PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Q&As

Professional Cloud DevOps Engineer

Pass Google PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/professional-cloud-devops-engineer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

QUESTION 1

You support a trading application written in Python and hosted on App Engine flexible environment. You want to customize the error information being sent to Stackdriver Error Reporting. What should you do?

A. Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.

B. Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.

C. Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.

D. Use the Stackdriver Error Reporting API to write errors from your application to ReportedErrorEvent, and then generate log entries with properly formatted error messages in Stackdriver Logging.

Correct Answer: D

If you\\'re using the Error Reporting API, you can report error events from your application by writing them to ReportedErrorEvent. Doing this generates log entries with properly formatted error messages in Cloud Logging. https://cloud.google.com/error-reporting/docs/formatting-error-messages

---

QUESTION 2

You have a pool of application servers running on Compute Engine. You need to provide a secure solution that requires the least amount of configuration and allows developers to easily access application logs for troubleshooting. How would you implement the solution on GCP?

A. 1. Deploy the Stackdriver logging agent to the application servers.

2. Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.

B. 1. Deploy the Stackdriver logging agent to the application servers.

2. Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.

C. 1. Deploy the Stackdriver monitoring agent to the application servers.

2. Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.

D. 1. Install the gsutil command line tool on your application servers.

2.

Write a script using gsutil to upload your application log to a Cloud Storage bucket, and then schedule it to run via cron every 5 minutes.

3.

Give the developers the IAM Object Viewer access to view the logs in the specified bucket.

Correct Answer: A

QUESTION 3

You support a web application that runs on App Engine and uses CloudSQL and Cloud Storage for data storage. After a short spike in website traffic, you notice a big increase in latency for all user requests, increase in CPU use, and the

number of processes running the application. Initial troubleshooting reveals:

After the initial spike in traffic, load levels returned to normal but users still experience high latency.

Requests for content from the CloudSQL database and images from Cloud Storage show the same high latency.

No changes were made to the website around the time the latency increased.

There is no increase in the number of errors to the users.

You expect another spike in website traffic in the coming days and want to make sure users don\\'t experience latency. What should you do?

A. Upgrade the GCS buckets to Multi-Regional.

B. Enable high availability on the CloudSQL instances.

C. Move the application from App Engine to Compute Engine.

D. Modify the App Engine configuration to have additional idle instances.

Correct Answer: D

Scaling App Engine scales the number of instances automatically in response to processing volume. This scaling factors in the automatic_scaling settings that are provided on a per-version basis in the configuration file. A service with basic scaling is configured by setting the maximum number of instances in the max_instances parameter of the basic_scaling setting. The number of live instances scales with the processing volume. You configure the number of instances of each version in that service\\'s configuration file. The number of instances usually corresponds to the size of a dataset being held in memory or the desired throughput for offline work. You can adjust the number of instances of a manually-scaled version very quickly, without stopping instances that are currently running, using the Modules API set_num_instances function.

https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed

QUESTION 4

You recently deployed your application in Google Kubernetes Engine (GKE) and now need to release a new version of the application. You need the ability to instantly roll back to the previous version of the application in case there are issues with the new version. Which deployment model should you use?

A. Perform a rolling deployment, and test your new application after the deployment is complete.

B. Perform A/B testing, and test your application periodically after the deployment is complete.

C. Perform a canary deployment, and test your new application periodically after the new version is deployed.

D. Perform a blue/green deployment, and test your new application after the deployment is complete.

Correct Answer: D

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

QUESTION 5

Your company operates in a highly regulated domain that requires you to store all organization logs for seven years. You want to minimize logging infrastructure complexity by using managed services. You need to avoid any future loss of log capture or stored logs due to misconfiguration or human error. What should you do?

A. Use Cloud Logging to configure an aggregated sink at the organization level to export all logs into a BigQuery dataset.

B. Use Cloud Logging to configure an aggregated sink at the organization level to export all logs into Cloud Storage with a seven-year retention policy and Bucket Lock.

C. Use Cloud Logging to configure an export sink at each project level to export all logs into a BigQuery dataset

D. Use Cloud Logging to configure an export sink at each project level to export all logs into Cloud Storage with a seven-year retention policy and Bucket Lock.

Correct Answer: B

QUESTION 6

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.

Correct Answer: A

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.

QUESTION 7

Your team is designing a new application for deployment into Google Kubernetes Engine (GKE). You need to set up

monitoring to collect and aggregate various application-level metrics in a centralized location. You want to use Google Cloud Platform services while minimizing the amount of work required to set up monitoring. What should you do?

A. Publish various metrics from the application directly to the Stackdriver Monitoring API, and then observe these custom metrics in Stackdriver.

B. Install the Cloud Pub/Sub client libraries, push various metrics from the application to various topics, and then observe the aggregated metrics in Stackdriver.

C. Install the OpenTelemetry client libraries in the application, configure Stackdriver as the export destination for the metrics, and then observe the application\\'s metrics in Stackdriver.

D. Emit all metrics in the form of application-specific log messages, pass these messages from the containers to the Stackdriver logging collector, and then observe metrics in Stackdriver.

Correct Answer: A

https://cloud.google.com/trace/docs/setup

QUESTION 8

You are working with a government agency that requires you to archive application logs for seven years. You need to configure Stackdriver to export and store the logs while minimizing costs of storage. What should you do?

