Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

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

NEW QUESTION 1
- (Exam Topic 1)
For this question, refer to the HipLocal case study.
Which Google Cloud product addresses HipLocal's business requirements for service level indicators and objectives?

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

Answer: B

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

NEW QUESTION 2
- (Exam Topic 1)
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 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 is configuring their access controls.
Which firewall configuration should they implement?

A. Block all traffic on port 443.
B. Allow all traffic into the network.
C. Allow traffic on port 443 for a specific tag.
D. Allow all traffic on port 443 into the network.

Answer: D

NEW QUESTION 5
- (Exam Topic 1)
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 6
- (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 7
- (Exam Topic 2)
You support an application that uses the Cloud Storage API. You review the logs and discover multiple HTTP 503 Service Unavailable error responses from the API. Your application logs the error and does not take any further action. You want to implement Google-recommended retry logic to improve success rates. Which approach should you take?

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

Answer: C

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

NEW QUESTION 8
- (Exam Topic 2)
Your company needs a database solution that stores customer purchase history and meets the following requirements:
Customers can query their purchase immediately after submission. Purchases can be sorted on a variety of fields.
Distinct record formats can be stored at the same time. Which storage option satisfies these requirements?

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

Answer: A

NEW QUESTION 9
- (Exam Topic 2)
You are developing an application hosted on Google Cloud that uses a MySQL relational database schema. The application will have a large volume of reads and writes to the database and will require backups and ongoing capacity planning. Your team does not have time to fully manage the database but can take on small administrative tasks. How should you host the database?

A. Configure Cloud SQL to host the database, and import the schema into Cloud SQL.
B. Deploy MySQL from the Google Cloud Marketplace to the database using a client, and import the schema.
C. Configure Bigtable to host the database, and import the data into Bigtable.
D. Configure Cloud Spanner to host the database, and import the schema into Cloud Spanner.
E. Configure Firestore to host the database, and import the data into Firestore.

Answer: A

Explanation:
https://cloud.google.com/spanner/docs/migrating-mysql-to-spanner#migration-process
Cloud SQL: Cloud SQL is a web service that allows you to create, configure, and use relational databases that live in Google's cloud. It is a fully-managed service that maintains, manages, and administers your databases, allowing you to focus on your applications and services.
https://cloud.google.com/sql/docs/mysql Cloud SQL for MySQL is a fully-managed database service that helps you set up, maintain, manage, and administer your MySQL relational databases on Google Cloud Platform.

NEW QUESTION 10
- (Exam Topic 2)
You are responsible for deploying a new API. That API will have three different URL paths:
• https://yourcompany.com/students
• https://yourcompany.com/teachers
• https:/yourcompany.com/classes
You need to configure each API URL path to invoke a different function in your code. What should you do?

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

Answer: D

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

NEW QUESTION 10
- (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 12
- (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 14
- (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 17
- (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 21
- (Exam Topic 2)
You have an application controlled by a managed instance group. When you deploy a new version of the application, costs should be minimized and the number of instances should not increase. You want to ensure that, when each new instance is created, the deployment only continues if the new instance is healthy. What should you do?

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

Answer: A

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

NEW QUESTION 25
- (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 30
- (Exam Topic 2)
Your application takes an input from a user and publishes it to the user's contacts. This input is stored in a table in Cloud Spanner. Your application is more sensitive to latency and less sensitive to consistency. How should you perform reads from Cloud Spanner for this application?

A. Perform Read-Only transactions.
B. Perform stale reads using single-read methods.
C. Perform strong reads using single-read methods.
D. Perform stale reads using read-write transactions.

Answer: D

Explanation:
Reference: https://cloud.google.com/solutions/best-practices-cloud-spanner-gaming-database

NEW QUESTION 34
- (Exam Topic 2)
You are writing a single-page web application with a user-interface that communicates with a third-party API for content using XMLHttpRequest. The data displayed on the UI by the API results is less critical than other data displayed on the same web page, so it is acceptable for some requests to not have the API data
displayed in the UI. However, calls made to the API should not delay rendering of other parts of the user interface. You want your application to perform well when the API response is an error or a timeout. What should you do?

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

Answer: A

NEW QUESTION 39
- (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 44
- (Exam Topic 2)
Your data is stored in Cloud Storage buckets. Fellow developers have reported that data downloaded from Cloud Storage is resulting in slow API performance. You want to research the issue to provide details to the GCP support team. Which command should you run?

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

Answer: B

Explanation:
Reference: https://groups.google.com/forum/#!topic/gce-discussion/xBl9Jq5HDsY

NEW QUESTION 49
- (Exam Topic 2)
Your company is planning to migrate their on-premises Hadoop environment to the cloud. Increasing storage cost and maintenance of data stored in HDFS is a major concern for your company. You also want to make minimal changes to existing data analytics jobs and existing architecture. How should you proceed with the migration?

