

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 deployment
- B. Provide developers with local MySQL instances to conduct testing on the application changes.
- C. Migrate data to Bigtable
- 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 Europe
- 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 instances

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 has connected their Hadoop infrastructure to GCP using Cloud Interconnect in order to query data stored on persistent disks.

Which IP strategy should they use?

- A. Create manual subnets.
- B. Create an auto mode subnet.
- C. Create multiple peered VPCs.
- D. Provision a single instance for NAT.

Answer: A

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.

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

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

Answer: D

Explanation:

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

NEW QUESTION 7

- (Exam Topic 1)

Which service should HipLocal use for their public APIs?

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

Answer: D

NEW QUESTION 8

- (Exam Topic 1)

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

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

Answer: BC

NEW QUESTION 9

- (Exam Topic 2)

You work for an organization that manages an ecommerce site. Your application is deployed behind a global HTTP(S) load balancer. You need to test a new product recommendation algorithm. You plan to use A/B testing to determine the new algorithm's effect on sales in a randomized way. How should you test this feature?

- A. Split traffic between versions using weights.
- B. Enable the new recommendation feature flag on a single instance.
- C. Mirror traffic to the new version of your application.
- D. Use HTTP header-based routing.

Answer: A

Explanation:

https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic_actions_weight-based_traffic Deploying a new version of an existing production service generally incurs some risk. Even if your tests pass in staging, you probably don't want to subject 100% of your users to the new version immediately. With traffic management, you can define percentage-based traffic splits across multiple backend services.

For example, you can send 95% of the traffic to the previous version of your service and 5% to the new version of your service. After you've validated that the new production version works as expected, you can gradually shift the percentages until 100% of the traffic reaches the new version of your service. Traffic splitting is typically used for deploying new versions, A/B testing, service migration, and similar processes.

https://cloud.google.com/traffic-director/docs/advanced-traffic-management#weight-based_traffic_splitting_for_https https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#split_the_traffic_https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic_actions_weight-based_traffic

NEW QUESTION 10

- (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 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 15

