



Google

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

NEW QUESTION 1

- (Exam Topic 1)

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

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

Answer: B

NEW QUESTION 2

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

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

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

Answer: A

NEW QUESTION 3

- (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 Engine
- 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 4

- (Exam Topic 1)

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

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

Answer: D

NEW QUESTION 5

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 6

- (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 7

- (Exam Topic 1)

HipLocal's.net-based auth service fails under intermittent load. What should they do?

- A. Use App Engine for autoscaling.
- B. Use Cloud Functions for autoscaling.
- C. Use a Compute Engine cluster for the service.
- D. Use a dedicated Compute Engine virtual machine instance for the service.

Answer: D

Explanation:

Reference: <https://www.qwiklabs.com/focuses/611?parent=catalog>

NEW QUESTION 8

- (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 9

- (Exam Topic 2)

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

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

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

Answer: A

NEW QUESTION 10

- (Exam Topic 2)

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 10

- (Exam Topic 2)

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

- A. Download the file using scp, change the file, and then upload the modified version

- B. Configure and log in to the Compute Engine instance through SSH, and change the file
- C. Configure and log in to the Compute Engine instance through the serial port, and change the file
- D. Configure and log in to the Compute Engine instance using a remote desktop client, and change the file

Answer: C

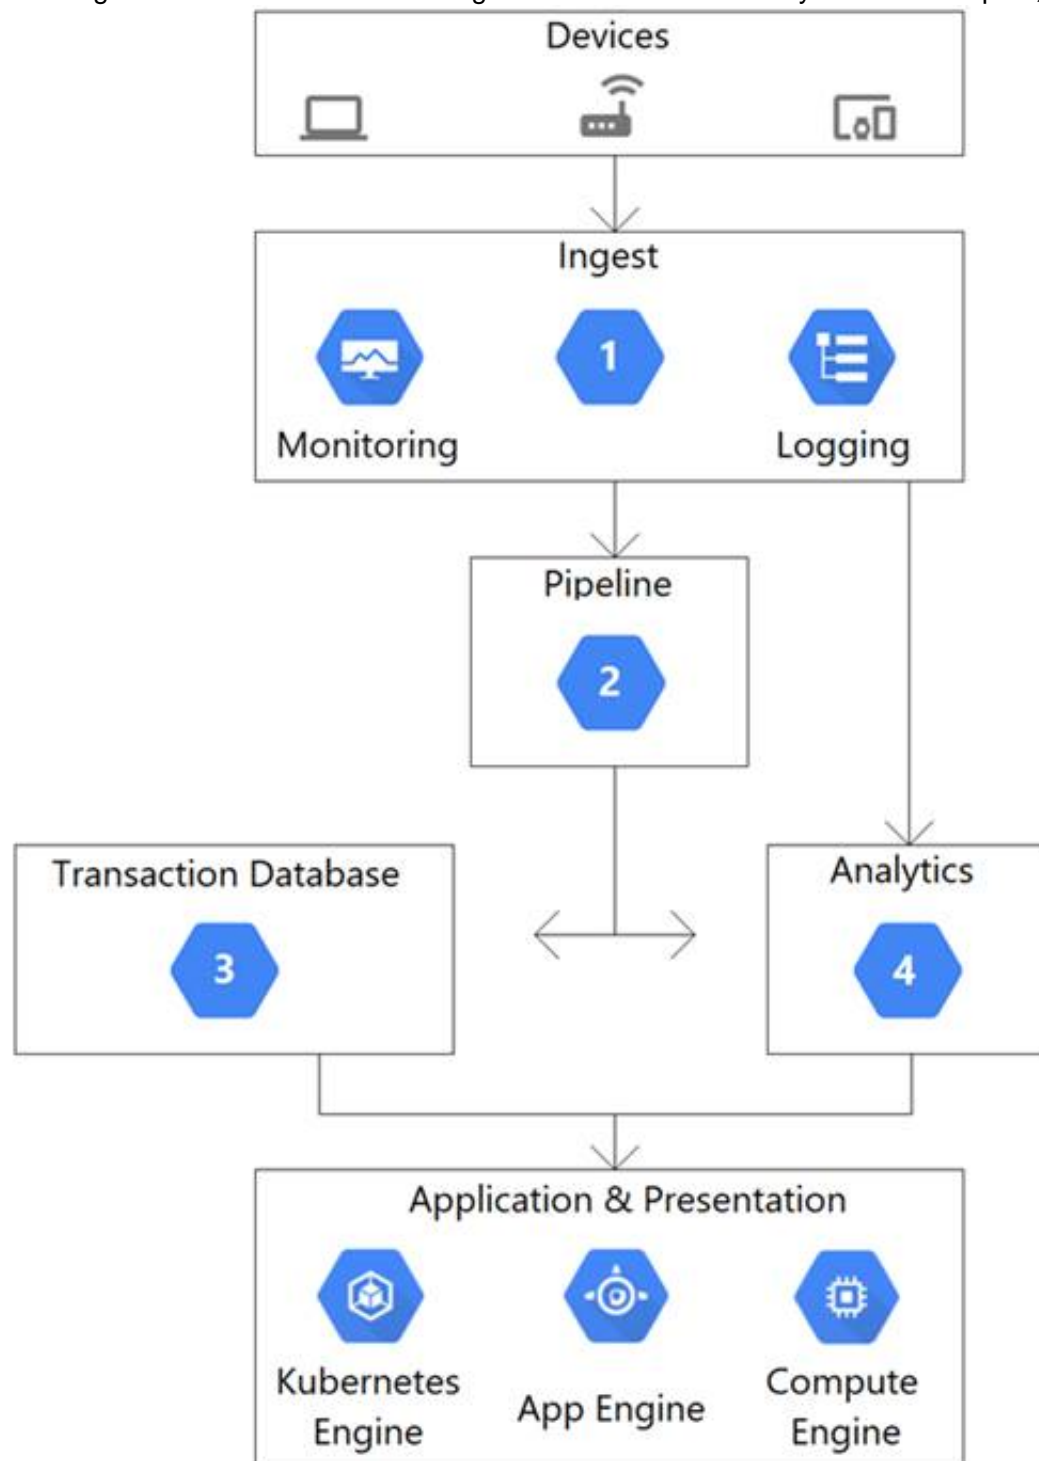
Explanation:

