

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer

<https://www.2passeasy.com/dumps/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)

Which database should HipLocal use for storing user activity?

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

Answer: A

NEW QUESTION 3

- (Exam Topic 1)

HipLocal's data science team wants to analyze user reviews. How should they prepare the data?

- A. Use the Cloud Data Loss Prevention API for redaction of the review dataset.
- B. Use the Cloud Data Loss Prevention API for de-identification of the review dataset.
- C. Use the Cloud Natural Language Processing API for redaction of the review dataset.
- D. Use the Cloud Natural Language Processing API for de-identification of the review dataset.

Answer: B

Explanation:

<https://cloud.google.com/dlp/docs/deidentify-sensitive-data>

NEW QUESTION 4

- (Exam Topic 1)

HipLocal is configuring their access controls.

Which firewall configuration should they implement?

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

Answer: D

NEW QUESTION 5

- (Exam Topic 1)

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

Answer: D

Explanation:

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

NEW QUESTION 7

- (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 8

- (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 9

- (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 10

- (Exam Topic 1)

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

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

Answer: C

NEW QUESTION 10

- (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 14

- (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 16

- (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 17

- (Exam Topic 2)

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

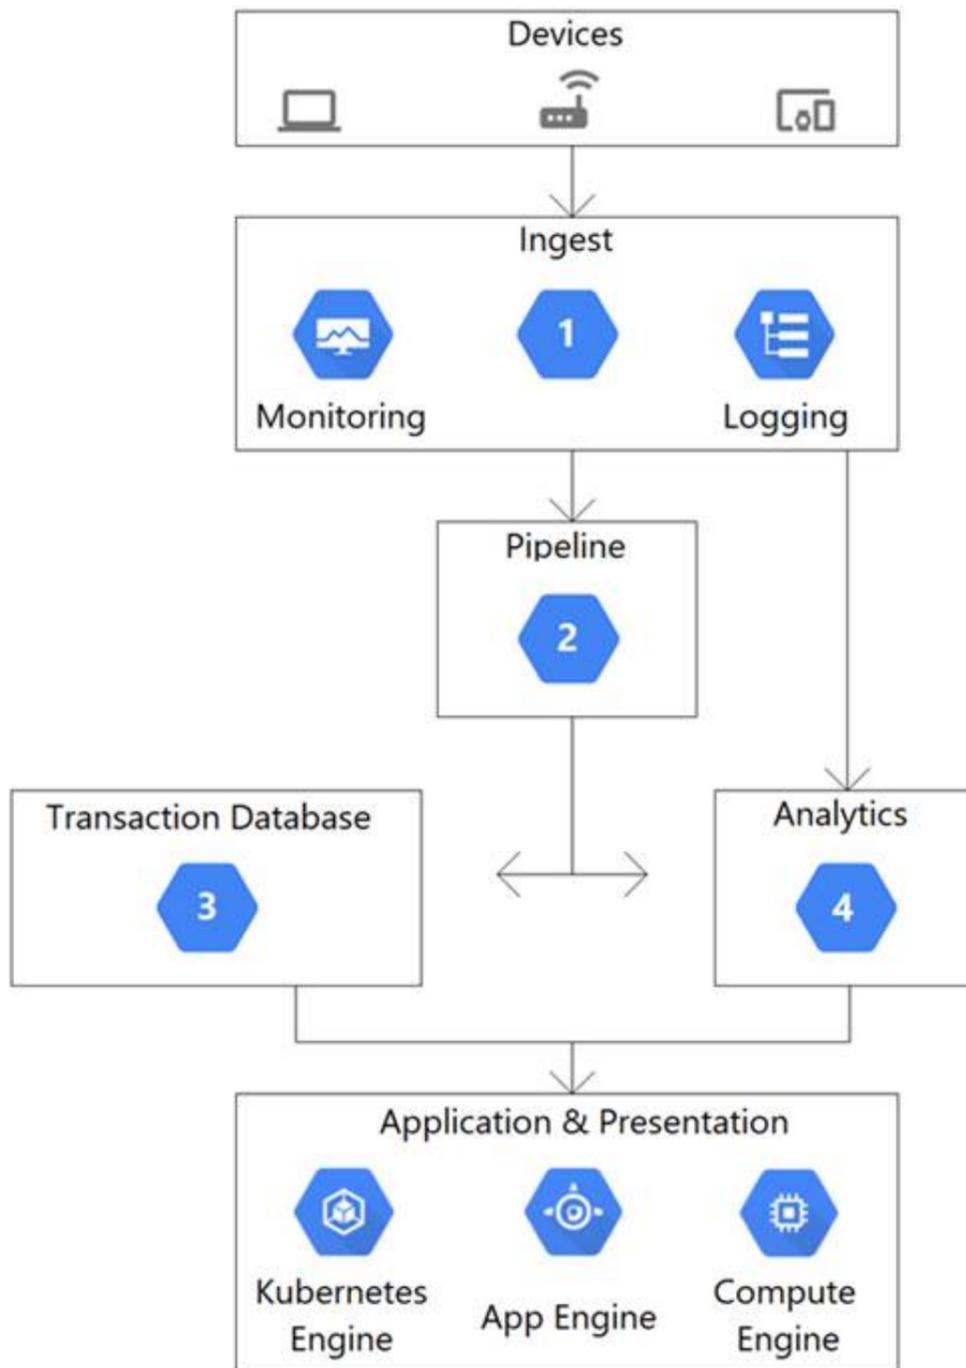
- A. Configure the application to send the file to the client as an email attachment.
- B. Generate and assign a Cloud Storage-signed URL for the file
- C. Make the URL available for the client to download.
- D. Create a temporary Cloud Storage bucket with time expiration specified, and give download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