- (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

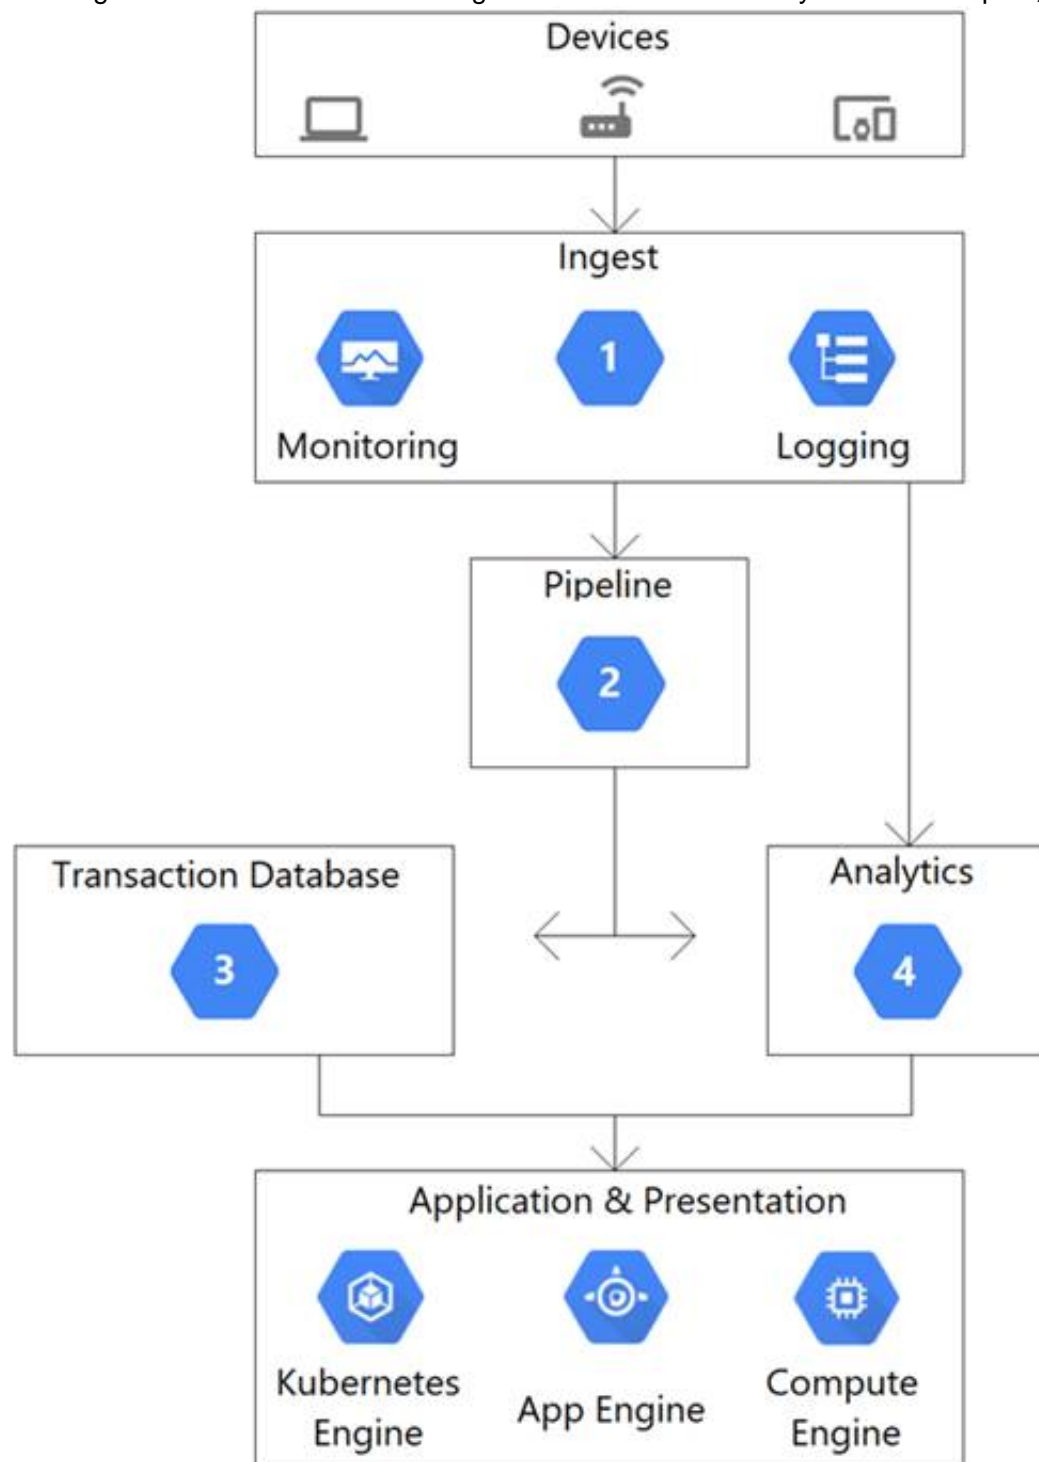
Explanation:

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

NEW QUESTION 19

- (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 20

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 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 26

- (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 file
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 file
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 file
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 file
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 30

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 31

- (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 35

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 38

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 41

- (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 44

- (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 47

- (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 50

- (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 55

- (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 56

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 58

- (Exam Topic 2)

You are building a new API. You want to minimize the cost of storing and reduce the latency of serving images. Which architecture should you use?

- A. App Engine backed by Cloud Storage
- B. Compute Engine backed by Persistent Disk
- C. Transfer Appliance backed by Cloud Filestore
- D. Cloud Content Delivery Network (CDN) backed by Cloud Storage

Answer: B

NEW QUESTION 62

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 64

- (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 68

- (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 73

- (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 77

- (Exam Topic 2)

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

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

application requests a Google Server to exchange the authorization code with an access token.4) The Google server responds with the access token that is used by the application to call the Google Cloud API.

Answer: D

Explanation:

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

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

NEW QUESTION 78

- (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 82

- (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 87

- (Exam Topic 2)

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

- A) Run `gcloud builds submit` in the directory that contains the application source code.
- B) Run `gcloud run deploy app-name --image gcr.io/$PROJECT_ID/app-name` in the directory that contains the application source code.
- C) Run `gcloud container images add-tag gcr.io/$PROJECT_ID/app-name gcr.io/$PROJECT_ID/app-name:latest` in the directory that contains the application source code.

