

Google

Exam Questions Professional-Cloud-Developer

Google Certified Professional - Cloud Developer



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)

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

Which IP strategy should they use?

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

Answer: A

NEW QUESTION 3

- (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 4

- (Exam Topic 1)

For this question, refer to the HipLocal case study.

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

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

Answer: D

Explanation:

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

NEW QUESTION 5

- (Exam Topic 1)

Which service should HipLocal use for their public APIs?

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

Answer: D

NEW QUESTION 6

- (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 7

- (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 8

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 9

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 10

- (Exam Topic 2)

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

- A. Download the file using scp, change the file, and then upload the modified version
- B. Configure and log in to the Compute Engine instance through SSH, and change the file
- C. Configure and log in to the Compute Engine instance through the serial port, and change the file
- D. Configure and log in to the Compute Engine instance using a remote desktop client, and change the file

Answer: C

Explanation:

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

NEW QUESTION 10

- (Exam Topic 2)

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

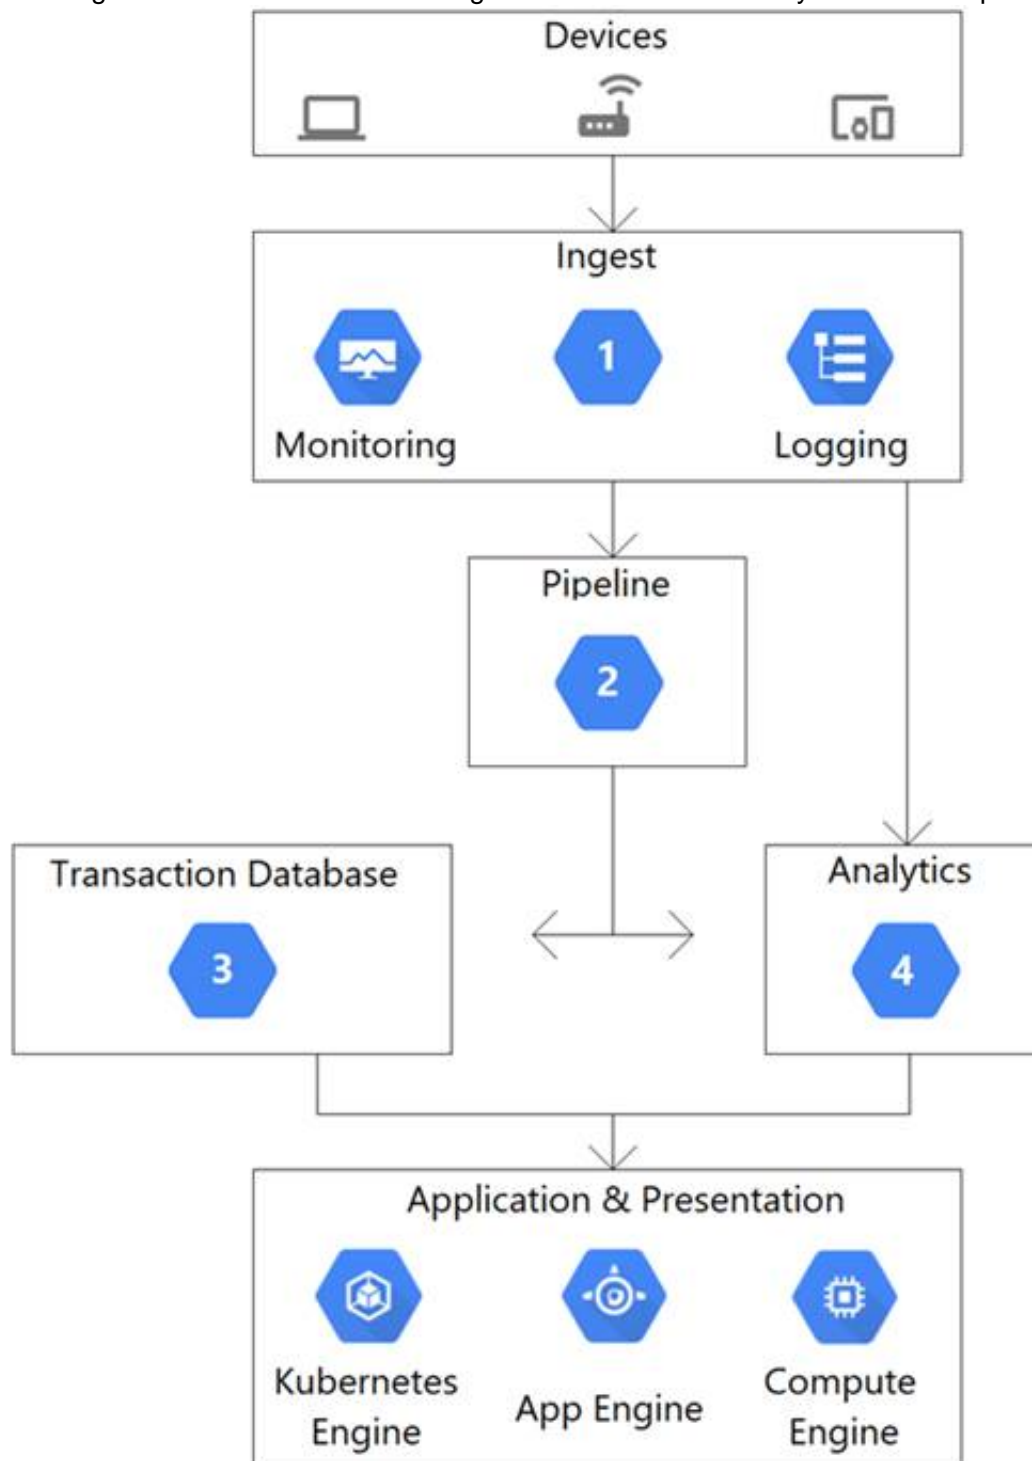
- A. Configure the application to send the file to the client as an email attachment.
- B. Generate and assign a Cloud Storage-signed URL for the file
- C. Make the URL available for the client to download.
- D. Create a temporary Cloud Storage bucket with time expiration specified, and give download 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 12