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

NEW QUESTION 14

- (Exam Topic 2)

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



- A. 1) App Engine2) Pub/Sub3) BigQuery4) Firestore
- B. 1) Dataflow2) Pub/Sub3) Firestore4) BigQuery
- C. 1) Pub/Sub2) Dataflow3) BigQuery4) Firestore
- D. 1) Pub/Sub2) Dataflow3) Firestore4) BigQuery

Answer: D

NEW QUESTION 15

- (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 16

- (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 21

- (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 the API 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 23

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 26

- (Exam Topic 2)

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

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

Answer: AC

NEW QUESTION 29

- (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 30

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 35

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 38

- (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 43

- (Exam Topic 2)

You are deploying your application to a Compute Engine virtual machine instance. Your application is configured to write its log files to disk. You want to view the logs in Stackdriver Logging without changing the application code. What should you do?

- A. Install the Stackdriver Logging Agent and configure it to send the application logs.
- B. Use a Stackdriver Logging Library to log directly from the application to Stackdriver Logging.
- C. Provide the log file folder path in the metadata of the instance to configure it to send the application logs.
- D. Change the application to log to /var/log so that its logs are automatically sent to Stackdriver Logging.

Answer: A

NEW QUESTION 47

- (Exam Topic 2)

You recently developed an application. You need to call the Cloud Storage API from a Compute Engine instance that doesn't have a public IP address. What should you do?

- A. Use Carrier Peering
- B. Use VPC Network Peering
- C. Use Shared VPC networks
- D. Use Private Google Access

Answer: D

Explanation:

<https://cloud.google.com/vpc/docs/private-google-access>

NEW QUESTION 48

- (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 49

- (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 51

- (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 53

- (Exam Topic 2)