Answer: B

NEW QUESTION 20

- (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 Engine 2) Pub/Sub 3) BigQuery 4) Firestore
- B. 1) Dataflow 2) Pub/Sub 3) Firestore 4) BigQuery
- C. 1) Pub/Sub 2) Dataflow 3) BigQuery 4) Firestore
- D. 1) Pub/Sub 2) Dataflow 3) Firestore 4) BigQuery

Answer: D

NEW QUESTION 24

- (Exam Topic 2)

Your team is developing a new application using a PostgreSQL database and Cloud Run. You are responsible for ensuring that all traffic is kept private on Google Cloud. You want to use managed services and follow Google-recommended best practices. What should you do?

- A. 1) Enable Cloud SQL and Cloud Run in the same project. 2) Configure a private IP address for Cloud SQL
- B. Enable private services access. 3) Create a Serverless VPC Access connector. 4) Configure Cloud Run to use the connector to connect to Cloud SQL.
- C. 1) Install PostgreSQL on a Compute Engine virtual machine (VM), and enable Cloud Run in the same project. 2) Configure a private IP address for the V
- D. Enable private services access. 3) Create a Serverless VPC Access connector. 4) Configure Cloud Run to use the connector to connect to the VM hosting PostgreSQL.
- E. 1) Use Cloud SQL and Cloud Run in different projects. 2) Configure a private IP address for Cloud SQL
- F. Enable private services access. 3) Create a Serverless VPC Access connector. 4) Set up a VPN connection between the two project
- G. Configure Cloud Run to use the connector to connect to Cloud SQL.
- H. 1) Install PostgreSQL on a Compute Engine VM, and enable Cloud Run in different projects. 2) Configure a private IP address for the V
- I. Enable private services access. 3) Create a Serverless VPC Access connector. 4) Set up a VPN connection between the two project
- J. Configure Cloud Run to use the connector to access the VM hosting PostgreSQL

Answer: A

Explanation:

<https://cloud.google.com/sql/docs/postgres/connect-run#private-ip>

NEW QUESTION 29

- (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 34

- (Exam Topic 2)

You have an application in production. It is deployed on Compute Engine virtual machine instances controlled by a managed instance group. Traffic is routed to the instances via a HTTP(s) load balancer. Your users are unable to access your application. You want to implement a monitoring technique to alert you when the application is unavailable.

Which technique should you choose?

- A. Smoke tests
- B. Stackdriver uptime checks
- C. Cloud Load Balancing - health checks
- D. Managed instance group - health checks

Answer: B

Explanation:

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

NEW QUESTION 35

- (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 37

- (Exam Topic 2)

Your existing application keeps user state information in a single MySQL database. This state information is very user-specific and depends heavily on how long a user has been using an application. The MySQL database is causing challenges to maintain and enhance the schema for various users.