- (Exam Topic 2)

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



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

Answer: D

NEW QUESTION 15

- (Exam Topic 2)

You are using Cloud Run to host a web application. You need to securely obtain the application project ID and region where the application is running and display this information to users. You want to use the most performant approach. What should you do?

- A. Use HTTP requests to query the available metadata server at the <http://metadata.google.internal/endpoint> with the Metadata-Flavor: Google header.
- B. In the Google Cloud console, navigate to the Project Dashboard and gather configuration details. Navigate to the Cloud Run "Variables & Secrets" tab, and add the desired environment variables in Key:Value format.
- C. In the Google Cloud console, navigate to the Project Dashboard and gather configuration detail

- D. Write the application configuration information to Cloud Run's in-memory container filesystem.
- E. Make an API call to the Cloud Asset Inventory API from the application and format the request to include instance metadata.

Answer: B

NEW QUESTION 18

- (Exam Topic 2)

Your company has deployed a new API to App Engine Standard environment. During testing, the API is not behaving as expected. You want to monitor the application over time to diagnose the problem within the application code without redeploying the application. Which tool should you use?

- A. Stackdriver Trace
- B. Stackdriver Monitoring
- C. Stackdriver Debug Snapshots
- D. Stackdriver Debug Logpoints

Answer: B

Explanation:

Reference: <https://rominirani.com/gcp-stackdriver-tutorial-debug-snapshots-traces-logging-and-logpoints-1ba49e4780e6>

NEW QUESTION 22

- (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 23

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 24

- (Exam Topic 2)

You want to create “fully baked” or “golden” Compute Engine images for your application. You need to bootstrap your application to connect to the appropriate database according to the environment the application is running on (test, staging, production). What should you do?