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

Answer: D

NEW QUESTION 54
- (Exam Topic 2)
You work for a web development team at a small startup. Your team is developing a Node.js application using Google Cloud services, including Cloud Storage and Cloud Build. The team uses a Git repository for version control. Your manager calls you over the weekend and instructs you to make an emergency update to one of the company's websites, and you're the only developer available. You need to access Google Cloud to make the update, but you don't have your work laptop. You are not allowed to store source code locally on a non-corporate computer. How should you set up your developer environment?

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

Answer: C

Explanation:
https://cloud.google.com/shell/docs

NEW QUESTION 56
- (Exam Topic 2)
You are developing a new public-facing application that needs to retrieve specific properties in the metadata of users' objects in their respective Cloud Storage buckets. Due to privacy and data residency requirements, you must retrieve only the metadata and not the object data. You want to maximize the performance of the retrieval process. How should you retrieve the metadata?

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

Answer: D

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

NEW QUESTION 60
- (Exam Topic 2)
You are developing an application that needs to store files belonging to users in Cloud Storage. You want each user to have their own subdirectory in Cloud Storage. When a new user is created, the corresponding empty subdirectory should also be created. What should you do?

A. Create an object with the name of the subdirectory ending with a trailing slash ('/') that is zero bytes in length.
B. Create an object with the name of the subdirectory, and then immediately delete the object within that subdirectory.
C. Create an object with the name of the subdirectory that is zero bytes in length and has WRITER access control list permission.
D. Create an object with the name of the subdirectory that is zero bytes in lengt
E. Set the Content-Type metadata to CLOUDSTORAGE_FOLDER.

Answer: A

Explanation:
https://cloud.google.com/storage/docs/folders
If you create an empty folder using the Google Cloud console, Cloud Storage creates a zero-byte object as a placeholder. For example, if you create a folder called folder in a bucket called my-bucket, a zero- byte object called gs://my-bucket/folder/ is created. This placeholder is discoverable by other tools when listing the objects in the bucket, for example when using the gsutil ls command.

NEW QUESTION 64
- (Exam Topic 2)
You are designing a schema for a table that will be moved from MySQL to Cloud Bigtable. The MySQL table is as follows:

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

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

A. Set Account_id as a key.
B. Set Account_id_Event_timestamp as a key.
C. Set Event_timestamp_Account_id as a key.
D. Set Event_timestamp as a key.

Answer: C

NEW QUESTION 66
- (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 71
- (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 73
- (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 77
- (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 80
- (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 81
- (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 84
- (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 85
- (Exam Topic 2)
You are building a mobile application that will store hierarchical data structures in a database. The application will enable users working offline to sync changes when they are back online. A backend service will enrich the data in the database using a service account. The application is expected to be very popular and needs to scale seamlessly and securely. Which database and IAM role should you use?

A. Use Cloud SQL, and assign the roles/cloudsql.editor role to the service account.
B. Use Bigtable, and assign the roles/bigtable.viewer role to the service account.
C. Use Firestore in Native mode and assign the roles/datastore.user role to the service account.
D. Use Firestore in Datastore mode and assign the roles/datastore.viewer role to the service account.

Answer: C

Explanation:
https://firebase.google.com/docs/firestore/manage-data/enable-offline
Cloud Firestore supports offline data persistence. This feature caches a copy of the Cloud Firestore data that your app is actively using, so your app can access the data when the device is offline. You can write, read, listen to, and query the cached data. When the device comes back online, Cloud Firestore synchronizes any local changes made by your app to the Cloud Firestore backend.

NEW QUESTION 88
- (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 89
- (Exam Topic 2)
You recently developed a new service on Cloud Run. The new service authenticates using a custom service and then writes transactional information to a Cloud Spanner database. You need to verify that your application can support up to 5,000 read and 1,000 write transactions per second while identifying any bottlenecks that occur. Your test infrastructure must be able to autoscale. What should you do?

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

Answer: B

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

NEW QUESTION 94
- (Exam Topic 2)
Your application is deployed in a Google Kubernetes Engine (GKE) cluster. When a new version of your application is released, your CI/CD tool updates the spec.template.spec.containers[0].image value to reference the Docker image of your new application version. When the Deployment object applies the change, you want to deploy at least 1 replica of the new version and maintain the previous replicas until the new replica is healthy.
Which change should you make to the GKE Deployment object shown below?

```
apiVersion: apps/v1
kind: Deployment
metadata:
    name: ecommerce-frontend-deployment
spec:
    replicas: 3
    selector:
        matchLabels:
            app: ecommerce-frontend
    template:
        metadata:
            labels:
                app: ecommerce-frontend
        spec:
            containers:
            - name: ecommerce-frontend-webapp
                image: ecommerce-frontend-webapp:1.7.9
                ports:
                - containerPort: 80
```