Which storage option should you choose?

- A. Cloud SQL
- B. Cloud Storage
- C. Cloud Spanner
- D. Cloud Datastore/Firestore

Answer: A

Explanation:

Reference: <https://cloud.google.com/solutions/migrating-mysql-to-cloudsql-concept>

NEW QUESTION 40

- (Exam Topic 2)

You are developing a marquee stateless web application that will run on Google Cloud. The rate of the incoming user traffic is expected to be unpredictable, with no traffic on some days and large spikes on other days. You need the application to automatically scale up and down, and you need to minimize the cost associated with running the application. What should you do?

- A. Build the application in Python with Firestore as the databas
- B. Deploy the application to Cloud Run.
- C. Build the application in C# with Firestore as the databas
- D. Deploy the application to App Engine flexible environment.
- E. Build the application in Python with CloudSQL as the databas
- F. Deploy the application to App Engine standard environment.
- G. Build the application in Python with Firestore as the databas
- H. Deploy the application to a Compute Engine managed instance group with autoscaling.

Answer: A

NEW QUESTION 43

- (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 46

- (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 48

- (Exam Topic 2)

Your data is stored in Cloud Storage buckets. Fellow developers have reported that data downloaded from Cloud Storage is resulting in slow API performance. You want to research the issue to provide details to the GCP support team. Which command should you run?

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

Answer: B

Explanation:

Reference: <https://groups.google.com/forum/#!topic/gce-discussion/xBI9Jq5HDsY>

NEW QUESTION 52

- (Exam Topic 2)

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

- A. Deploy the application to Compute Engine and turn on autoscaling.
- B. Replace the application's features with appropriate microservices in phases.
- C. Refactor the monolithic application with appropriate microservices in a single effort and deploy it.
- D. Build a new application with the appropriate microservices separate from the monolith and replace it when it is complete.

Answer: C

Explanation:

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

NEW QUESTION 57

- (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 62

- (Exam Topic 2)

You are designing a resource-sharing policy for applications used by different teams in a Google Kubernetes Engine cluster. You need to ensure that all applications can access the resources needed to run. What should you do? (Choose two.)

- A. Specify the resource limits and requests in the object specifications.
- B. Create a namespace for each team, and attach resource quotas to each namespace.
- C. Create a LimitRange to specify the default compute resource requirements for each namespace.
- D. Create a Kubernetes service account (KSA) for each application, and assign each KSA to the namespace.
- E. Use the Anthos Policy Controller to enforce label annotations on all namespace
- F. Use taints and tolerations to allow resource sharing for namespaces.

Answer: BC

Explanation:

<https://kubernetes.io/docs/concepts/policy/resource-quotas/> <https://kubernetes.io/docs/concepts/policy/limit-range/>
<https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-resource-requests-and-l>

NEW QUESTION 65

- (Exam Topic 2)

You are running an application on App Engine that you inherited. You want to find out whether the application is using insecure binaries or is vulnerable to XSS attacks. Which service should you use?

- A. Cloud Amor
- B. Stackdriver Debugger
- C. Cloud Security Scanner
- D. Stackdriver Error Reporting

Answer: C

Explanation:

Reference: <https://cloud.google.com/security-scanner>

NEW QUESTION 69

- (Exam Topic 2)

Your application requires service accounts to be authenticated to GCP products via credentials stored on its host Compute Engine virtual machine instances. You want to distribute these credentials to the host instances as securely as possible. What should you do?

- A. Use HTTP signed URLs to securely provide access to the required resources.
- B. Use the instance's service account Application Default Credentials to authenticate to the required resources.
- C. Generate a P12 file from the GCP Console after the instance is deployed, and copy the credentials to the host instance before starting the application.
- D. Commit the credential JSON file into your application's source repository, and have your CI/CD process package it with the software that is deployed to the instance.

Answer: B

Explanation:

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

NEW QUESTION 74

- (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 77

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 78

- (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 83

- (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 88

- (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 90

- (Exam Topic 2)

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

- A. Package each component in a separate container
- B. Implement readiness and liveness probes.
- C. Package the application in a single container
- D. Use a process management tool to manage each component.
- E. Package each component in a separate container
- F. Use a script to orchestrate the launch of the components.
- G. Package the application in a single container
- H. Use a bash script as an entrypoint to the container, and then spawn each component as a background job.