- A. Embed the appropriate database connection string in the image
- B. Create a different image for each environment.
- C. When creating the Compute Engine instance, add a tag with the name of the database to be connected. In your application, query the Compute Engine API to pull the tags for the current instance, and use the tag to construct the appropriate database connection string.
- D. When creating the Compute Engine instance, create a metadata item with a key of “DATABASE” and a value for the appropriate database connection string
- E. In your application, read the “DATABASE” environment variable, and use the value to connect to the appropriate database.
- F. When creating the Compute Engine instance, create a metadata item with a key of “DATABASE” and a value for the appropriate database connection string
- G. In your application, query the metadata server for the “DATABASE” value, and use the value to connect to the appropriate database.

Answer: C

NEW QUESTION 29

- (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 database
- B. Deploy the application to Cloud Run.
- C. Build the application in C# with Firestore as the database

- D. Deploy the application to App Engine flexible environment.
- E. Build the application in Python with CloudSQL as the database.
- F. Deploy the application to App Engine standard environment.
- G. Build the application in Python with Firestore as the database.
- H. Deploy the application to a Compute Engine managed instance group with autoscaling.

Answer: A

NEW QUESTION 33

- (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 35

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 39

- (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 43

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 46

- (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 48

- (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 development.
- D. Send your source code updates as pull requests.
- E. Use a Cloud Storage bucket to store the source code that you need to edit.
- F. Mount the bucket to a public computer as a drive, and use a code editor to update the code.
- 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 49

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 52

- (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 API.
- 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 API.
- 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 API.
- 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 API.
- 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 55

- (Exam Topic 2)

You are creating an App Engine application that writes a file to any user's Google Drive. How should the application authenticate to the Google Drive API?

- A. With an OAuth Client ID that uses the <https://www.googleapis.com/auth/drive.file> scope to obtain an access token for each user.
- B. With an OAuth Client ID with delegated domain-wide authority.
- C. With the App Engine service account and <https://www.googleapis.com/auth/drive.file> scope that generates a signed JWT.
- D. With the App Engine service account with delegated domain-wide authority.

Answer: B

Explanation:

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

NEW QUESTION 60

- (Exam Topic 2)

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

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

Answer: A

Explanation:

<https://cloud.google.com/storage/docs/folders>

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

NEW QUESTION 65

- (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 69

- (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 70

- (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 73

- (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-bette>

NEW QUESTION 77

- (Exam Topic 2)

Your company has deployed a new API to a Compute Engine instance. During testing, the API is not behaving as expected. You want to monitor the application over 12 hours to diagnose the problem within the application code without redeploying the application. Which tool should you use?

- A. Cloud Trace
- B. Cloud Monitoring
- C. Cloud Debugger logpoints
- D. Cloud Debugger snapshots

Answer: C

Explanation:

<https://cloud.google.com/debugger/docs/using/logpoints>

Logpoints allow you to inject logging into running services without restarting or interfering with the normal function of the service

NEW QUESTION 81

- (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 86

- (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 87

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 89

- (Exam Topic 2)

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

should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 92

- (Exam Topic 2)

You want to notify on-call engineers about a service degradation in production while minimizing development time. What should you do?

- A. Use Cloud Function to monitor resources and raise alerts.
- B. Use Cloud Pub/Sub to monitor resources and raise alerts.
- C. Use Stackdriver Error Reporting to capture errors and raise alerts.
- D. Use Stackdriver Monitoring to monitor resources and raise alerts.

Answer: A

NEW QUESTION 94

- (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 99

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 102

- (Exam Topic 2)

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

```

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