A. Set the Deployment strategy to RollingUpdate with maxSurge set to 0, maxUnavailable set to 1.
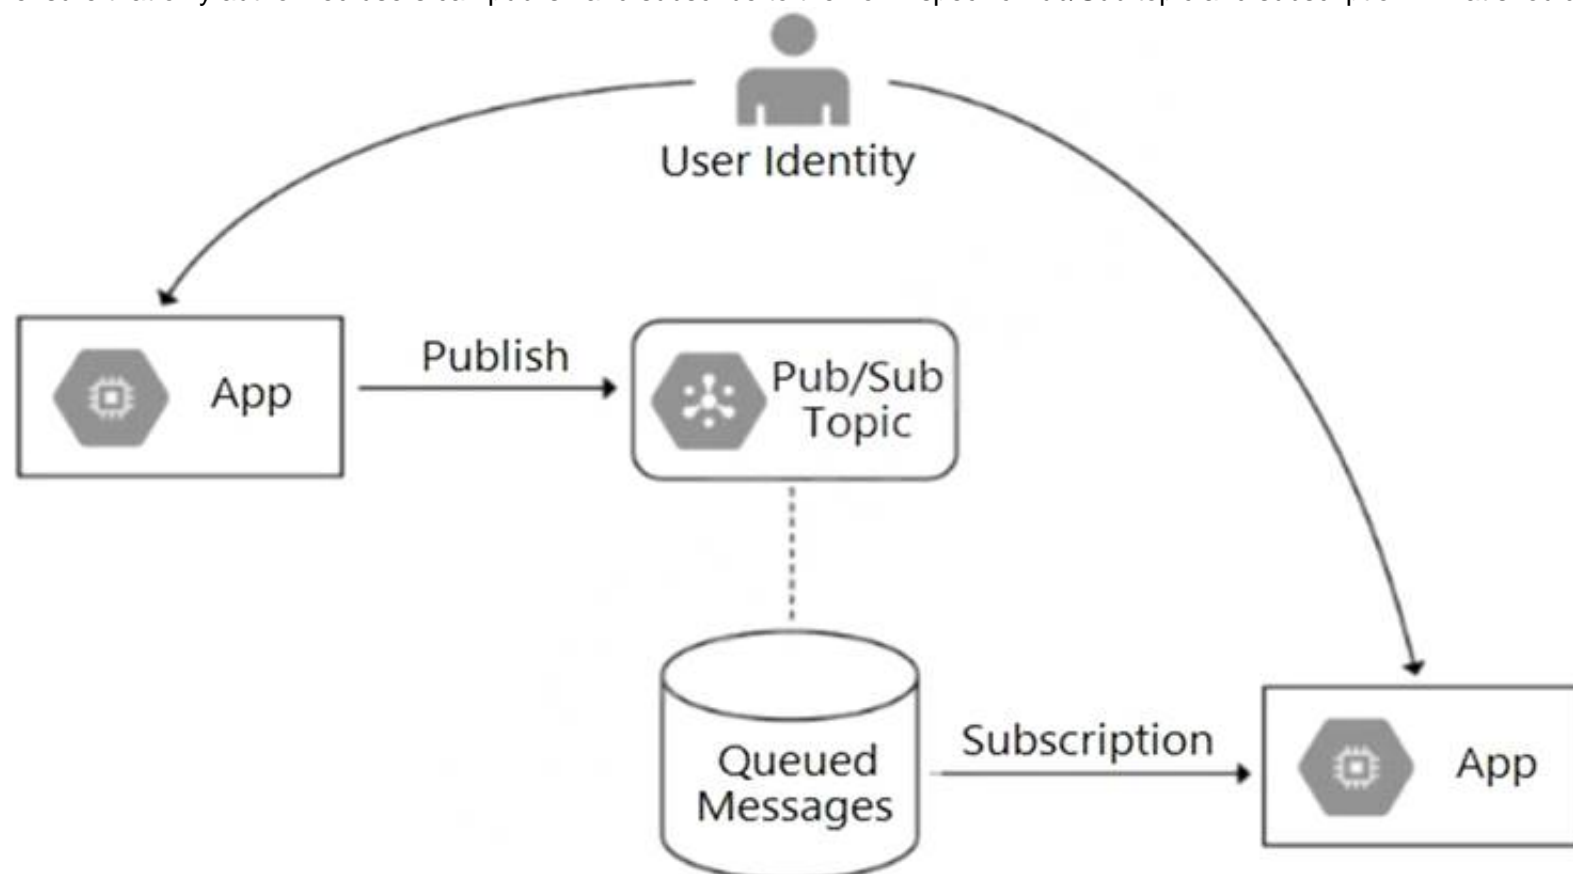B. Set the Deployment strategy to RollingUpdate with maxSurge set to 1, maxUnavailable set to 0.