Answer: A

Explanation:

<https://cloud.google.com/blog/products/containers-kubernetes/7-best-practices-for-building-containers> <https://cloud.google.com/architecture/best-practices-for-building-containers>

"classic Apache/MySQL/PHP stack: you might be tempted to run all the components in a single container. However, the best practice is to use two or three different containers: one for Apache, one for MySQL, and potentially one for PHP if you are running PHP-FPM."

NEW QUESTION 95

- (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 100

- (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 102

- (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 107

- (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 110

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 111

- (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 112

- (Exam Topic 2)

You are building a mobile application that will store hierarchical data structures in a database. The application will enable users working offline to sync changes when they are back online. A backend service will enrich the data in the database using a service account. The application is expected to be very popular and needs to scale seamlessly and securely. Which database and IAM role should you use?

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

Answer: C

Explanation:

<https://firebase.google.com/docs/firestore/manage-data/enable-offline>

Cloud Firestore supports offline data persistence. This feature caches a copy of the Cloud Firestore data that your app is actively using, so your app can access the data when the device is offline. You can write, read, listen to, and query the cached data. When the device comes back online, Cloud Firestore synchronizes any local changes made by your app to the Cloud Firestore backend.

NEW QUESTION 115

- (Exam Topic 2)

You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. The application exposes an HTTP-based health check at /healthz. You want to use this health check endpoint to determine whether traffic should be routed to the pod by the load balancer. Which code snippet should you include in your Pod configuration?

- A.

```
livenessProbe:
  httpGet:
    path: /healthz
    port: 80
```
- B.

```
readinessProbe:
  httpGet:
    path: /healthz
    port: 80
```
- C.

```
loadbalancerHealthCheck:
  httpGet:
    path: /healthz
    port: 80
```
- D.

```
healthCheck:
  httpGet:
    path: /healthz
    port: 80
```

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

Answer: B

Explanation:

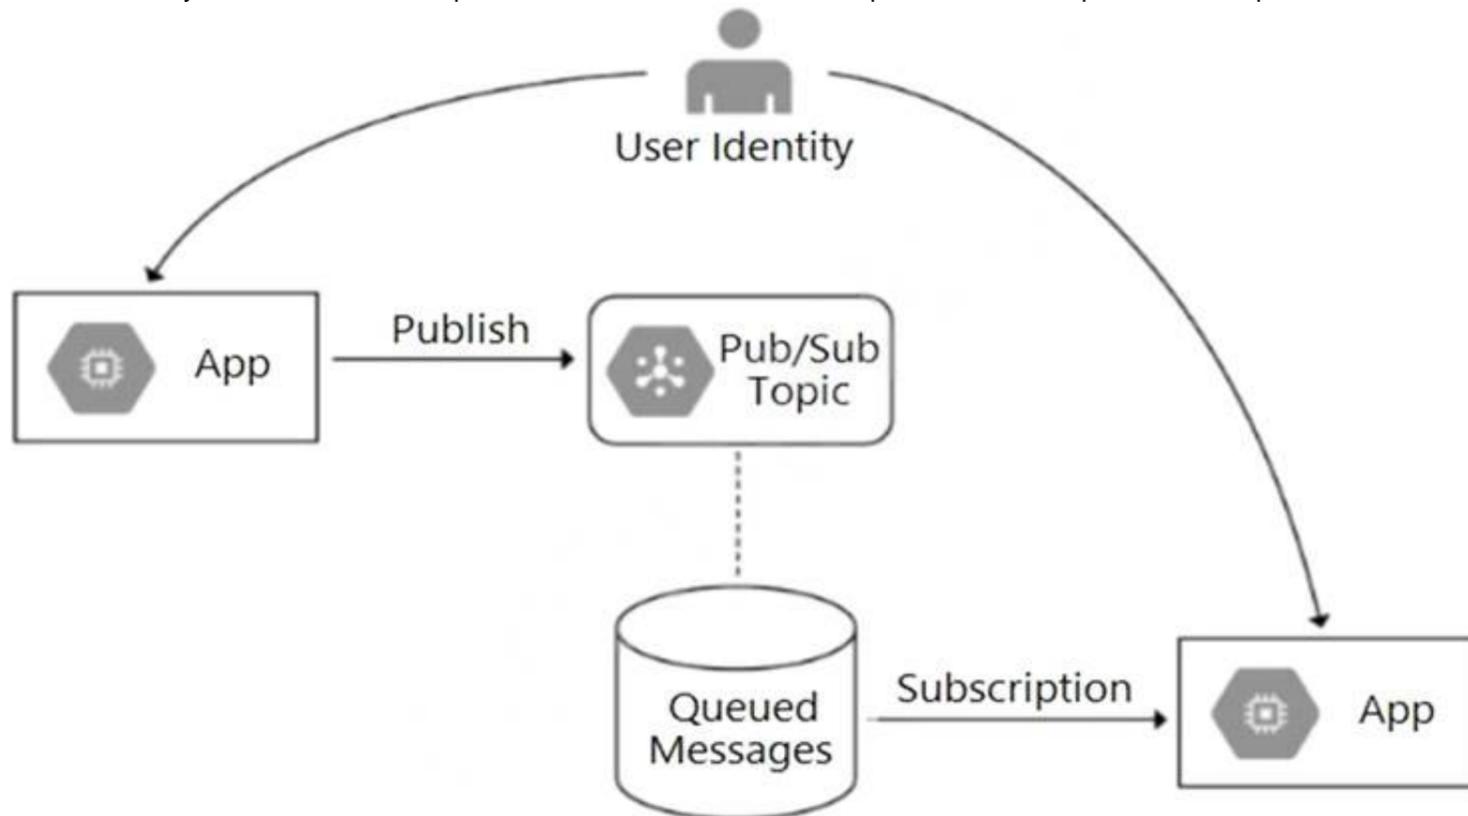
For the GKE ingress controller to use your readinessProbes as health checks, the Pods for an Ingress must exist at the time of Ingress creation. If your replicas