```

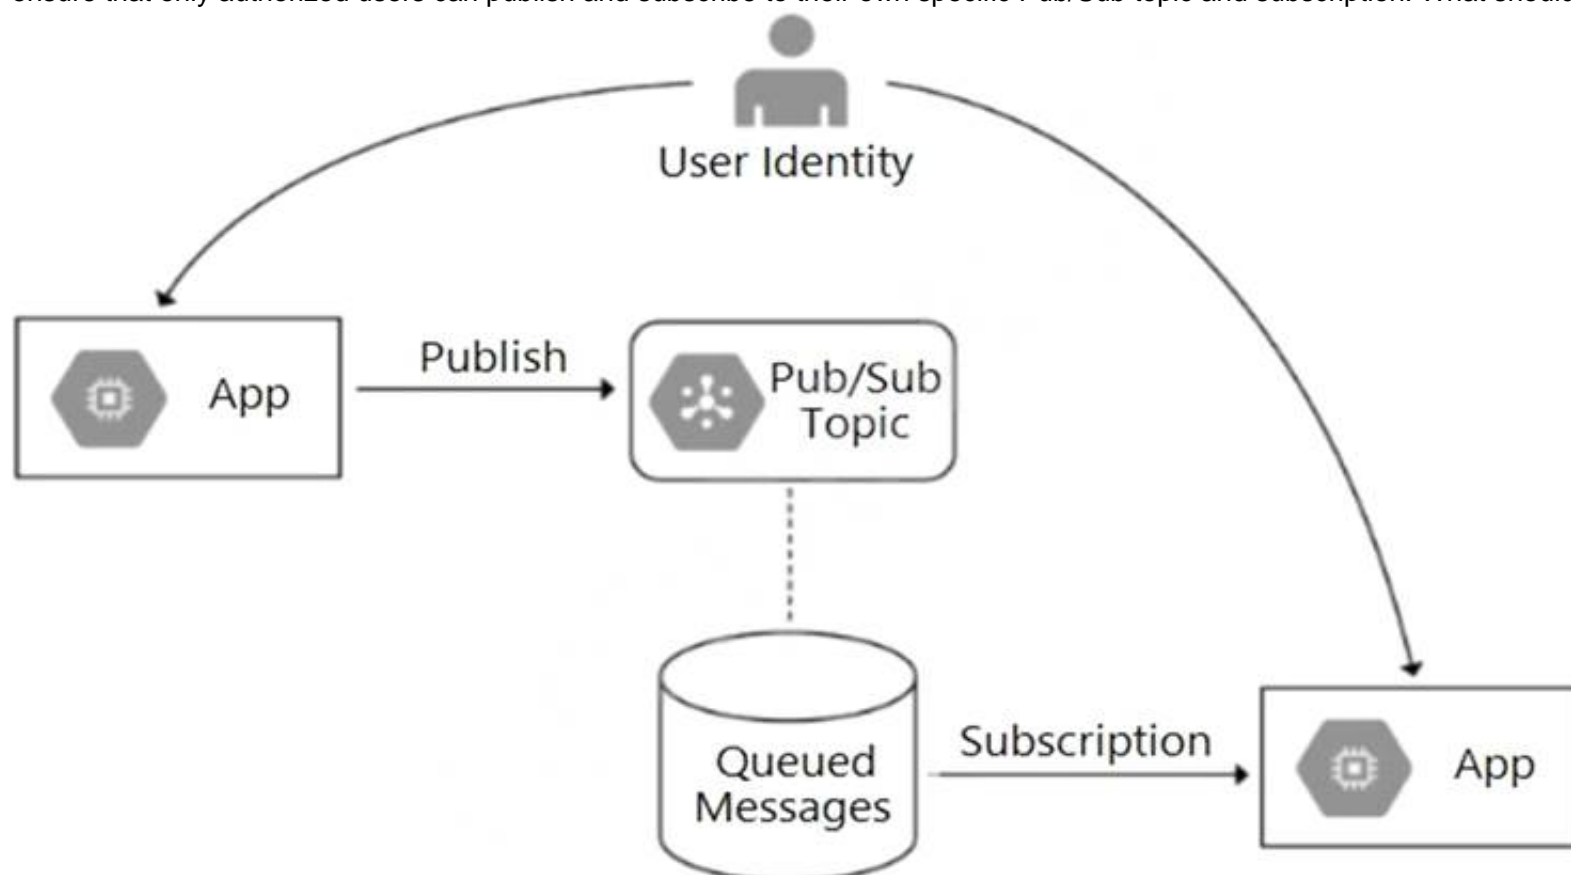
- A. Set the Deployment strategy to RollingUpdate with maxSurge set to 0, maxUnavailable set to 1.
- B. Set the Deployment strategy to RollingUpdate with maxSurge set to 1, maxUnavailable set to 0.
- C. Set the Deployment strategy to Recreate with maxSurge set to 0, maxUnavailable set to 1.
- D. Set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

Answer: D

NEW QUESTION 103

- (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 104

- (Exam Topic 2)

Your team is writing a backend application to implement the business logic for an interactive voice response (IVR) system that will support a payroll application.