A. Create a Cloud Storage bucket and develop your application to send logs directly to the bucket.

B. Develop an App Engine application that pulls the logs from Stackdriver and saves them in BigQuery.

C. Create an export in Stackdriver and configure Cloud Pub/Sub to store logs in permanent storage for seven years.

D. Create a sink in Stackdriver, name it, create a bucket on Cloud Storage for storing archived logs, and then select the bucket as the log export destination.

Correct Answer: D

https://cloud.google.com/logging/docs/routing/overview

QUESTION 9

You support an application deployed on Compute Engine. The application connects to a Cloud SQL instance to store and retrieve data. After an update to the application, users report errors showing database timeout messages. The number of concurrent active users remained stable. You need to find the most probable cause of the database timeout. What should you do?

A. Check the serial port logs of the Compute Engine instance.

B. Use Stackdriver Profiler to visualize the resources utilization throughout the application.

C. Determine whether there is an increased number of connections to the Cloud SQL instance.

D. Use Cloud Security Scanner to see whether your Cloud SQL is under a Distributed Denial of Service (DDoS) attack.

Correct Answer: B

QUESTION 10

You are designing a new Google Cloud organization for a client. Your client is concerned with the risks associated with long-lived credentials created in Google Cloud. You need to design a solution to completely eliminate the risks associated with the use of JSON service account keys while minimizing operational overhead. What should you do?

A. Apply the constraints/iam.disableServiceAccountKevCreation constraint to the organization.

B. Use custom versions of predefined roles to exclude all iam.serviceAccountKeys.* service account role permissions.

C. Apply the constraints/iam.disableServiceAccountKeyUpload constraint to the organization.

D. Grant the roles/iam.serviceAccountKeyAdmin IAM role to organization administrators only.

Correct Answer: A

constraints/iam.disableServiceAccountKeyCreation

QUESTION 11

Your company runs applications in Google Kubernetes Engine (GKE). Several applications rely on ephemeral volumes. You noticed some applications were unstable due to the DiskPressure node condition on the worker nodes. You need to identify which Pods are causing the issue, but you do not have execute access to workloads and nodes. What should you do?

A. Check the node/ephemeral_storage/used_bytes metric by using Metrics Explorer.

B. Check the container/ephemeral_storage/used_bytes metric by using Metrics Explorer.

C. Locate all the Pods with emptyDir volumes. Use the df -h command to measure volume disk usage.

D. Locate all the Pods with emptyDir volumes. Use the df -sh * command to measure volume disk usage.

Correct Answer: A

node/ephemeral_storage/used_bytes GA Ephemeral storage usage GAUGE, INT64, By k8s_node Local ephemeral storage bytes used by the node. Sampled every 60 seconds.
https://cloud.google.com/monitoring/api/metrics_kubernetes

QUESTION 12

You are configuring connectivity across Google Kubernetes Engine (GKE) clusters in different VPCs. You notice that the nodes in Cluster A are unable to access the nodes in Cluster B. You suspect that the workload access issue is due to the network configuration. You need to troubleshoot the issue but do not have execute access to workloads and nodes. You want to identify the layer at which the network connectivity is broken. What should you do?

A. Install a toolbox container on the node in Cluster Confirm that the routes to Cluster B are configured appropriately.

B. Use Network Connectivity Center to perform a Connectivity Test from Cluster A to Cluster B.

C. Use a debug container to run the traceroute command from Cluster A to Cluster B and from Cluster B to Cluster A. Identify the common failure point.

D. Enable VPC Flow Logs in both VPCs, and monitor packet drops.

Correct Answer: B

QUESTION 13

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.

Correct Answer: A

"The filter_record_transformer filter plugin mutates/transforms incoming event streams in a versatile manner. If there is a need to add/delete/modify events, this plugin is the first filter to try. It is included in the Fluentd\\'s core."

https://docs.fluentd.org/filter/record_transformer

QUESTION 14

You use Terraform to manage an application deployed to a Google Cloud environment. The application runs on instances deployed by a managed instance group. The Terraform code is deployed by using a CI/CD pipeline. When you change the machine type on the instance template used by the managed instance group, the pipeline fails at the terraform apply stage with the following error message:

```
Error waiting for Deleting Instance Template: The instance_template resource
'project/my-project/global/instanceTemplates/my-it-20220101010100000000001'
is already being used by 'projects/my-project/regions/us-
central1/instanceGroupManagers/my-mig'
```

You need to update the instance template and minimize disruption to the application and the number of pipeline runs. What should you do?

A. Delete the managed instance group, and recreate it after updating the instance template.

B. Add a new instance template, update the managed instance group to use the new instance template, and delete the old instance template.

C. Remove the managed instance group from the Terraform state file, update the instance template, and reimport the managed instance group.

D. Set the create_before_destroy meta-argument to true in the lifecycle block on the instance template.

Correct Answer: D

QUESTION 15

You are investigating issues in your production application that runs on Google Kubernetes Engine (GKE). You determined that the source of the issue is a recently updated container image, although the exact change in code was not identified. The deployment is currently pointing to the latest tag. You need to update your cluster to run a version of the container that functions as intended. What should you do?

A. Create a new tag called stable that points to the previously working container, and change the deployment to point to the new tag.

B. Alter the deployment to point to the sha256 digest of the previously working container.

C. Build a new container from a previous Git tag, and do a rolling update on the deployment to the new container.

D. Apply the latest tag to the previous container image, and do a rolling update on the deployment.

Correct Answer: B

https://cloud.google.com/kubernetes-engine/docs/concepts/about-container-images