are scaled to 0, the default health check will apply.

NEW QUESTION 120

- (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 121

- (Exam Topic 2)

You recently developed a new application. You want to deploy the application on Cloud Run without a Dockerfile. Your organization requires that all container images are pushed to a centrally managed container repository. How should you build your container using Google Cloud services? (Choose two.)

- A. Push your source code to Artifact Registry.
- B. Submit a Cloud Build job to push the image.
- C. Use the pack build command with pack CLI.
- D. Include the --source flag with the gcloud run deploy CLI command.
- E. Include the --platform=kubernetes flag with the gcloud run deploy CLI command.

Answer: AC

Explanation:

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

NEW QUESTION 124

- (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 127

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 128

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 129

- (Exam Topic 2)

You are deploying your application on a Compute Engine instance that communicates with Cloud SQL. You will use Cloud SQL Proxy to allow your application to communicate to the database using the service account associated with the application's instance. You want to follow the Google-recommended best practice of providing minimum access for the role assigned to the service account. What should you do?

- A. Assign the Project Editor role.
- B. Assign the Project Owner role.
- C. Assign the Cloud SQL Client role.
- D. Assign the Cloud SQL Editor role.

Answer: C

Explanation:

Reference: <https://cloud.google.com/sql/docs/mysql/sql-proxy>

NEW QUESTION 131

- (Exam Topic 2)

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

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

Answer: D

Explanation:

<https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#file-system> Whenever possible, avoid storing service account keys on a file system. If you can't avoid storing keys on disk, make sure to restrict access to the key file, configure file access auditing, and encrypt the underlying disk.
<https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys#software-keystore> In situations where using a hardware-based key store isn't viable, use a software-based key store to manage service account keys. Similar to hardware-based options, a software-based key store lets users or applications use service account keys without revealing the private key. Software-based key store solutions can help you control key access in a fine-grained manner and can also ensure that each key access is logged.

NEW QUESTION 135

- (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 139

- (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 140

- (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 repository
- 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 141

- (Exam Topic 2)

Your company has a new security initiative that requires all data stored in Google Cloud to be encrypted by customer-managed encryption keys. You plan to use Cloud Key Management Service (KMS) to configure access to the keys. You need to follow the "separation of duties" principle and Google-recommended best practices. What should you do? (Choose two.)