The IVR system has the following technical characteristics:

- Each customer phone call is associated with a unique IVR session.
- The IVR system creates a separate persistent gRPC connection to the backend for each session.
- If the connection is interrupted, the IVR system establishes a new connection, causing a slight latency for that call.

You need to determine which compute environment should be used to deploy the backend application. Using current call data, you determine that:

- Call duration ranges from 1 to 30 minutes.
- Calls are typically made during business hours.
- There are significant spikes of calls around certain known dates (e.g., pay days), or when large payroll changes occur.

You want to minimize cost, effort, and operational overhead. Where should you deploy the backend application?

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

Answer: D

Explanation:

This page shows Cloud Run-specific details for developers who want to use gRPC to connect a Cloud Run service with other services, for example, to provide simple, high performance communication between internal microservices. You can use all gRPC types, streaming or unary, with Cloud Run.

Possible use cases include:

Communication between internal microservices.

High loads of data (gRPC uses protocol buffers, which are up to seven times faster than REST calls). Only a simple service definition is needed, you don't want to write a full client library.

Use streaming gRPCs in your gRPC server to build more responsive applications and APIs. <https://cloud.google.com/run/docs/tutorials/secure-services#:~:text=The%20backend%20service%20is%20priva>

NEW QUESTION 107

- (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 111

- (Exam Topic 2)

Before promoting your new application code to production, you want to conduct testing across a variety of different users. Although this plan is risky, you want to test the new version of the application with production users and you want to control which users are forwarded to the new version of the application based on their operating system. If bugs are discovered in the new version, you want to roll back the newly deployed version of the application as quickly as possible.

What should you do?

- A. Deploy your application on Cloud Ru
- B. Use traffic splitting to direct a subset of user traffic to the new version based on the revision tag.
- C. Deploy your application on Google Kubernetes Engine with Anthos Service Mes
- D. Use traffic splitting to direct a subset of user traffic to the new version based on the user-agent header.
- E. Deploy your application on App Engin
- F. Use traffic splitting to direct a subset of user traffic to the new version based on the IP address.
- G. Deploy your application on Compute Engin
- H. Use Traffic Director to direct a subset of user traffic to the new version based on predefined weights.

Answer: B

NEW QUESTION 115

- (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 117

- (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 121

- (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 125

- (Exam Topic 2)

You migrated some of your applications to Google Cloud. You are using a legacy monitoring platform deployed on-premises for both on-premises and cloud-deployed applications. You discover that your notification system is responding slowly to time-critical problems in the cloud applications. What should you do?

- A. Replace your monitoring platform with Cloud Monitoring.
- B. Install the Cloud Monitoring agent on your Compute Engine instances.
- C. Migrate some traffic back to your old platform
- D. Perform A/B testing on the two platforms concurrently.
- E. Use Cloud Logging and Cloud Monitoring to capture logs, monitor, and send alert
- F. Send them to your existing platform.

Answer: D

NEW QUESTION 129

- (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 database
- 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 availability
- 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 availability
- 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 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 132

- (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 136

- (Exam Topic 2)

You are monitoring a web application that is written in Go and deployed in Google Kubernetes Engine. You notice an increase in CPU and memory utilization. You need to determine which source code is consuming the most CPU and memory resources. What should you do?

- A. Download, install, and start the Snapshot Debugger agent in your V
- B. Take debug snapshots of the functions that take the longest tim
- C. Review the call stack frame, and identify the local variables at that level in the stack.
- D. Import the Cloud Profiler package into your application, and initialize the Profiler agen
- E. Review the generated flame graph in the Google Cloud console to identify time-intensive functions.
- F. Import OpenTelemetry and Trace export packages into your application, and create the trace provider. Review the latency data for your application on the Trace overview page, and identify where bottlenecks are occurring.
- G. Create a Cloud Logging query that gathers the web application's log
- H. Write a Python script that calculates the difference between the timestamps from the beginning and the end of the application's longest functions to identity time-intensive functions.