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

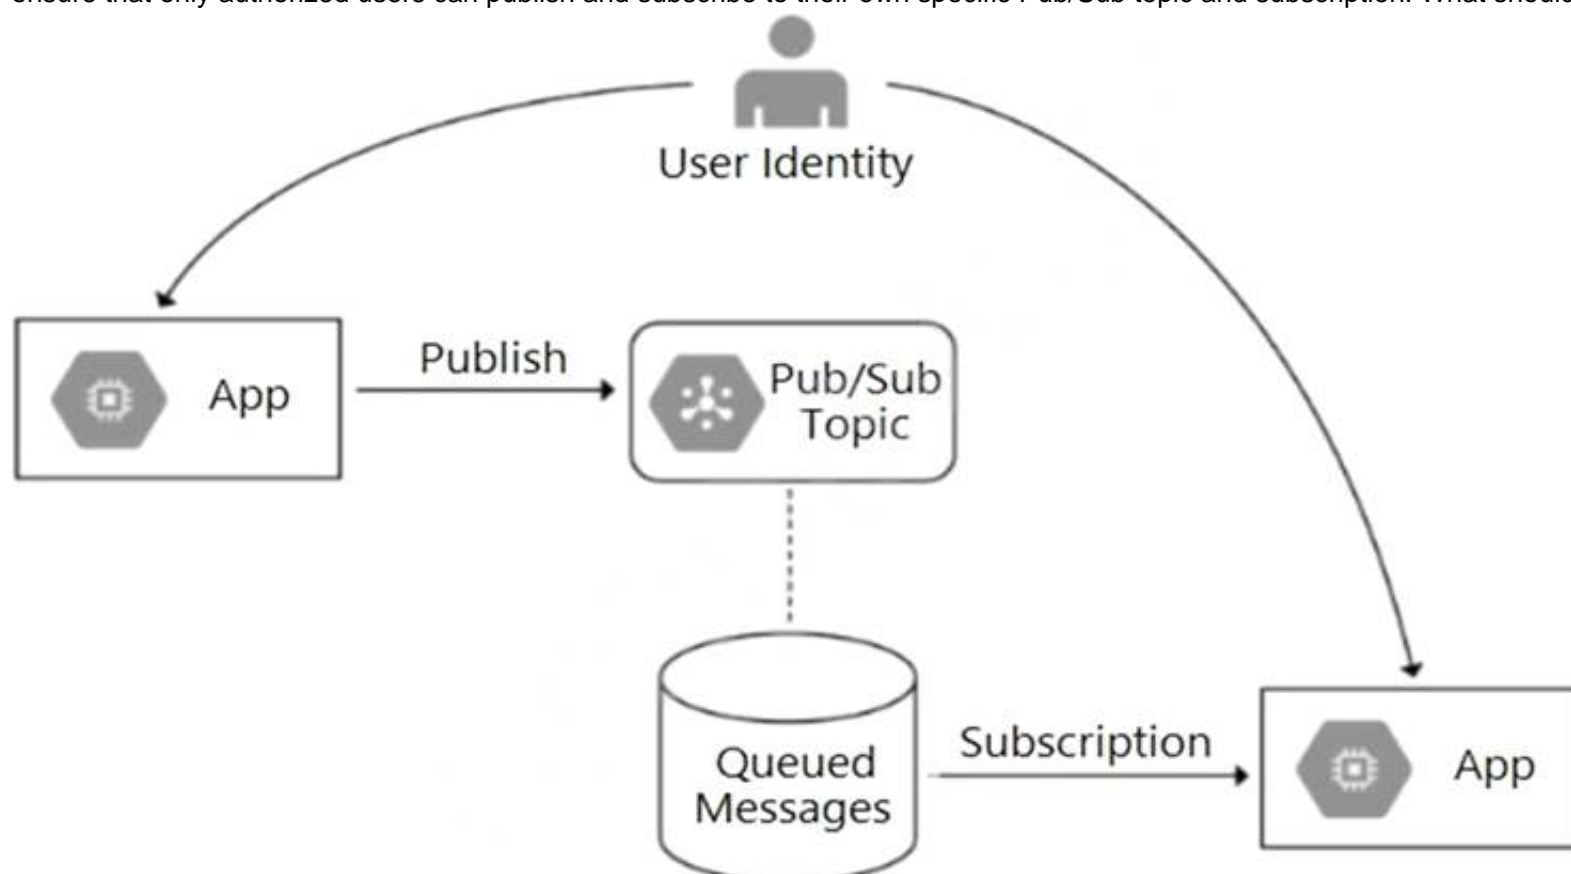
- A. Create a Cloud SQL table that has TransactionId and CustomerId configured as primary key
- B. Use an incremental number for the TransactionId.
- C. Create a Cloud SQL table that has TransactionId and CustomerId configured as primary key
- D. Use a random string (UUID) for the TransactionId.
- E.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

Answer: C

NEW QUESTION 57

- (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 60

- (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 62

- (Exam Topic 2)

You have a container deployed on Google Kubernetes Engine. The container can sometimes be slow to launch, so you have implemented a liveness probe. You notice that the liveness probe occasionally fails on launch. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 63

- (Exam Topic 2)

You have an application that uses an HTTP Cloud Function to process user activity from both desktop browser and mobile application clients. This function will serve as the endpoint for all metric submissions using HTTP POST.

