Google

Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

   All examinations will be up to date.

* 24/7 Quality Support

   We will provide service round the clock.

* 100% Pass Rate

   Our guarantee that you will pass the exam.

* Unique Gurantee

   If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
You use Cloud Build to build your application. You want to reduce the build time while minimizing cost and development effort. What should you do?

A. Use Cloud Storage to cache intermediate artifacts.
B. Run multiple Jenkins agents to parallelize the build.
C. Use multiple smaller build steps to minimize execution time.
D. Use larger Cloud Build virtual machines (VMs) by using the machine-type option.

Answer: C

Explanation:
https://cloud.google.com/storage/docs/best-practices
https://cloud.google.com/build/docs/speeding-up-builds#caching_directories_with_google_cloud_storage Caching directories with Google Cloud Storage To increase the speed of a build, reuse the results from a
previous build. You can copy the results of a previous build to a Google Cloud Storage bucket, use the results for faster calculation, and then copy the new results back to the bucket. Use this method when your build takes a long time and produces a small number of files that does not take time to copy to and from Google Cloud Storage.
upvoted 2 times

NEW QUESTION 2
Your application images are built using Cloud Build and pushed to Google Container Registry (GCR). You want to be able to specify a particular version of your application for deployment based on the release version tagged in source control. What should you do when you push the image?

A. Reference the image digest in the source control tag.
B. Supply the source control tag as a parameter within the image name.
C. Use Cloud Build to include the release version tag in the application image.
D. Use GCR digest versioning to match the image to the tag in source control.

Answer: B

Explanation:
https://cloud.google.com/container-registry/docs/pushing-and-pulling

NEW QUESTION 3
You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. All PII entries begin with the text userinfo. You want to capture these log entries in a secure location for later review and prevent them from leaking to Stackdriver Logging. What should you do?

A. Create a basic log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.
B. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, and then copy the entries to a Cloud Storage bucket.
C. Create an advanced log filter matching userinfo, configure a log export in the Stackdriver console with Cloud Storage as a sink, and then configure a tog exclusion with userinfo as a filter.
D. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, create an advanced log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.

Answer: B

Explanation:
https://medium.com/google-cloud/fluentd-filter-plugin-for-google-cloud-data-loss-prevention-api-42bbb1308e7

NEW QUESTION 4
Some of your production services are running in Google Kubernetes Engine (GKE) in the eu-west-1 region. Your build system runs in the us-west-1 region. You want to push the container images from your build system to a scalable registry to maximize the bandwidth for transferring the images to the cluster. What should you do?

A. Push the images to Google Container Registry (GCR) using the gcr.io hostname.
B. Push the images to Google Container Registry (GCR) using the us.gcr.io hostname.
C. Push the images to Google Container Registry (GCR) using the eu.gcr.io hostname.
D. Push the images to a private image registry running on a Compute Engine instance in the eu-west-1 region.

Answer: C

Explanation:
Hostname Storage location gcr.io Stores images in data centers in the United States asia.gcr.io Stores images in data centers in Asia eu.gcr.io Stores images in data centers within member states of the European Union us.gcr.io Stores images in data centers in the United States

NEW QUESTION 5
You support a high-traffic web application that runs on Google Cloud Platform (GCP). You need to measure application reliability from a user perspective without making any engineering changes to it. What should you do?
Choose 2 answers

A. Review current application metrics and add new ones as needed.
B. Modify the code to capture additional information for user interaction.
C. Analyze the web proxy logs only and capture response time of each request.
D. Create new synthetic clients to simulate a user journey using the application.
E. Use current and historic Request Logs to trace customer interaction with the application.

Answer: CE

Explanation:
https://cloud.google.com/architecture/adopting-slos?hl=en

NEW QUESTION 6
You currently store the virtual machine (VM) utilization logs in Stackdriver. You need to provide an easy-to-share interactive VM utilization dashboard that is updated in real time and contains information aggregated on a quarterly basis. You want to use Google Cloud Platform solutions. What should you do?

A. * 1. Export VM utilization logs from Stackdriver to BigQuery.* 2. Create a dashboard in Data Studio.* 3. Share the dashboard with your stakeholders.
B. * 1. Export VM utilization logs from Stackdriver to Cloud Pub/Sub.* 2. From Cloud Pub/Sub, send the logs to a Security Information and Event Management (SIEM) system.* 3. Build the dashboards in the SIEM system and share with your stakeholders.
C. * 1. Export VM utilization logs (rom Stackdriver to BigQuery.* 2. From BigQuer
D. export the logs to a CSV file.* 3. Import the CSV file into Google Sheets.* 4. Build a dashboard in Google Sheets and share it with your stakeholders.
E. * 1. Export VM utilization logs from Stackdriver to a Cloud Storage bucket.* 2. Enable the Cloud Storage API to pull the logs programmatically.* 3. Build a custom data visualization application.* 4. Display the pulled logs in a custom dashboard.