C. Set the Deployment strategy to Recreate with maxSurge set to 0, maxUnavailable set to 1.
D. Set the Deployment strategy to Recreate with maxSurge set to 1, maxUnavailable set to 0.

Answer: D

NEW QUESTION 98
- (Exam Topic 2)
Your team is developing an application in Google Cloud that executes with user identities maintained by Cloud Identity. Each of your application's users will have an associated Pub/Sub topic to which messages are published, and a Pub/Sub subscription where the same user will retrieve published messages. You need to ensure that only authorized users can publish and subscribe to their own specific Pub/Sub topic and subscription. What should you do?

A. Bind the user identity to the pubsub.publisher and pubsub.subscriber roles at the resource level.
B. Grant the user identity the pubsub.publisher and pubsub.subscriber roles at the project level.
C. Grant the user identity a custom role that contains the pubsub.topics.create and pubsub.subscriptions.create permissions.
D. Configure the application to run as a service account that has the pubsub.publisher and pubsub.subscriber roles.

Answer: C

NEW QUESTION 99
- (Exam Topic 2)
Your team is writing a backend application to implement the business logic for an interactive voice response (IVR) system that will support a payroll application.
The IVR system has the following technical characteristics:
• Each customer phone call is associated with a unique IVR session.
• The IVR system creates a separate persistent gRPC connection to the backend for each session.
• If the connection is interrupted, the IVR system establishes a new connection, causing a slight latency for that call.
You need to determine which compute environment should be used to deploy the backend application. Using current call data, you determine that:
• Call duration ranges from 1 to 30 minutes.
• Calls are typically made during business hours.
• There are significant spikes of calls around certain known dates (e.g., pay days), or when large payroll changes occur.
You want to minimize cost, effort, and operational overhead. Where should you deploy the backend application?

A. Compute Engine
B. Google Kubernetes Engine cluster in Standard mode
C. Cloud Functions
D. Cloud Run

Answer: D

Explanation:
This page shows Cloud Run-specific details for developers who want to use gRPC to connect a Cloud Run service with other services, for example, to provide simple, high performance communication between internal microservices. You can use all gRPC types, streaming or unary, with Cloud Run.
Possible use cases include:
Communication between internal microservices.
High loads of data (gRPC uses protocol buffers, which are up to seven times faster than REST calls). Only a simple service definition is needed, you don't want to write a full client library.
Use streaming gRPCs in your gRPC server to build more responsive applications and APIs. https://cloud.google.com/run/docs/tutorials/secure-services#:~:text=The%20backend%20service%20is%20priva

NEW QUESTION 102
- (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 Ru
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 Mes
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 Engin
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 Engin
H. Use Traffic Director to direct a subset of user traffic to the new version based on predefined weights.

Answer: B

NEW QUESTION 106
- (Exam Topic 2)
Your web application is deployed to the corporate intranet. You need to migrate the web application to Google Cloud. The web application must be available only to company employees and accessible to employees as they travel. You need to ensure the security and accessibility of the web application while minimizing application changes. What should you do?

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

Answer: B

NEW QUESTION 109
- (Exam Topic 2)
You are designing an application that consists of several microservices. Each microservice has its own RESTful API and will be deployed as a separate Kubernetes Service. You want to ensure that the consumers of these APIs aren't impacted when there is a change to your API, and also ensure that third-party systems aren't interrupted when new versions of the API are released. How should you configure the connection to the application following Google-recommended best practices?

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

Answer: A

NEW QUESTION 110
- (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 Registr
D. Review the vulnerability results before each deployment.
E. Enable Container Analysis, and upload new container images to Artifact Registr
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 114
- (Exam Topic 2)
You need to configure a Deployment on Google Kubernetes Engine (GKE). You want to include a check that verifies that the containers can connect to the database. If the Pod is failing to connect, you want a script on the container to run to complete a graceful shutdown. How should you configure the Deployment?

A. Create two jobs: one that checks whether the container can connect to the database, and another that runs the shutdown script if the Pod is failing.
B. Create the Deployment with a livenessProbe for the container that will fail if the container can't connect to the databas
C. Configure a Prestop lifecycle handler that runs the shutdown script if the container is failing.
D. Create the Deployment with a PostStart lifecycle handler that checks the service availabilit
E. Configure a PreStop lifecycle handler that runs the shutdown script if the container is failing.
F. Create the Deployment with an initContainer that checks the service availabilit
G. Configure a Prestop lifecycle handler that runs the shutdown script if the Pod is failing.

Answer: B

Explanation:
https://cloud.google.com/architecture/best-practices-for-running-cost-effective-kubernetes-applications-on-gke#

NEW QUESTION 115
- (Exam Topic 2)
Your team is building an application for a financial institution. The application's frontend runs on Compute Engine, and the data resides in Cloud SQL and one Cloud Storage bucket. The application will collect data containing PII, which will be stored in the Cloud SQL database and the Cloud Storage bucket. You need to secure the PII data. What should you do?

