

Professional-Cloud-Developer Dumps

Google Certified Professional - Cloud Developer

<https://www.certleader.com/Professional-Cloud-Developer-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

HipLocal is expanding into new locations. They must capture additional data each time the application is launched in a new European country. This is causing delays in the development process due to constant schema changes and a lack of environments for conducting testing on the application changes. How should they resolve the issue while meeting the business requirements?

- A. Create new Cloud SQL instances in Europe and North America for testing and 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 2

- (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 Splitting
- E. Do not move QA traffic to the new versions if errors are found.

Answer: B

NEW QUESTION 3

- (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 4

- (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 5

- (Exam Topic 2)

Users are complaining that your Cloud Run-hosted website responds too slowly during traffic spikes. You want to provide a better user experience during traffic peaks. What should you do?

- A. Read application configuration and static data from the database on application startup.
- B. Package application configuration and static data into the application image during build time.
- C. Perform as much work as possible in the background after the response has been returned to the user.
- D. Ensure that timeout exceptions and errors cause the Cloud Run instance to exit quickly so a replacement instance can be started.

Answer: C

NEW QUESTION 6

- (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 7

- (Exam Topic 2)

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

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

Answer: D

Explanation:


Reference:

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

NEW QUESTION 8

- (Exam Topic 2)

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 9

- (Exam Topic 2)

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

Which technique should you choose?

- A. Smoke tests
- B. Stackdriver uptime checks
- C. Cloud Load Balancing - 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 10

- (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 10

- (Exam Topic 2)

Your website is deployed on Compute Engine. Your marketing team wants to test conversion rates between 3 different website designs.

Which approach should you use?

- A. Deploy the website on App Engine and use traffic splitting.
- B. Deploy the website on App Engine as three separate services.
- C. Deploy the website on Cloud Functions and use traffic splitting.
- D. Deploy the website on Cloud Functions as three separate functions.

Answer: A

Explanation:

Reference: <https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

NEW QUESTION 15

- (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 18

- (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 22

- (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 27

- (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 31

- (Exam Topic 2)

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

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

Answer: C

Explanation:

<https://cloud.google.com/shell/docs>

NEW QUESTION 32

- (Exam Topic 2)

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

- A. Package each component in a separate containe
- B. Implement readiness and liveness probes.
- C. Package the application in a single containe
- D. Use a process management tool to manage each component.
- E. Package each component in a separate containe
- F. Use a script to orchestrate the launch of the components.

- G. Package the application in a single 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 35

- (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 37

- (Exam Topic 2)

You are developing an application that will be launched on Compute Engine instances into multiple distinct projects, each corresponding to the environments in your software development process (development, QA, staging, and production). The instances in each project have the same application code but a different configuration. During deployment, each instance should receive the application's configuration based on the environment it serves. You want to minimize the number of steps to configure this flow. What should you do?

- A. When creating your instances, configure a startup script using the gcloud command to determine the project name that indicates the correct environment.
- B. In each project, configure a metadata key "environment" whose value is the environment it serve
- C. Use your deployment tool to query the instance metadata and configure the application based on the "environment" value.
- D. Deploy your chosen deployment tool on an instance in each projec
- E. Use a deployment job to retrieve the appropriate configuration file from your version control system, and apply the configuration when deploying the application on each instance.
- F. During each instance launch, configure an instance custom-metadata key named "environment" whose value is the environment the instance serve
- G. Use your deployment tool to query the instance metadata, and configure the application based on the "environment" value.

Answer: B

Explanation:

Reference: <https://cloud.google.com/compute/docs/metadata/overview>

NEW QUESTION 38

- (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 43

- (Exam Topic 2)

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

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

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

Answer: D

NEW QUESTION 46

- (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 50

- (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 52

- (Exam Topic 2)

You are using Cloud Build for your CI/CD pipeline to complete several tasks, including copying certain files to Compute Engine virtual machines. Your pipeline requires a flat file that is generated in one builder in the pipeline to be accessible by subsequent builders in the same pipeline. How should you store the file so that all the builders in the pipeline can access it?

- A. Store and retrieve the file contents using Compute Engine instance metadata.
- B. Output the file contents to a file in /workspac
- C. Read from the same /workspace file in the subsequent build step.
- D. Use gsutil to output the file contents to a Cloud Storage objec
- E. Read from the same object in the subsequent build step.
- F. Add a build argument that runs an HTTP POST via curl to a separate web server to persist the value in one builde
- G. Use an HTTP GET via curl from the subsequent build step to read the value.

Answer: B

Explanation:

<https://cloud.google.com/build/docs/build-config-file-schema>

NEW QUESTION 53

- (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 54

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 56

- (Exam Topic 2)

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

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 58

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 61

- (Exam Topic 2)

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 62

- (Exam Topic 2)

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

- A. Reboot the machine.

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

Answer: A

NEW QUESTION 64

- (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 66

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 68

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 71

- (Exam Topic 2)

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

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

Answer: D

Explanation:

<https://cloud.google.com/storage/docs/access-control/signed-urls#should-you-use>

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

NEW QUESTION 74

- (Exam Topic 2)

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

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

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

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

Answer: CE

Explanation:

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

NEW QUESTION 75

- (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 79

- (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 80

- (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 83

- (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 87

- (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 92

- (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-cluster> 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 97

- (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 102

- (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 106

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 108

- (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 112

- (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 117

- (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 cluster.
- 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 image.
- 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 startup.
- 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 122

- (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 123

- (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 124

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Professional-Cloud-Developer Exam with Our Prep Materials Via below:

<https://www.certleader.com/Professional-Cloud-Developer-dumps.html>