Answer: A

NEW QUESTION 7
You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. You want to prevent these fields from being written in new log entries as quickly as possible. What should you do?

A. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
B. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.
C. Wait for the application developers to patch the application, and then verify that the log entries are no longer exposing PII.
D. Stage log entries to Cloud Storage, and then trigger a Cloud Function to remove the fields and write the entries to Stackdriver via the Stackdriver Logging API.

Answer: A

NEW QUESTION 8
You support a service with a well-defined Service Level Objective (SLO). Over the previous 6 months, your service has consistently met its SLO and customer satisfaction has been consistently high. Most of your service's operations tasks are automated and few repetitive tasks occur frequently. You want to optimize the balance between reliability and deployment velocity while following site reliability engineering best practices. What should you do? (Choose two.)

A. Make the service's SLO more strict.
B. Increase the service's deployment velocity and/or risk.
C. Shift engineering time to other services that need more reliability.
D. Get the product team to prioritize reliability work over new features.
E. Change the implementation of your Service Level Indicators (SLIs) to increase coverage.

Answer: BC

Explanation:
(https://sre.google/workbook/implementing-slos/#slo-decision-matrix)

NEW QUESTION 9
You are responsible for creating and modifying the Terraform templates that define your Infrastructure. Because two new engineers will also be working on the same code, you need to define a process and adopt a tool that will prevent you from overwriting each other's code. You also want to ensure that you capture all updates in the latest version. What should you do?

A. • Store your code in a Git-based version control system.• Establish a process that allows developers to merge their own changes at the end of each day.• Package and upload code lo a versioned Cloud Storage bucket as the latest master version.
B. • Store your code in a Git-based version control system.• Establish a process that includes code reviews by peers and unit testing to ensure integrity and functionality before integration of code.• Establish a process where the fully integrated code in the repository becomes the latest master version.
C. • Store your code as text files in Google Drive in a defined folder structure that organizes the files.• At the end of each da
D. confirm that all changes have been captured in the files within the folder structure.• Rename the folder structure with a predefined naming convention that increments the version.
E. • Store your code as text files in Google Drive in a defined folder structure that organizes the files.• At the end of each day, confirm that all changes have been captured in the files within the folder structure and create a new .zip archive with a predefined naming convention.• Upload the .zip archive to a versioned Cloud Storage bucket and accept it as the latest version.

Answer: B

NEW QUESTION 10
Your application services run in Google Kubernetes Engine (GKE). You want to make sure that only images from your centrally-managed Google Container Registry (GCR) image registry in the altostrat-images project can be deployed to the cluster while minimizing development time. What should you do?

A. Create a custom builder for Cloud Build that will only push images to gcr.io/altostrat-images.
B. Use a Binary Authorization policy that includes the whitelist name pattern gcr.io/attostrat-images/.
C. Add logic to the deployment pipeline to check that all manifests contain only images from gcr.io/altostrat-images.
D. Add a tag to each image in gcr.io/altostrat-images and check that this tag is present when the image is deployed.

Answer: B

NEW QUESTION 10

You are responsible for the reliability of a high-volume enterprise application. A large number of users report that an important subset of the application's functionality – a data intensive reporting feature – is consistently failing with an HTTP 500 error. When you investigate your application's dashboards, you notice a strong correlation between the failures and a metric that represents the size of an internal queue used for generating reports. You trace the failures to a reporting backend that is experiencing high I/O wait times. You quickly fix the issue by resizing the backend's persistent disk (PD). How you need to create an availability Service Level Indicator (SLI) for the report generation feature. How would you define it?

A. As the I/O wait times aggregated across all report generation backends
B. As the proportion of report generation requests that result in a successful response
C. As the application's report generation queue size compared to a known-good threshold
D. As the reporting backend PD throughout capacity compared to a known-good threshold

Answer: B

Explanation:
According to SRE Workbook, one of potential SLI is as below:
* Type of service: Request-driven
* Type of SLI: Availability
* Description: The proportion of requests that resulted in a successful response. https://sre.google/workbook/implementing-slos/

NEW QUESTION 14
You support a stateless web-based API that is deployed on a single Compute Engine instance in the
europe-west2-a zone . The Service Level Indicator (SLI) for service availability is below the specified Service Level Objective (SLO). A postmortem has revealed that requests to the API regularly time out. The time outs are due to the API having a high number of requests and running out memory. You want to improve service availability. What should you do?