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

Answer: C

NEW QUESTION 117
- (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 118
- (Exam Topic 2)
Your teammate has asked you to review the code below, which is adding a credit to an account balance in Cloud Datastore. Which improvement should you suggest your teammate make?

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

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

Answer: A

NEW QUESTION 122
- (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 127
- (Exam Topic 2)
You are designing an application that will subscribe to and receive messages from a single Pub/Sub topic and insert corresponding rows into a database. Your application runs on Linux and leverages preemptible virtual machines to reduce costs. You need to create a shutdown script that will initiate a graceful shutdown. What should you do?

A. Write a shutdown script that uses inter-process signals to notify the application process to disconnect from the database.
B. Write a shutdown script that broadcasts a message to all signed-in users that the Compute Engine instance is going down and instructs them to save current work and sign out.
C. Write a shutdown script that writes a file in a location that is being polled by the application once every five minute
D. After the file is read, the application disconnects from the database.
E. Write a shutdown script that publishes a message to the Pub/Sub topic announcing that a shutdown is in progres
F. After the application reads the message, it disconnects from the database.

Answer: D

NEW QUESTION 132
- (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 133
- (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 134
- (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 136
- (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 137
- (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 departmen
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 departmen
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%20securit "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 140
- (Exam Topic 2)
You need to redesign the ingestion of audit events from your authentication service to allow it to handle a large increase in traffic. Currently, the audit service and the authentication system run in the same Compute Engine virtual machine. You plan to use the following Google Cloud tools in the new architecture:
Multiple Compute Engine machines, each running an instance of the authentication service Multiple Compute Engine machines, each running an instance of the audit service
Pub/Sub to send the events from the authentication services.
How should you set up the topics and subscriptions to ensure that the system can handle a large volume of messages and can scale efficiently?

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

Answer: A

Explanation:
https://cloud.google.com/pubsub/docs/subscriber "Multiple subscribers can make pull calls to the same "shared" subscription. Each subscriber will receive a subset of the messages."

NEW QUESTION 142
- (Exam Topic 2)
You are using Cloud Build build to promote a Docker image to Development, Test, and Production environments. You need to ensure that the same Docker image is deployed to each of these environments. How should you identify the Docker image in your build?

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

Answer: D

NEW QUESTION 146
- (Exam Topic 2)
You are in the final stage of migrating an on-premises data center to Google Cloud. You are quickly approaching your deadline, and discover that a web API is running on a server slated for decommissioning. You need to recommend a solution to modernize this API while migrating to Google Cloud. The modernized web API must meet the following requirements:
• Autoscales during high traffic periods at the end of each month
• Written in Python 3.x
• Developers must be able to rapidly deploy new versions in response to frequent code changes
You want to minimize cost, effort, and operational overhead of this migration. What should you do?

A. Modernize and deploy the code on App Engine flexible environment.
B. Modernize and deploy the code on App Engine standard environment.
C. Deploy the modernized application to an n1-standard-1 Compute Engine instance.
D. Ask the development team to re-write the application to run as a Docker container on Google Kubernetes Engine.

Answer: B

Explanation:
https://cloud.google.com/appengine/docs/standard

NEW QUESTION 148
- (Exam Topic 2)
You are running a containerized application on Google Kubernetes Engine. Your container images are stored in Container Registry. Your team uses CI/CD practices. You need to prevent the deployment of containers with known critical vulnerabilities. What should you do?

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

Answer: D

Explanation:
https://cloud.google.com/binary-authorization/docs/creating-attestations-kritis
https://cloud.google.com/container-analysis/docs/os-overview

NEW QUESTION 152
- (Exam Topic 2)
You want to view the memory usage of your application deployed on Compute Engine. What should you do?

A. Install the Stackdriver Client Library.
B. Install the Stackdriver Monitoring Agent.
C. Use the Stackdriver Metrics Explorer.
D. Use the Google Cloud Platform Console.

Answer: C

Explanation:
Reference:
https://stackoverflow.com/questions/43991246/google-cloud-platform-how-to-monitor-memory-usage-of-vm-in

NEW QUESTION 154
- (Exam Topic 2)
You are developing a web application that contains private images and videos stored in a Cloud Storage bucket. Your users are anonymous and do not have Google Accounts. You want to use your
application-specific logic to control access to the images and videos. How should you configure access?

A. Cache each web application user's IP address to create a named IP table using Google Cloud Armor.Create a Google Cloud Armor security policy that allows users to access the backend bucket.
B. Grant the Storage Object Viewer IAM role to allUser
C. Allow users to access the bucket after authenticating through your web application.
D. Configure Identity-Aware Proxy (IAP) to authenticate users into the web applicatio
E. Allow users to access the bucket after authenticating through IAP.
F. Generate a signed URL that grants read access to the bucke
G. Allow users to access the URL after authenticating through your web application.