- A. Provision Cloud KMS in its own project.
- B. Do not assign an owner to the Cloud KMS project.
- C. Provision Cloud KMS in the project where the keys are being used.
- D. Grant the roles/cloudkms.admin role to the owner of the project where the keys from Cloud KMS are being used.
- E. Grant an owner role for the Cloud KMS project to a different user than the owner of the project where the keys from Cloud KMS are being used.

Answer: AB

Explanation:

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

NEW QUESTION 146

- (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 151

- (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 156

- (Exam Topic 2)

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

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

Answer: C

Explanation:

<https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview>

Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to exercise as many user inputs and event handlers as possible.

NEW QUESTION 157

- (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 161

- (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 166

- (Exam Topic 2)

Your teammate has asked you to review the code below. Its purpose is to efficiently add a large number of small rows to a BigQuery table.

```
BigQuery service = BigQueryOptions.newBuilder().build().getService();

public void writeToBigQuery(Collection<Map<String, String>> rows){
    for(Map<String, String> row : rows) {
        InsertAllRequest insertRequest = InsertAllRequest.newBuilder(
            "datasetId", "tableId",
            InsertAllRequest.RowToInsert.of(row)).build();
        service.insertAll(insertRequest);
    }
}
```

Which improvement should you suggest your teammate make?

- A. Include multiple rows with each request.
- B. Perform the inserts in parallel by creating multiple threads.
- C. Write each row to a Cloud Storage object, then load into BigQuery.
- D. Write each row to a Cloud Storage object in parallel, then load into BigQuery.

Answer: B

NEW QUESTION 169

- (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 172

- (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 175

- (Exam Topic 2)

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

What should you do?

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

Answer: C

NEW QUESTION 178

- (Exam Topic 2)

You want to migrate an on-premises container running in Knative to Google Cloud. You need to make sure that the migration doesn't affect your application's deployment strategy, and you want to use a fully managed service. Which Google Cloud service should you use to deploy your container?

- A. Cloud Run
- B. Compute Engine
- C. Google Kubernetes Engine
- D. App Engine flexible environment

Answer: A

Explanation:

<https://cloud.google.com/blog/products/serverless/knative-based-cloud-run-services-are-ga>

NEW QUESTION 183

- (Exam Topic 2)

You have an application deployed in Google Kubernetes Engine (GKE). You need to update the application to make authorized requests to Google Cloud managed services. You want this to be a one-time setup, and you need to follow security best practices of auto-rotating your security keys and storing them in an encrypted store. You already created a service account with appropriate access to the Google Cloud service. What should you do next?

- A. Assign the Google Cloud service account to your GKE Pod using Workload Identity.
- B. Export the Google Cloud service account, and share it with the Pod as a Kubernetes Secret.
- C. Export the Google Cloud service account, and embed it in the source code of the application.
- D. Export the Google Cloud service account, and upload it to HashiCorp Vault to generate a dynamic service account for your application.

Answer: A

Explanation:

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

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 184

- (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 189

- (Exam Topic 2)

Your company uses Cloud Logging to manage large volumes of log data. You need to build a real-time log analysis architecture that pushes logs to a third-party application for processing. What should you do?

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

Answer: B

NEW QUESTION 192

- (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 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 194

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

Answer: B

NEW QUESTION 197

- (Exam Topic 2)

Your development team has been tasked with maintaining a .NET legacy application. The application incurs occasional changes and was recently updated. Your goal is to ensure that the application provides consistent results while moving through the CI/CD pipeline from environment to environment. You want to minimize the cost of deployment while making sure that external factors and dependencies between hosting environments are not problematic. Containers are not yet approved in your organization. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 198

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 199

- (Exam Topic 2)

You have an application deployed in production. When a new version is deployed, some issues don't arise until the application receives traffic from users in production. You want to reduce both the impact and the number of users affected.

Which deployment strategy should you use?

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

Answer: A

Explanation:

Reference: <https://thenewstack.io/deployment-strategies/>

NEW QUESTION 203

- (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 207

- (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 209

- (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 211

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 215

- (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 219

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 221

- (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 224

- (Exam Topic 2)

Your App Engine standard configuration is as follows: service: production
instance_class: B1

You want to limit the application to 5 instances. Which code snippet should you include in your configuration?

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

Answer: C

NEW QUESTION 229

- (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 232

- (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 236

- (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 237

- (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 is complete.
- D. complete.
- E. Perform a canary deployment, and test your new application periodically after the new version is deployed.