Due to legacy restrictions, the function must be mapped to a domain that is separate from the domain requested by users on web or mobile sessions. The domain for the Cloud Function is <https://fn.example.com>. Desktop and mobile clients use the domain <https://www.example.com>. You need to add a header to the function's HTTP response so that only those browser and mobile sessions can submit metrics to the Cloud Function. Which response header should you add?

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

Answer: D

NEW QUESTION 64

- (Exam Topic 2)

You need to migrate an internal file upload API with an enforced 500-MB file size limit to App Engine. What should you do?

- A. Use FTP to upload files.
- B. Use CPanel to upload files.
- C. Use signed URLs to upload files.
- D. Change the API to be a multipart file upload API.

Answer: C

Explanation:

Reference: https://wiki.christophchamp.com/index.php?title=Google_Cloud_Platform

NEW QUESTION 67

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

Applications running on GKE might need access to Google Cloud APIs such as Compute Engine API, BigQuery Storage API, or Machine Learning APIs.

Workload Identity allows a Kubernetes service account in your GKE cluster to act as an IAM service account. Pods that use the configured Kubernetes service account automatically authenticate as the IAM service account when accessing Google Cloud APIs. Using Workload Identity allows you to assign distinct, fine-grained identities and authorization for each application in your cluster.

NEW QUESTION 71

- (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 72

- (Exam Topic 2)

You have recently instrumented a new application with OpenTelemetry, and you want to check the latency of your application requests in Trace. You want to ensure that a specific request is always traced. What should you do?

- A. Wait 10 minutes, then verify that Trace captures those types of requests automatically.
- B. Write a custom script that sends this type of request repeatedly from your dev project.
- C. Use the Trace API to apply custom attributes to the trace.
- D. Add the X-Cloud-Trace-Context header to the request with the appropriate parameters.

Answer: D

Explanation:

<https://cloud.google.com/trace/docs/setup#force-trace>

Cloud Trace doesn't sample every request. To force a specific request to be traced, add an X-Cloud-Trace-Context header to the request.

NEW QUESTION 73

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 74

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 78

- (Exam Topic 2)

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

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

Answer: C

Explanation:

Reference:

<https://medium.com/google-cloud/how-to-securely-invoke-a-cloud-function-from-google-kubernetes-engine-run>

NEW QUESTION 83

- (Exam Topic 2)

Your team is developing an ecommerce platform for your company. Users will log in to the website and add items to their shopping cart. Users will be automatically logged out after 30 minutes of inactivity. When users log back in, their shopping cart should be saved. How should you store users' session and shopping cart information while following Google-recommended best practices?

- A. Store the session information in Pub/Sub, and store the shopping cart information in Cloud SQL.
- B. Store the shopping cart information in a file on Cloud Storage where the filename is the SESSION ID.
- C. Store the session and shopping cart information in a MySQL database running on multiple Compute Engine instances.
- D. Store the session information in Memorystore for Redis or Memorystore for Memcached, and store the shopping cart information in Firestore.

Answer: D

NEW QUESTION 86

- (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 minutes
- 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 progress
- F. After the application reads the message, it disconnects from the database.

Answer: D

NEW QUESTION 88

- (Exam Topic 2)

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 90

- (Exam Topic 2)

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

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

Answer: C

Explanation:

Reference: <https://docs.docker.com/docker-hub/builds/>

NEW QUESTION 92

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 95

- (Exam Topic 2)

Your company wants to expand their users outside the United States for their popular application. The company wants to ensure 99.999% availability of the database for their application and also wants to minimize the read latency for their users across the globe.

Which two actions should they take? (Choose two.)

- A. Create a multi-regional Cloud Spanner instance with "nam-asia-eur1" configuration.
- B. Create a multi-regional Cloud Spanner instance with "nam3" configuration.
- C. Create a cluster with at least 3 Spanner nodes.
- D. Create a cluster with at least 1 Spanner node.
- E. Create a minimum of two Cloud Spanner instances in separate regions with at least one node.
- F. Create a Cloud Dataflow pipeline to replicate data across different databases.