Answer: D

Explanation:
https://cloud.google.com/storage/docs/access-control/signed-urls#should-you-use
In some scenarios, you might not want to require your users to have a Google account in order to access Cloud Storage, but you still want to control access using your application-specific logic. The typical way to address this use case is to provide a signed URL to a user, which gives the user read, write, or delete access to that resource for a limited time. You specify an expiration time when you create the signed URL. Anyone who knows the URL can access the resource until the expiration time for the URL is reached or the key used to sign the URL is rotated.

NEW QUESTION 158
- (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 162
- (Exam Topic 2)
Your company uses Cloud Logging to manage large volumes of log data. You need to build a real-time log analysis architecture that pushes logs to a third-party application for processing. What should you do?

A. Create a Cloud Logging log export to Pub/Sub.
B. Create a Cloud Logging log export to BigQuery.
C. Create a Cloud Logging log export to Cloud Storage.
D. Create a Cloud Function to read Cloud Logging log entries and send them to the third-party application.

Answer: B

NEW QUESTION 167
- (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 emulato
B. Generate random content in the publishing service, and publish to the emulator.
C. Create a service to publish messages to your applicatio
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 emulato
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 emulato
H. Publish a standard set of testing messages from the publishing service to the emulator.

Answer: D

NEW QUESTION 168
- (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 169
- (Exam Topic 2)
You recently migrated an on-premises monolithic application to a microservices application on Google Kubernetes Engine (GKE). The application has dependencies on backend services on-premises, including a CRM system and a MySQL database that contains personally identifiable information (PII). The backend services must remain on-premises to meet regulatory requirements.
You established a Cloud VPN connection between your on-premises data center and Google Cloud. You notice that some requests from your microservices application on GKE to the backend services are failing due to latency issues caused by fluctuating bandwidth, which is causing the application to crash. How should you address the latency issues?

A. Use Memorystore to cache frequently accessed PII data from the on-premises MySQL database
B. Use Istio to create a service mesh that includes the microservices on GKE and the on-premises services
C. Increase the number of Cloud VPN tunnels for the connection between Google Cloud and the on-premises services

D. Decrease the network layer packet size by decreasing the Maximum Transmission Unit (MTU) value from its default value on Cloud VPN

Answer: C

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

NEW QUESTION 172
- (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 175
- (Exam Topic 2)
You have an application deployed in production. When a new version is deployed, some issues don't arise until the application receives traffic from users in production. You want to reduce both the impact and the number of users affected.
Which deployment strategy should you use?

A. Blue/green deployment
B. Canary deployment
C. Rolling deployment
D. Recreate deployment

Answer: A

Explanation:
Reference: https://thenewstack.io/deployment-strategies/

NEW QUESTION 178
- (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 180
- (Exam Topic 2)
You recently deployed a Go application on Google Kubernetes Engine (GKE). The operations team has noticed that the application's CPU usage is high even when there is low production traffic. The operations team has asked you to optimize your application's CPU resource consumption. You want to determine which Go functions consume the largest amount of CPU. What should you do?

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

Answer: D

Explanation:
https://cloud.google.com/profiler/docs/about-profiler
Cloud Profiler is a statistical, low-overhead profiler that continuously gathers CPU usage and
memory-allocation information from your production applications. It attributes that information to the source code that generated it, helping you identify the parts of

your application that are consuming the most resources, and otherwise illuminating your applications performance characteristics.
https://cloud.google.com/profiler/docs

NEW QUESTION 183
- (Exam Topic 2)
You are building a highly available and globally accessible application that will serve static content to users. You need to configure the storage and serving components. You want to minimize management overhead and latency while maximizing reliability for users. What should you do?

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

Answer: D

NEW QUESTION 185
- (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 189
- (Exam Topic 2)
Your company has a BigQuery data mart that provides analytics information to hundreds of employees. One user of 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 192
- (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.updateStrategv.type to RollingUpdat
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 194
- (Exam Topic 2)
You are creating a Google Kubernetes Engine (GKE) cluster and run this command:

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

The command fails with the error:

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

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

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

Answer: A

NEW QUESTION 199
- (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 201
- (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 202
- (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 DN
C. Clients should use the name of the A record to connect to the service.
D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/.
E. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[API_NAME]/[API_VERSION]/.

Answer: D

NEW QUESTION 203
- (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 205
- (Exam Topic 2)
You are developing a new application that has the following design requirements: Creation and changes to the application infrastructure are versioned and auditable.
The application and deployment infrastructure uses Google-managed services as much as possible. The application runs on a serverless compute platform. How should you design the application's architecture?