Answer: B

NEW QUESTION 141

- (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 142

- (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 147

- (Exam Topic 2)

You are evaluating developer tools to help drive Google Kubernetes Engine adoption and integration with your development environment, which includes VS Code and IntelliJ. What should you do?

- A. Use Cloud Code to develop applications.
- B. Use the Cloud Shell integrated Code Editor to edit code and configuration files.
- C. Use a Cloud Notebook instance to ingest and process data and deploy models.
- D. Use Cloud Shell to manage your infrastructure and applications from the command line.

Answer: A

Explanation:

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

NEW QUESTION 150

- (Exam Topic 2)

You are developing a microservice-based application that will run on Google Kubernetes Engine (GKE). Some of the services need to access different Google Cloud APIs. How should you set up authentication of these services in the cluster following Google-recommended best practices? (Choose two.)

- A. Use the service account attached to the GKE node.
- B. Enable Workload Identity in the cluster via the gcloud command-line tool.
- C. Access the Google service account keys from a secret management service.
- D. Store the Google service account keys in a central secret management service.
- E. Use gcloud to bind the Kubernetes service account and the Google service account using roles/iam.workloadIdentity.

Answer: BE

Explanation:

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

NEW QUESTION 153

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 156

- (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 159

- (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 163

- (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 167

- (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 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 175

- (Exam Topic 2)

You are working on a social media application. You plan to add a feature that allows users to upload images. These images will be 2 MB – 1 GB in size. You want to minimize their infrastructure operations overhead for this feature. What should you do?

- A. Change the application to accept images directly and store them in the database that stores other user information.
- B. Change the application to create signed URLs for Cloud Storage
- C. Transfer these signed URLs to the client application to upload images to Cloud Storage.
- D. Set up a web server on GCP to accept user images and create a file store to keep uploaded file
- E. Change the application to retrieve images from the file store.
- F. Create a separate bucket for each user in Cloud Storage
- G. Assign a separate service account to allow write access on each bucket
- H. Transfer service account credentials to the client application based on user information
- I. The application uses this service account to upload images to Cloud Storage.

Answer: B

Explanation:

Reference:

<https://cloud.google.com/blog/products/storage-data-transfer/uploading-images-directly-to-cloud-storage-by-usi>

NEW QUESTION 178

- (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 179

- (Exam Topic 2)

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

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

Answer: AD

Explanation:

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

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

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

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

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

NEW QUESTION 182

- (Exam Topic 2)

You are using Cloud Build to promote a Docker image to Development, Test, and Production environments. You need to ensure that the same Docker image is deployed to each of these environments. How should you identify the Docker image in your build?

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

Answer: D

NEW QUESTION 184

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 187

- (Exam Topic 2)

You are designing a deployment technique for your new applications on Google Cloud. As part of your deployment planning, you want to use live traffic to gather

performance metrics for both new and existing applications. You need to test against the full production load prior to launch. What should you do?

- A. Use canary deployment
- B. Use blue/green deployment
- C. Use rolling updates deployment
- D. Use A/B testing with traffic mirroring during deployment

Answer: A

Explanation:

Reference: <https://cloud.google.com/architecture/application-deployment-and-testing-strategies>

NEW QUESTION 188

- (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 191

- (Exam Topic 2)

You are in the final stage of migrating an on-premises data center to Google Cloud. You are quickly approaching your deadline, and discover that a web API is running on a server slated for decommissioning. You need to recommend a solution to modernize this API while migrating to Google Cloud. The modernized web API must meet the following requirements:

- Autoscales during high traffic periods at the end of each month
- Written in Python 3.x
- Developers must be able to rapidly deploy new versions in response to frequent code changes

You want to minimize cost, effort, and operational overhead of this migration. What should you do?