Answer: BF

NEW QUESTION 98

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 100

- (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 installed
- D. Configure Apache web server to host the application.
- E. Containerize your application first
- F. Deploy this container to Google Kubernetes Engine (GKE) and assign an 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 103

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 106

- (Exam Topic 2)

You need to redesign the ingestion of audit events from your authentication service to allow it to handle a large increase in traffic. Currently, the audit service and the authentication system run in the same Compute Engine virtual machine. You plan to use the following Google Cloud tools in the new architecture:

Multiple Compute Engine machines, each running an instance of the authentication service
Multiple Compute Engine machines, each running an instance of the audit service

Pub/Sub to send the events from the authentication services.

How should you set up the topics and subscriptions to ensure that the system can handle a large volume of messages and can scale efficiently?

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

Answer: A

Explanation:

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

NEW QUESTION 111

- (Exam Topic 2)

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

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

Answer: AD

Explanation:

<https://cloud.google.com/storage/docs/bucket-lock>

This page discusses the Bucket Lock feature, which allows you to configure a data retention policy for a Cloud Storage bucket that governs how long objects in the bucket must be retained. The feature also allows you to lock the data retention policy, permanently preventing the policy from being reduced or removed.

<https://cloud.google.com/storage/docs/storage-classes#archive>

Archive storage is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Unlike the "coldest" storage services offered by other Cloud providers, your data is available within milliseconds, not hours or days.

Archive storage is the best choice for data that you plan to access less than once a year.

NEW QUESTION 116

- (Exam Topic 2)

You manage an application that runs in a Compute Engine instance. You also have multiple backend services executing in stand-alone Docker containers running in Compute Engine instances. The Compute Engine instances supporting the backend services are scaled by managed instance groups in multiple regions. You want your calling application to be loosely coupled. You need to be able to invoke distinct service implementations that are chosen based on the value of an HTTP header found in the request. Which Google Cloud feature should you use to invoke the backend services?

- A. Traffic Director
- B. Service Directory
- C. Anthos Service Mesh
- D. Internal HTTP(S) Load Balancing

Answer: D

NEW QUESTION 117

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 119

- (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 123

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 126

- (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 131

- (Exam Topic 2)

You are a developer working on an internal application for payroll processing. You are building a component of the application that allows an employee to submit a timesheet, which then initiates several steps:

- An email is sent to the employee and manager, notifying them that the timesheet was submitted.
- A timesheet is sent to payroll processing for the vendor's API.
- A timesheet is sent to the data warehouse for headcount planning.

These steps are not dependent on each other and can be completed in any order. New steps are being considered and will be implemented by different development teams. Each development team will implement the error handling specific to their step. What should you do?

- A. Deploy a Cloud Function for each step that calls the corresponding downstream system to complete the required action.
- B. Create a Pub/Sub topic for each step
- C. Create a subscription for each downstream development team to subscribe to their step's topic.
- D. Create a Pub/Sub topic for timesheet submission
- E. Create a subscription for each downstream development team to subscribe to the topic.
- F. Create a timesheet microservice deployed to Google Kubernetes Engine
- G. The microservice calls each downstream step and waits for a successful response before calling the next step.

Answer: C

NEW QUESTION 136

- (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 139

- (Exam Topic 2)

Your company has a data warehouse that keeps your application information in BigQuery. The BigQuery data warehouse keeps 2 PBs of user data. Recently, your company expanded your user base to include EU users and needs to comply with these requirements:

Your company must be able to delete all user account information upon user request. All EU user data must be stored in a single region specifically for EU users.

Which two actions should you take? (Choose two.)

- A. Use BigQuery federated queries to query data from Cloud Storage.
- B. Create a dataset in the EU region that will keep information about EU users only.
- C. Create a Cloud Storage bucket in the EU region to store information for EU users only.
- D. Re-upload your data using to a Cloud Dataflow pipeline by filtering your user records out.
- E. Use DML statements in BigQuery to update/delete user records based on their requests.