A. Change the specified SLO to match the measured SLI.
B. Move the service to higher-specification compute instances with more memory.
C. Set up additional service instances in other zones and load balance the traffic between all instances.
D. Set up additional service instances in other zones and use them as a failover in case the primary instance is unavailable.

Answer: C

NEW QUESTION 17
You are running an experiment to see whether your users like a new feature of a web application. Shortly after deploying the feature as a canary release, you receive a spike in the number of 500 errors sent to users, and your monitoring reports show increased latency. You want to quickly minimize the negative impact on users.
What should you do first?

A. Roll back the experimental canary release.
B. Start monitoring latency, traffic, errors, and saturation.
C. Record data for the postmortem document of the incident.
D. Trace the origin of 500 errors and the root cause of increased latency.

Answer: A

NEW QUESTION 18
Your team is designing a new application for deployment both inside and outside Google Cloud Platform (GCP). You need to collect detailed metrics such as system resource utilization. You want to use centralized GCP services while minimizing the amount of work required to set up this collection system. What should you do?

A. Import the Stackdriver Profiler package, and configure it to relay function timing data to Stackdriver for further analysis.
B. Import the Stackdriver Debugger package, and configure the application to emit debug messages with timing information.
C. Instrument the code using a timing library, and publish the metrics via a health check endpoint that is scraped by Stackdriver.
D. Install an Application Performance Monitoring (APM) tool in both locations, and configure an export to a central data storage location for analysis.

Answer: A

NEW QUESTION 20
You deploy a new release of an internal application during a weekend maintenance window when there is minimal user traffic. After the window ends, you learn that one of the new features isn't working as expected in the production environment. After an extended outage, you roll back the new release and deploy a fix. You want to modify your release process to reduce the mean time to recovery so you can avoid extended outages in the future. What should you do?
Choose 2 answers

A. Before merging new code, require 2 different peers to review the code changes.
B. Adopt the blue/green deployment strategy when releasing new code via a CD server.
C. Integrate a code linting tool to validate coding standards before any code is accepted into the repository.
D. Require developers to run automated integration tests on their local development environments before release.
E. Configure a CI serve
F. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.

Answer: BE

NEW QUESTION 23
You are managing an application that exposes an HTTP endpoint without using a load balancer. The latency of the HTTP responses is important for the user experience. You want to understand what HTTP latencies all of your users are experiencing. You use Stackdriver Monitoring. What should you do?

A. • In your application, create a metric with a metricKind set to DELTA and a valueType set to DOUBLE.• In Stackdriver's Metrics Explorer, use a Slacked Bar graph to visualize the metric.

B. • In your application, create a metric with a metricKind set to CUMULATIVE and a valueType set to DOUBLE.• In Stackdriver's Metrics Explorer, use a Line graph to visualize the metric.
C. • In your application, create a metric with a metricKind set to gauge and a valueType set to distribution.• In Stackdriver's Metrics Explorer, use a Heatmap graph to visualize the metric.
D. • In your application, create a metric with a metricKin
E. set toMETRIc_KIND_UNSPECIFIEDanda valueType set to INT64.• In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Answer: C

Explanation:
 https://sre.google/workbook/implementing-slos/ https://cloud.google.com/architecture/adopting-slos/
Latency is commonly measured as a distribution. Given a distribution, you can measure various percentiles.
For example, you might measure the number of requests that are slower than the historical 99th percentile.

NEW QUESTION 27
You need to deploy a new service to production. The service needs to automatically scale using a Managed Instance Group (MIG) and should be deployed over multiple regions. The service needs a large number of resources for each instance and you need to plan for capacity. What should you do?

A. Use the n1-highcpu-96 machine type in the configuration of the MIG.
B. Monitor results of Stackdriver Trace to determine the required amount of resources.
C. Validate that the resource requirements are within the available quota limits of each region.
D. Deploy the service in one region and use a global load balancer to route traffic to this region.

Answer: C

Explanation:
 https://cloud.google.com/compute/quotas#understanding_quotas https://cloud.google.com/compute/quotas

NEW QUESTION 32
......

Relate Links

100% Pass Your Professional-Cloud-DevOps-Engineer Exam with Exambible Prep Materials

https://www.exambible.com/Professional-Cloud-DevOps-Engineer-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/