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

Answer: B

Explanation:

<https://cloud.google.com/appengine/docs/standard>

NEW QUESTION 195

- (Exam Topic 2)

You want to view the memory usage of your application deployed on Compute Engine. What should you do?

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

Answer: C

Explanation:

Reference:

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

NEW QUESTION 200

- (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 205

- (Exam Topic 2)

You are developing an application that will handle requests from end users. You need to secure a Cloud Function called by the application to allow authorized end users to authenticate to the function via the application while restricting access to unauthorized users. You will integrate Google Sign-In as part of the solution and want to follow Google-recommended best practices. What should you do?

- A. Deploy from a source code repository and grant users the roles/cloudfunctions.viewer role.
- B. Deploy from a source code repository and grant users the roles/cloudfunctions.invoker role
- C. Deploy from your local machine using gcloud and grant users the roles/cloudfunctions.admin role
- D. Deploy from your local machine using gcloud and grant users the roles/cloudfunctions.developer role

Answer: C

NEW QUESTION 209

- (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 211

- (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 214

- (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 Run
- B. Use revisions to separate the environments.
- C. Use Cloud Build to deploy the application as a new Compute Engine image for each build
- 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 cluster
- 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 218

- (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 221

- (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 226

- (Exam Topic 2)

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

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

Answer: D

Explanation:

<https://cloud.google.com/profiler/docs/about-profiler>

Cloud Profiler is a statistical, low-overhead profiler that continuously gathers CPU usage and

memory-allocation information from your production applications. It attributes that information to the source code that generated it, helping you identify the parts of your application that are consuming the most resources, and otherwise illuminating your applications performance characteristics.

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

NEW QUESTION 229

- (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 231

- (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 232

- (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 236

- (Exam Topic 2)

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

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

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

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

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

Answer: AC

Explanation:

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

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

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

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

NEW QUESTION 240

- (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 242

- (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 246

- (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 249

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 253

- (Exam Topic 2)

Your company has created an application that uploads a report to a Cloud Storage bucket. When the report is uploaded to the bucket, you want to publish a message to a Cloud Pub/Sub topic. You want to implement a solution that will take a small amount of effort to implement. What should you do?

- A. Configure the Cloud Storage bucket to trigger Cloud Pub/Sub notifications when objects are modified.
- B. Create an App Engine application to receive the file; when it is received, publish a message to the Cloud Pub/Sub topic.
- C. Create a Cloud Function that is triggered by the Cloud Storage bucket.
- D. In the Cloud Function, publish a message to the Cloud Pub/Sub topic.
- E. Create an application deployed in a Google Kubernetes Engine cluster to receive the file; when it is received, publish a message to the Cloud Pub/Sub topic.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/pubsub-notifications>

NEW QUESTION 258

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 262

- (Exam Topic 2)

Your organization has recently begun an initiative to replatform their legacy applications onto Google Kubernetes Engine. You need to decompose a monolithic application into microservices. Multiple instances have read and write access to a configuration file, which is stored on a shared file system. You want to minimize the effort required to manage this transition, and you want to avoid rewriting the application code. What should you do?

- A. Create a new Cloud Storage bucket, and mount it via FUSE in the container.
- B. Create a new persistent disk, and mount the volume as a shared PersistentVolume.
- C. Create a new Filestore instance, and mount the volume as an NFS PersistentVolume.
- D. Create a new ConfigMap and volumeMount to store the contents of the configuration file.

Answer: D

Explanation:

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

ConfigMaps bind non-sensitive configuration artifacts such as configuration files, command-line arguments, and environment variables to your Pod containers and system components at runtime.

A ConfigMap separates your configurations from your Pod and components, which helps keep your workloads portable. This makes their configurations easier to change and manage, and prevents hardcoding configuration data to Pod specifications.

NEW QUESTION 265

- (Exam Topic 2)

You are designing a schema for a Cloud Spanner customer database. You want to store a phone number array field in a customer table. You also want to allow users to search customers by phone number. How should you design this schema?

