a Cloud Spanner database instance running in production. Your manager has asked you to optimize the Spanner instance to reduce cost while maintaining high reliability and availability of the database. What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 132
- (Exam Topic 2)
You plan to deploy a new application revision with a Deployment resource to Google Kubernetes Engine (GKE) in production. The container might not work correctly. You want to minimize risk in case there are issues after deploying the revision. You want to follow Google-recommended best practices. What should you do?

A. Perform a rolling update with a PodDisruptionBudget of 80%.
B. Perform a rolling update with a HorizontalPodAutoscaler scale-down policy value of 0.
C. Convert the Deployment to a StatefulSet, and perform a rolling update with a PodDisruptionBudget of 80%.
D. Convert the Deployment to a StatefulSet, and perform a rolling update with a HorizontalPodAutoscaler scale-down policy value of 0.

Answer: A

Explanation:
https://cloud.google.com/blog/products/containers-kubernetes/ensuring-reliability-and-uptime-for-your-gke-clus Setting PodDisruptionBudget ensures that your workloads have a sufficient number of replicas, even during maintenance. Using the PDB, you can define a number (or percentage) of pods that can be terminated, even if terminating them brings the current replica count below the desired value. With PDB configured, Kubernetes will drain a node following the configured disruption schedule. New pods will be deployed on other available nodes. This approach ensures Kubernetes schedules workloads in an optimal way while controlling the disruption based on the PDB configuration.
https://blog.knoldus.com/how-to-avoid-outages-in-your-kubernetes-cluster-using-pdb/

NEW QUESTION 133
- (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 135
- (Exam Topic 2)
You are planning to deploy hundreds of microservices in your Google Kubernetes Engine (GKE) cluster. How should you secure communication between the microservices on GKE using a managed service?

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

Answer: D

Explanation:
https://cloud.google.com/service-mesh/docs/overview#security_benefits
- Ensures encryption in transit. Using mTLS for authentication also ensures that all TCP communications are encrypted in transit.

NEW QUESTION 138
- (Exam Topic 2)
You are deploying your applications on Compute Engine. One of your Compute Engine instances failed to launch. What should you do? (Choose two.)

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

Answer: AD

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

NEW QUESTION 140
- (Exam Topic 2)
Your application is running in multiple Google Kubernetes Engine clusters. It is managed by a Deployment in each cluster. The Deployment has created multiple replicas of your Pod in each cluster. You want to view the logs sent to stdout for all of the replicas in your Deployment in all clusters. Which command should you use?

A. kubectl logs [PARAM]
B. gcloud logging read [PARAM]
C. kubectl exec –it [PARAM] journalctl
D. gcloud compute ssh [PARAM] —-command= "sudo journalctl"

Answer: D

NEW QUESTION 145
- (Exam Topic 2)
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.

Answer: D

NEW QUESTION 146
- (Exam Topic 2)
You are a SaaS provider deploying dedicated blogging software to customers in your Google Kubernetes Engine (GKE) cluster. You want to configure a secure multi-tenant platform to ensure that each customer has access to only their own blog and can't affect the workloads of other customers. What should you do?

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

Answer: B

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

NEW QUESTION 148
- (Exam Topic 2)
You are building an API that will be used by Android and iOS apps The API must:
• Support HTTPs
• Minimize bandwidth cost
• Integrate easily with mobile apps Which API architecture should you use?

A. RESTful APIs
B. MQTT for APIs
C. gRPC-based APIs
D. SOAP-based APIs

Answer: A

Explanation:
Reference: https://www.devteam.space/blog/how-to-build-restful-api-for-your-mobile-app/

NEW QUESTION 149
- (Exam Topic 2)
Your operations team has asked you to create a script that lists the Cloud Bigtable, Memorystore, and Cloud SQL databases running within a project. The script should allow users to submit a filter expression to limit the results presented. How should you retrieve the data?

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

Answer: D

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

NEW QUESTION 152
- (Exam Topic 2)
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 projec

E. Create Compute engine instances in a different project.Create a VPN between these two projects to allow internal access to CloudSQL.

Answer: C

Explanation:
Reference: https://cloud.google.com/sql/docs/mysql/connect-external-app

NEW QUESTION 154
- (Exam Topic 2)
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.

Answer: A

Explanation:
https://cloud.google.com/build/docs/automating-builds/create-manage-triggers#:~:text=A%20Cloud%20Build%

NEW QUESTION 158
- (Exam Topic 2)
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

Answer: B

Explanation:
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.

NEW QUESTION 163
- (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 Servic
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 Servic
D. Clients should use the service name in the URL to connect to the service.
E. Define a GKE Endpoin
F. Clients should get the endpoint name from the appropriate environment variable in the client container.
G. Define a GKE Endpoin
H. Clients should get the endpoint name from Cloud DNS.

Answer: C

NEW QUESTION 165
- (Exam Topic 2)
You have containerized a legacy application that stores its configuration on an NFS share. You need to deploy this application to Google Kubernetes Engine (GKE) and do not want the application serving traffic until after the configuration has been retrieved. What should you do?

A. Use the gsutil utility to copy files from within the Docker container at startup, and start the service using an ENTRYPOINT script.
B. Create a PersistentVolumeClaim on the GKE cluste
C. Access the configuration files from the volume, and start the service using an ENTRYPOINT script.
D. Use the COPY statement in the Dockerfile to load the configuration into the container imag
E. Verify that the configuration is available, and start the service using an ENTRYPOINT script.
F. Add a startup script to the GKE instance group to mount the NFS share at node startu
G. Copy the configuration files into the container, and start the service using an ENTRYPOINT script.

Answer: D

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

NEW QUESTION 170

- (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 175
- (Exam Topic 2)
You are developing an application that will allow users to read and post comments on news articles. You want to configure your application to store and display user-submitted comments using Firestore. How should you design the schema to support an unknown number of comments and articles?

A. Store each comment in a subcollection of the article.
B. Add each comment to an array property on the article.
C. Store each comment in a document, and add the comment's key to an array property on the article.
D. Store each comment in a document, and add the comment's key to an array property on the user profile.

Answer: D

NEW QUESTION 178
- (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 179
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

Professional-Cloud-Developer Practice Exam Features:

* Professional-Cloud-Developer Questions and Answers Updated Frequently

* Professional-Cloud-Developer Practice Questions Verified by Expert Senior Certified Staff

* Professional-Cloud-Developer Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Professional-Cloud-Developer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year