A. * 1. Store the application and infrastructure source code in a Git repository.* 2. Use Cloud Build to deploy the application infrastructure with Terraform.* 3. Deploy the application to a Cloud Function as a pipeline step.
B. * 1. Deploy Jenkins from the Google Cloud Marketplace, and define a continuous integration pipeline in Jenkins.* 2. Configure a pipeline step to pull the application source code from a Git repository.* 3. Deploy the application source code to App Engine as a pipeline step.
C. * 1. Create a continuous integration pipeline on Cloud Build, and configure the pipeline to deploy the application infrastructure using Deployment Manager templates.* 2. Configure a pipeline step to create a container with the latest application source code.* 3. Deploy the container to a Compute Engine instance as a pipeline step.
D. * 1. Deploy the application infrastructure using gcloud commands.* 2. Use Cloud Build to define a continuous integration pipeline for changes to the application source code.* 3. Configure a pipeline step to pull the application source code from a Git repository, and create a containerized application.* 4. Deploy the new container on Cloud Run as a pipeline step.

Answer: D

Explanation:
Reference: https://cloud.google.com/docs/ci-cd

NEW QUESTION 209
- (Exam Topic 2)
You have deployed an HTTP(s) Load Balancer with the gcloud commands shown below.

```
export NAME-load-balancer

# create network
gcloud compute networks create ${NAME}

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

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

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

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

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

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

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

Answer: C

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

NEW QUESTION 212
- (Exam Topic 2)
You recently deployed your application in Google Kubernetes Engine, and now need to release a new version of your application. You need the ability to instantly roll back to the previous version in case there are issues with the new version. Which deployment model should you use?

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

Answer: C

NEW QUESTION 216
- (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 218
- (Exam Topic 2)
You are designing a schema for a Cloud Spanner customer database. You want to store a phone number array field in a customer table. You also want to allow users to search customers by phone number. How should you design this schema?

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

Answer: C

NEW QUESTION 220
- (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 224
- (Exam Topic 2)
You are developing an ecommerce web application that uses App Engine standard environment and Memorystore for Redis. When a user logs into the app, the application caches the user's information (e.g., session, name, address, preferences), which is stored for quick retrieval during checkout.
While testing your application in a browser, you get a 502 Bad Gateway error. You have determined that the application is not connecting to Memorystore. What is the reason for this error?

A. Your Memorystore for Redis instance was deployed without a public IP address.
B. You configured your Serverless VPC Access connector in a different region than your App Engine instance.
C. The firewall rule allowing a connection between App Engine and Memorystore was removed during an infrastructure update by the DevOps team.
D. You configured your application to use a Serverless VPC Access connector on a different subnet in a different availability zone than your App Engine instance.

Answer: A

Explanation:
Reference: https://cloud.google.com/endpoints/docs/openapi/troubleshoot-response-errors

NEW QUESTION 227
- (Exam Topic 2)
Your application is deployed on hundreds of Compute Engine instances in a managed instance group (MIG) in multiple zones. You need to deploy a new instance template to fix a critical vulnerability immediately but must avoid impact to your service. What setting should be made to the MIG after updating the instance template?

A. Set the Max Surge to 100%.

B. Set the Update mode to Opportunistic.
C. Set the Maximum Unavailable to 100%.
D. Set the Minimum Wait time to 0 seconds.