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

Answer: C

NEW QUESTION 267

- (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 270

- (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 273

- (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 276

- (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 277

- (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 281

- (Exam Topic 2)

You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

- A. Deploy each microservice as a Deployment
- B. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.
- C. Deploy each microservice as a Deployment
- D. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.
- E. Deploy each microservice as a Pod
- F. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.
- G. Deploy each microservice as a Pod
- H. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

Answer: A

NEW QUESTION 285

- (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 289

- (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 294

- (Exam Topic 2)

You are a SaaS provider deploying dedicated blogging software to customers in your Google Kubernetes Engine (GKE) cluster. You want to configure a secure multi-tenant platform to ensure that each customer has access to only their own blog and can't affect the workloads of other customers. What should you do?

- A. Enable Application-layer Secrets on the GKE cluster to protect the cluster.
B. Deploy a namespace per tenant and use Network Policies in each blog deployment.
C. Use GKE Audit Logging to identify malicious containers and delete them on discovery.
D. Build a custom image of the blogging software and use Binary Authorization to prevent untrusted image deployments.

Answer: B

Explanation:

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

NEW QUESTION 295

- (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 300

- (Exam Topic 2)

You are building a CI/CD pipeline that consists of a version control system, Cloud Build, and Container Registry. Each time a new tag is pushed to the repository, a Cloud Build job is triggered, which runs unit tests on the new code builds a new Docker container image, and pushes it into Container Registry. The last step of your pipeline should deploy the new container to your production Google Kubernetes Engine (GKE) cluster. You need to select a tool and deployment strategy that meets the following requirements:

- Zero downtime is incurred
 - Testing is fully automated
 - Allows for testing before being rolled out to users
 - Can quickly rollback if needed
- What should you do?

- A. Trigger a Spinnaker pipeline configured as an A/B test of your new code and, if it is successful, deploy the container to production.
B. Trigger a Spinnaker pipeline configured as a canary test of your new code and, if it is successful, deploy the container to production.
C. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a canary test.
D. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a shadow test.

Answer: D

Explanation:

https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#perform_a_shad With a shadow test, you test the new version of your application by mirroring user traffic from the current application version without impacting the user requests.

NEW QUESTION 302

- (Exam Topic 2)

You are building an API that will be used by Android and iOS apps The API must:

- Support HTTPS
 - Minimize bandwidth cost
 - Integrate easily with mobile apps
- Which API architecture should you use?

- A. RESTful APIs
B. MQTT for APIs
C. gRPC-based APIs
D. SOAP-based APIs

Answer: A

Explanation:

Reference: <https://www.devteam.space/blog/how-to-build-restful-api-for-your-mobile-app/>

NEW QUESTION 306

- (Exam Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 307

- (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 312

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 317

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 318

- (Exam Topic 2)

Your company just experienced a Google Kubernetes Engine (GKE) API outage due to a zone failure. You want to deploy a highly available GKE architecture that minimizes service interruption to users in the event of a future zone failure. What should you do?

- A. Deploy Zonal clusters
- B. Deploy Regional clusters
- C. Deploy Multi-Zone clusters

D. Deploy GKE on-premises clusters

Answer: B

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/types-of-clusters#regional_clusters

A regional cluster has multiple replicas of the control plane, running in multiple zones within a given region. Nodes in a regional cluster can run in multiple zones or a single zone depending on the configured node locations. By default, GKE replicates each node pool across three zones of the control plane's region. When you create a cluster or when you add a new node pool, you can change the default configuration by specifying the zone(s) in which the cluster's nodes run. All zones must be within the same region as the control plane.

NEW QUESTION 319

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 322

- (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 323

- (Exam Topic 2)

You are developing an application that will allow users to read and post comments on news articles. You want to configure your application to store and display user-submitted comments using Firestore. How should you design the schema to support an unknown number of comments and articles?

- A. Store each comment in a subcollection of the article.
- B. Add each comment to an array property on the article.
- C. Store each comment in a document, and add the comment's key to an array property on the article.
- D. Store each comment in a document, and add the comment's key to an array property on the user profile.

Answer: D

NEW QUESTION 328

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 330

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Professional-Cloud-Developer Practice Exam Features:

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

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