Answer: CE

Explanation:

Reference: <https://cloud.google.com/solutions/bigquery-data-warehouse>

NEW QUESTION 140

- (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 143

- (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 group

- B. Replicate the static content across the virtual machines (VMs)2) Create an external HTTP(S) load balancer.3) Enable Cloud CDN, and send traffic to the managed instance group.
- C. 1) Create an unmanaged instance group
- D. Replicate the static content across the VMs.2) Create an external HTTP(S) load balancer3) Enable Cloud CDN, and send traffic to the unmanaged instance group.
- E. 1) Create a Standard storage class, regional Cloud Storage bucket
- 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 bucket
- 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 144

- (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 149

- (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 151

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 156

- (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 160

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 162

- (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/](https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/).
- E. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[API_NAME\]/\[API_VERSION\]/](https://[API_NAME]/[API_VERSION]/).

Answer: D

NEW QUESTION 166

- (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 168

- (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 171

- (Exam Topic 2)

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

you do?

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

Answer: A

Explanation:

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

<https://blog.knoldus.com/how-to-avoid-outages-in-your-kubernetes-cluster-using-pdb/>

NEW QUESTION 175

- (Exam Topic 2)

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

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

Answer: D

Explanation:

https://cloud.google.com/service-mesh/docs/overview#security_benefits

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

NEW QUESTION 180

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 181

- (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 185

- (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 190

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 191

- (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 192

- (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 196

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 201

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 204

- (Exam Topic 2)

You have an application running in App Engine. Your application is instrumented with Stackdriver Trace. The /product-details request reports details about four known unique products at /sku-details as shown below. You want to reduce the time it takes for the request to complete. What should you do?

Timeline



- A. Increase the size of the instance class.
- B. Change the Persistent Disk type to SSD.
- C. Change /product-details to perform the requests in parallel.
- D. Store the /sku-details information in a database, and replace the webservice call with a database query.

Answer: C

NEW QUESTION 208

- (Exam Topic 2)

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

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

Answer: D

Explanation:

Reference:

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

NEW QUESTION 210

- (Exam Topic 2)

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

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

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

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

Answer: AC

NEW QUESTION 214

- (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 218

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 219

- (Exam Topic 2)

You have an on-premises application that authenticates to the Cloud Storage API using a user-managed service account with a user-managed key. The application connects to Cloud Storage using Private Google Access over a Dedicated Interconnect link. You discover that requests from the application to access objects in the Cloud Storage bucket are failing with a 403 Permission Denied error code. What is the likely cause of this issue?

- A. The folder structure inside the bucket and object paths have changed.
- B. The permissions of the service account's predefined role have changed.
- C. The service account key has been rotated but not updated on the application server.
- D. The Interconnect link from the on-premises data center to Google Cloud is experiencing a temporary outage.

Answer: C

NEW QUESTION 221

- (Exam Topic 2)

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

Which method should they use?

- A. Use Cloud Build with a trigger configured for each source code commit.
- B. Use Jenkins deployed via the Google Cloud Platform Marketplace, configured to watch for source code commits.

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

Answer: A

Explanation:

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

NEW QUESTION 223

- (Exam Topic 2)

You manage your company's ecommerce platform's payment system, which runs on Google Cloud. Your company must retain user logs for 1 year for internal auditing purposes and for 3 years to meet compliance requirements. You need to store new user logs on Google Cloud to minimize on-premises storage usage and ensure that they are easily searchable. You want to minimize effort while ensuring that the logs are stored correctly. What should you do?

- A. Store the logs in a Cloud Storage bucket with bucket lock turned on.
- B. Store the logs in a Cloud Storage bucket with a 3-year retention period.
- C. Store the logs in Cloud Logging as custom logs with a custom retention period.
- D. Store the logs in a Cloud Storage bucket with a 1-year retention period.
- E. After 1 year, move the logs to another bucket with a 2-year retention period.

Answer: C

Explanation:

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

NEW QUESTION 225

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 226

- (Exam Topic 2)

Your development team has been asked to refactor an existing monolithic application into a set of composable microservices. Which design aspects should you implement for the new application? (Choose two.)

- A. Develop the microservice code in the same programming language used by the microservice caller.
- B. Create an API contract agreement between the microservice implementation and microservice caller.
- C. Require asynchronous communications between all microservice implementations and microservice callers.
- D. Ensure that sufficient instances of the microservice are running to accommodate the performance requirements.
- E. Implement a versioning scheme to permit future changes that could be incompatible with the current interface.

Answer: BE

NEW QUESTION 228

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 230

- (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 233

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

Professional-Cloud-Developer Practice Exam Features:

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

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