Answer: B

Explanation:
https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups#type Alternatively, if an automated update is potentially too disruptive, you can choose to perform an opportunistic update. The MIG applies an opportunistic update only when you manually initiate the update on selected instances or when new instances are created. New instances can be created when you or another service, such as an autoscaler, resizes the MIG. Compute Engine does not actively initiate requests to apply opportunistic updates on existing instances.

NEW QUESTION 228
- (Exam Topic 2)
Your team develops services that run on Google Kubernetes Engine. You need to standardize their log data using Google-recommended practices and make the data more useful in the fewest number of steps. What should you do? (Choose two.)

A. Create aggregated exports on application logs to BigQuery to facilitate log analytics.
B. Create aggregated exports on application logs to Cloud Storage to facilitate log analytics.
C. Write log output to standard output (stdout) as single-line JSON to be ingested into Cloud Logging as structured logs.
D. Mandate the use of the Logging API in the application code to write structured logs to Cloud Logging.
E. Mandate the use of the Pub/Sub API to write structured data to Pub/Sub and create a Dataflow streaming pipeline to normalize logs and write them to BigQuery for analytics.

Answer: AC

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

NEW QUESTION 233
- (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 238
- (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 239
- (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 242
- (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.
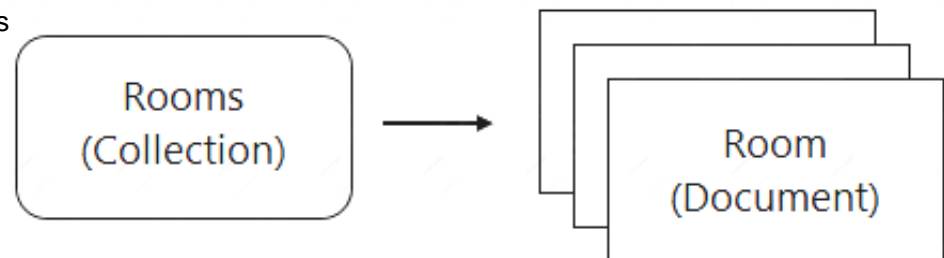
Answer: C

Explanation:
Reference: https://cloud.google.com/sql/docs/mysql/connect-external-app
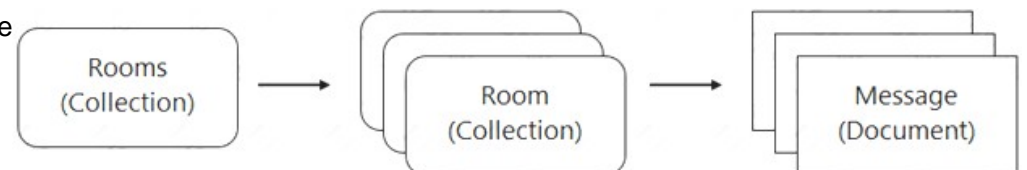
NEW QUESTION 245
- (Exam Topic 2)
You are designing a chat room application that will host multiple rooms and retain the message history for each room. You have selected Firestore as your
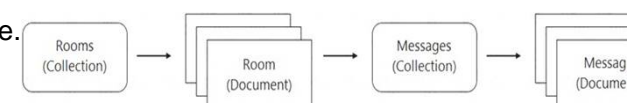database. How should you represent the data in Firestore?

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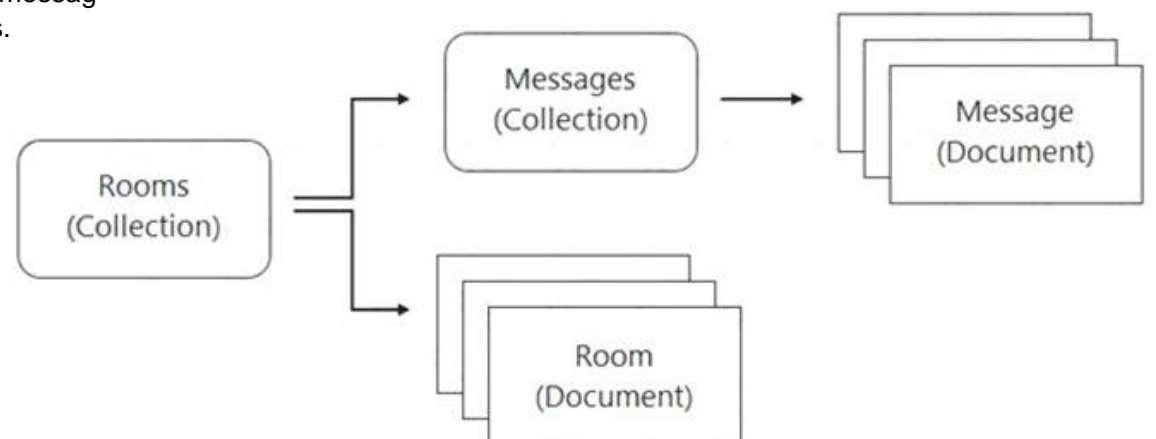B. For each room, create a document that lists the contents of the messages

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

E. Create a collection for the room
F. For each room, create a document that contains a collection for documents, each of which contains a message.

G. Create a collection for the rooms, and create a document for each roo
H. Create a separate collection for messages, with one document per messag
I. Each room's document contains a list of references to the messages.

Answer: C

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

NEW QUESTION 249
- (Exam Topic 2)
Your team is developing unit tests for Cloud Function code. The code is stored in a Cloud Source Repositories repository. You are responsible for implementing
the tests. Only a specific service account has the necessary permissions to deploy the code to Cloud Functions. You want to ensure that the code cannot be
deployed without first passing the tests. How should you configure the unit testing process?

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

G. Configure Cloud Build to run the unit tests, using the specific service account as the build agen
H. If the code passes the tests, Cloud Build deploys the Cloud Function.

Answer: D

NEW QUESTION 253
- (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 254
- (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 255
- (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 257
- (Exam Topic 2)
Your analytics system executes queries against a BigQuery dataset. The SQL query is executed in batch and passes the contents of a SQL file to the BigQuery CLI. Then it redirects the BigQuery CLI output to another process. However, you are getting a permission error from the BigQuery CLI when the queries are executed. You want to resolve the issue. What should you do?

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

Answer: B

NEW QUESTION 259
- (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 261
- (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 266
......

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