Answer: C

NEW QUESTION 240

- (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 242

- (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 customer
- H. Create a secondary index on the Array field.
- I. Create a table named Customers as a parent table
- J. Create a table named Phones, and interleave this table into the Customer table
- K. Create an index on the phone number field in the Phones table.

Answer: C

NEW QUESTION 246

- (Exam Topic 2)

You are deploying your application to a Compute Engine virtual machine instance with the Stackdriver Monitoring Agent installed. Your application is a unix process on the instance. You want to be alerted if the unix process has not run for at least 5 minutes. You are not able to change the application to generate metrics or logs. Which alert condition should you configure?

- A. Uptime check
- B. Process health
- C. Metric absence
- D. Metric threshold

Answer: B

Explanation:

Reference: <https://cloud.google.com/monitoring/alerts/concepts-indepth>

NEW QUESTION 250

- (Exam Topic 2)

Your service adds text to images that it reads from Cloud Storage. During busy times of the year, requests to Cloud Storage fail with an HTTP 429 "Too Many Requests" status code. How should you handle this error?

- A. Add a cache-control header to the objects.
- B. Request a quota increase from the GCP Console.
- C. Retry the request with a truncated exponential backoff strategy.
- D. Change the storage class of the Cloud Storage bucket to Multi-regional.

Answer: C

Explanation:

Reference: <https://developers.google.com/gmail/api/v1/reference/quota>

NEW QUESTION 255

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 256

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 260

- (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 263

- (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 265

- (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 266

- (Exam Topic 2)

You are using Cloud Build to build and test application source code stored in Cloud Source Repositories. The build process requires a build tool not available in the Cloud Build environment. What should you do?

- A. Download the binary from the internet during the build process.
- B. Build a custom cloud builder image and reference the image in your build steps.
- C. Include the binary in your Cloud Source Repositories repository and reference it in your build scripts.
- D. Ask to have the binary added to the Cloud Build environment by filing a feature request against the Cloud Build public Issue Tracker.

Answer: B

NEW QUESTION 269

- (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 270

- (Exam Topic 2)

You are planning to migrate a MySQL database to the managed Cloud SQL database for Google Cloud. You have Compute Engine virtual machine instances that will connect with this Cloud SQL instance. You do not want to whitelist IPs for the Compute Engine instances to be able to access Cloud SQL. What should you do?

- A. Enable private IP for the Cloud SQL instance.
- B. Whitelist a project to access Cloud SQL, and add Compute Engine instances in the whitelisted project.
- C. Create a role in Cloud SQL that allows access to the database from external instances, and assign the Compute Engine instances to that role.
- D. Create a CloudSQL instance on one project
- E. Create Compute engine instances in a different project. Create a VPN between these two projects to allow internal access to CloudSQL.

Answer: C

Explanation:

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

NEW QUESTION 275

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 278

- (Exam Topic 2)

Your company's development teams want to use various open source operating systems in their Docker builds. When images are created in published containers in your company's environment, you need to scan them for Common Vulnerabilities and Exposures (CVEs). The scanning process must not impact software development agility. You want to use managed services where possible. What should you do?

- A. Enable the Vulnerability scanning setting in the Container Registry.
- B. Create a Cloud Function that is triggered on a code check-in and scan the code for CVEs.
- C. Disallow the use of non-commercially supported base images in your development environment.
- D. Use Cloud Monitoring to review the output of Cloud Build to determine whether a vulnerable version has been used.

Answer: A

Explanation:

<https://cloud.google.com/container-analysis/docs/os-overview>

NEW QUESTION 280

- (Exam Topic 2)

Your API backend is running on multiple cloud providers. You want to generate reports for the network latency of your API. Which two steps should you take? (Choose two.)

- A. Use Zipkin collector to gather data.
- B. Use Fluentd agent to gather data.
- C. Use Stackdriver Trace to generate reports.
- D. Use Stackdriver Debugger to generate report.
- E. Use Stackdriver Profiler to generate report.

Answer: AC

Explanation:

<https://cloud.google.com/trace/docs/zipkin>

"receive traces from Zipkin clients and forward those traces to Cloud Trace for analysis." https://cloud.google.com/trace/docs/quickstart#analysis_reports_window

NEW QUESTION 282

- (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 283

.....

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