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

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

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

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

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

Answer: AD

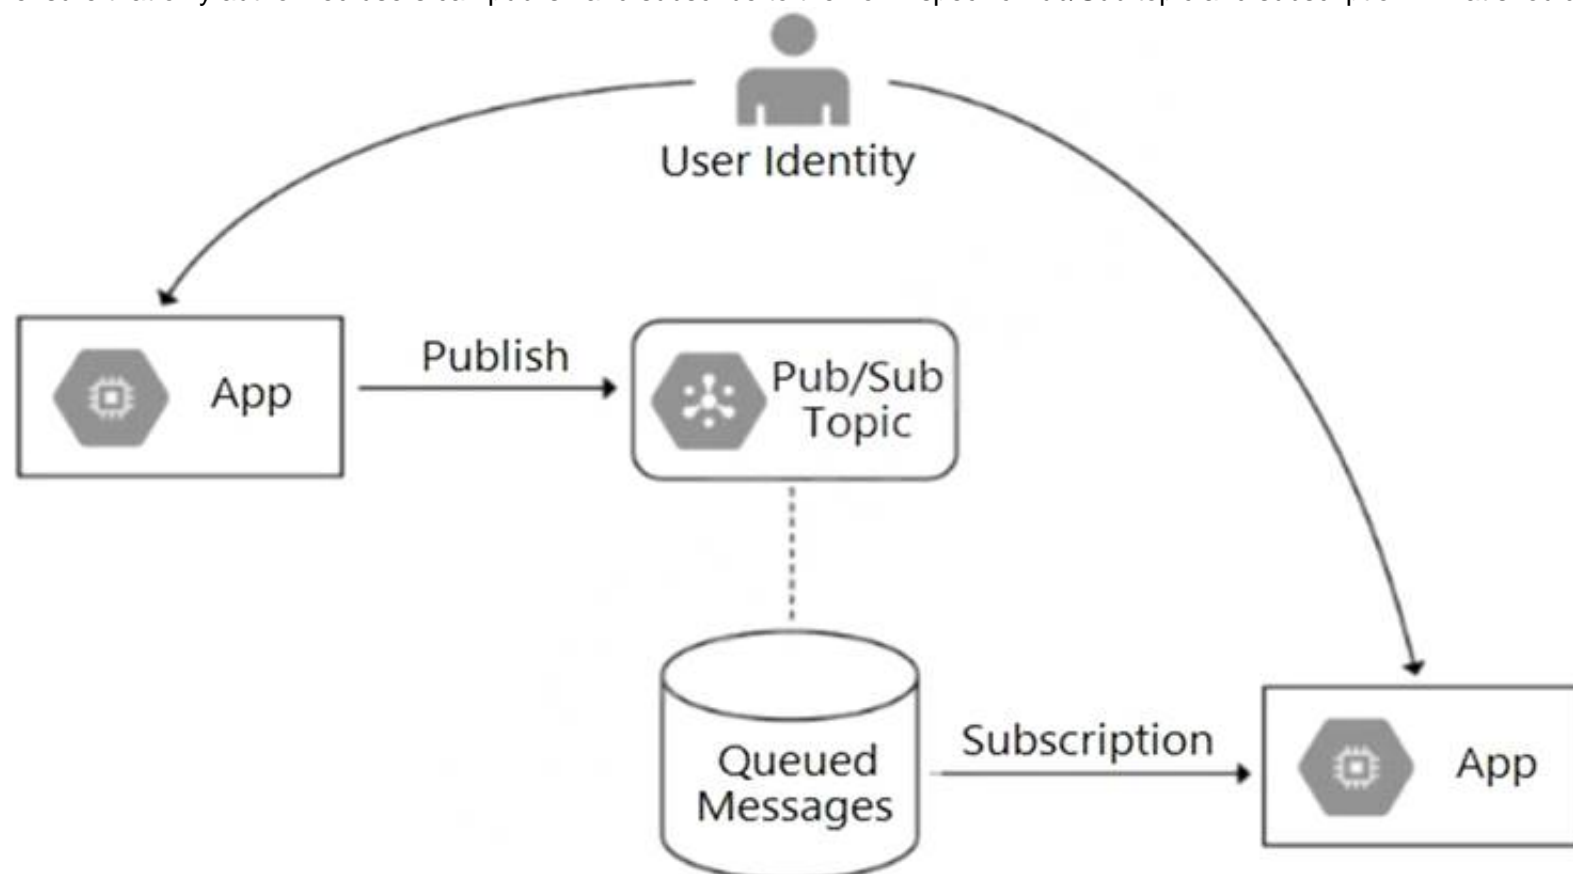
Explanation:

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

NEW QUESTION 89

- (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 94

- (Exam Topic 2)

You have two tables in an ANSI-SQL compliant database with identical columns that you need to quickly combine into a single table, removing duplicate rows from the result set.

What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 97

- (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 99

- (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 102

- (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 103

- (Exam Topic 2)

You are developing an ecommerce application that stores customer, order, and inventory data as relational tables inside Cloud Spanner. During a recent load test, you discover that Spanner performance is not scaling linearly as expected. Which of the following is the cause?

- A. The use of 64-bit numeric types for 32-bit numbers.
- B. The use of the STRING data type for arbitrary-precision values.
- C. The use of Version 1 UUIDs as primary keys that increase monotonically.
- D. The use of LIKE instead of STARTS_WITH keyword for parameterized SQL queries.

Answer: C

NEW QUESTION 106

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 109

- (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 113

- (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 patter
- C. Use a Git tag convention for new releases.
- D. Create a build trigger on a Git branch name patter
- 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 trigge
- G. Use this repository for new releases only.

Answer: C

Explanation:

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

NEW QUESTION 116

- (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 117

- (Exam Topic 2)

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

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

Answer: C

Explanation:

<https://cloud.google.com/pubsub/docs/pull>

NEW QUESTION 120

- (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 123

- (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 124

- (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 126

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 129

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 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 134

- (Exam Topic 2)

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

What should you do?

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

Answer: B

NEW QUESTION 138

- (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 142

- (Exam Topic 2)

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

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

G. Connect to the web instances using the bastion host.

Answer: C

Explanation:

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

NEW QUESTION 147

- (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 149

- (Exam Topic 2)

You have an application deployed in production. When a new version is deployed, you want to ensure that all production traffic is routed to the new version of your application. You also want to keep the previous version deployed so that you can revert to it if there is an issue with the new version. Which deployment strategy should you use?

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

Answer: A

NEW QUESTION 154

- (Exam Topic 2)

You are configuring a continuous integration pipeline using Cloud Build to automate the deployment of new container images to Google Kubernetes Engine (GKE). The pipeline builds the application from its source code, runs unit and integration tests in separate steps, and pushes the container to Container Registry. The application runs on a Python web server.

The Dockerfile is as follows: FROM python:3.7-alpine - COPY . /app WORKDIR /app

RUN pip install -r requirements.txt CMD ["unicorn", "-w 4", "main:app"]

You notice that Cloud Build runs are taking longer than expected to complete. You want to decrease the build time. What should you do? (Choose two.)

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

Answer: AC

Explanation:

<https://cloud.google.com/build/docs/optimize-builds/increase-vcpu-for-builds>

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

https://cloud.google.com/build/docs/optimize-builds/speeding-up-builds#using_a_cached_docker_image

The easiest way to increase the speed of your Docker image build is by specifying a cached image that can be used for subsequent builds. You can specify the cached image by adding the --cache-from argument in your build config file, which will instruct Docker to build using that image as a cache source.

NEW QUESTION 159

- (Exam Topic 2)

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

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

The command fails with the error:


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

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

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

Answer: A

NEW QUESTION 161

- (Exam Topic 2)

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

- A. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Clients should use this IP address to connect to the service.
- B. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Then, define an A record in Cloud DNS.
- C. Clients should use the name of the A record to connect to the service.
- D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url `https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/`.
- 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 164

- (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 167

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 170

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

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

Answer: C

NEW QUESTION 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 182

- (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 186

- (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 187

- (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 190

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 192

- (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 194

- (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 197

- (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 202

- (Exam Topic 2)

You want to use the Stackdriver Logging Agent to send an application's log file to Stackdriver from a Compute Engine virtual machine instance. After installing the Stackdriver Logging Agent, what should you do first?

- A. Enable the Error Reporting API on the project.
- B. Grant the instance full access to all Cloud APIs.
- C. Configure the application log file as a custom source.
- D. Create a Stackdriver Logs Export Sink with a filter that matches the application's log entries.

Answer: B

NEW QUESTION 207

- (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 Function.
- 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 agent.
- 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 agent.
- H. If the code passes the tests, Cloud Build deploys the Cloud Function.

Answer: D

NEW QUESTION 212

- (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 216

- (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 217

- (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 220

.....

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