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 Spinnaker to deploy your application and have created a canary deployment stage in the pipeline. Your application has an in-memory cache that loads objects at start time. You want to automate the comparison of the canary version against the production version. How should you configure the canary analysis?

A. Compare the canary with a new deployment of the current production version.
B. Compare the canary with a new deployment of the previous production version.
C. Compare the canary with the existing deployment of the current production version.
D. Compare the canary with the average performance of a sliding window of previous production versions.

Answer: A

Explanation:
https://cloud.google.com/architecture/automated-canary-analysis-kubernetes-engine-spinnaker https://spinnaker.io/guides/user/canary/best-practices/#compare-canary-against-baseline-not-against-production

NEW QUESTION 2
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 3
Your application runs on Google Cloud Platform (GCP). You need to implement Jenkins for deploying application releases to GCP. You want to streamline the release process, lower operational toil, and keep user data secure. What should you do?

A. Implement Jenkins on local workstations.
B. Implement Jenkins on Kubernetes on-premises
C. Implement Jenkins on Google Cloud Functions.
D. Implement Jenkins on Compute Engine virtual machines.

Answer: D

Explanation:
Your application runs on Google Cloud Platform (GCP). You need to implement Jenkins for deploying application releases to GCP. You want to streamline the release process, lower operational toil, and keep user data secure. What should you do?
https://plugins.jenkins.io/google-compute-engine/

NEW QUESTION 4
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.

Answer: B

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
Your organization recently adopted a container-based workflow for application development. Your team develops numerous applications that are deployed continuously through an automated build pipeline to a Kubernetes cluster in the production environment. The security auditor is concerned that developers or operators could circumvent automated testing and push code changes to production without approval. What should you do to enforce approvals?

A. Configure the build system with protected branches that require pull request approval.
B. Use an Admission Controller to verify that incoming requests originate from approved sources.
C. Leverage Kubernetes Role-Based Access Control (RBAC) to restrict access to only approved users.
D. Enable binary authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.

Answer: D

Explanation:
The keywords here is "developers or operators". Option A the operators could push images to production without approval (operators could touch the cluster directly and the cluster cannot do any action against them). Rest same as francisco_guerra.

NEW QUESTION 7
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 8
You are running an application in a virtual machine (VM) using a custom Debian image. The image has the Stackdriver Logging agent installed. The VM has the cloud-platform scope. The application is logging information via syslog. You want to use Stackdriver Logging in the Google Cloud Platform Console to visualize the logs. You notice that syslog is not showing up in the "All logs" dropdown list of the Logs Viewer. What is the first thing you should do?

A. Look for the agent's test log entry in the Logs Viewer.
B. Install the most recent version of the Stackdriver agent.
C. Verify the VM service account access scope includes the monitoring.write scope.
D. SSH to the VM and execute the following commands on your VM: ps ax I grep fluentd

Answer: D

Explanation:
https://cloud.google.com/compute/docs/access/service-accounts#associating_a_service_account_to_an_instance

NEW QUESTION 9
You support an application that stores product information in cached memory. For every cache miss, an entry is logged in Stackdriver Logging. You want to visualize how often a cache miss happens over time. What should you do?

A. Link Stackdriver Logging as a source in Google Data Studi
B. Filler (he logs on the cache misses.
C. Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
D. Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
E. Configure BigOuery as a sink for Stackdriver Loggin
F. Create a scheduled query to filter the cache miss logs and write them to a separate table

Answer: C

Explanation:
https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric

NEW QUESTION 10
Your team uses Cloud Build for all CI/CO pipelines. You want to use the kubectl builder for Cloud Build to deploy new images to Google Kubernetes Engine (GKE). You need to authenticate to GKE while minimizing development effort. What should you do?

A. Assign the Container Developer role to the Cloud Build service account.
B. Specify the Container Developer role for Cloud Build in the cloudbuild.yaml file.
C. Create a new service account with the Container Developer role and use it to run Cloud Build.
D. Create a separate step in Cloud Build to retrieve service account credentials and pass these to kubectl.

Answer: A

Explanation:
https://cloud.google.com/build/docs/deploying-builds/deploy-gke https://cloud.google.com/build/docs/securing-builds/configure-user-specified-service-accounts

NEW QUESTION 10
You support a Node.js application running on Google Kubernetes Engine (GKE) in production. The application makes several HTTP requests to dependent

applications. You want to anticipate which dependent applications might cause performance issues. What should you do?

A. Instrument all applications with Stackdriver Profiler.
B. Instrument all applications with Stackdriver Trace and review inter-service HTTP requests.
C. Use Stackdriver Debugger to review the execution of logic within each application to instrument all applications.
D. Modify the Node.js application to log HTTP request and response times to dependent application
E. Use Stackdriver Logging to find dependent applications that are performing poorly.

Answer: B

NEW QUESTION 13
You are deploying an application that needs to access sensitive information. You need to ensure that this information is encrypted and the risk of exposure is minimal if a breach occurs. What should you do?

A. Store the encryption keys in Cloud Key Management Service (KMS) and rotate the keys frequently
B. Inject the secret at the time of instance creation via an encrypted configuration management system.
C. Integrate the application with a Single sign-on (SSO) system and do not expose secrets to the application
D. Leverage a continuous build pipeline that produces multiple versions of the secret for each instance of the application.

Answer: A

Explanation:
https://cloud.google.com/security-key-management

NEW QUESTION 16
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 18
You support a large service with a well-defined Service Level Objective (SLO). The development team deploys new releases of the service multiple times a week.
If a major incident causes the service to miss its SLO, you want the development team to shift its focus from working on features to improving service reliability.
What should you do before a major incident occurs?

A. Develop an appropriate error budget policy in cooperation with all service stakeholders.
B. Negotiate with the product team to always prioritize service reliability over releasing new features.
C. Negotiate with the development team to reduce the release frequency to no more than once a week.
D. Add a plugin to your Jenkins pipeline that prevents new releases whenever your service is out of SLO.

Answer: A

Explanation:
Reason : Incident has not occurred yet, even when development team is already pushing new features multiple times a week. The option A says, to define an error budget "policy", not to define error budget(It is already present). Just simple means to bring in all stakeholders, and decide how to consume the error budget effectively that could bring balance between feature deployment and reliability.
The goals of this policy are to: -- Protect customers from repeated SLO misses -- Provide an incentive to balance reliability with other features
https://sre.google/workbook/error-budget-policy/

NEW QUESTION 23
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 26
Your organization wants to implement Site Reliability Engineering (SRE) culture and principles. Recently, a service that you support had a limited outage. A manager on another team asks you to provide a formal explanation of what happened so they can action remediations. What should you do?

A. Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action item
B. Share it with the manager only.
C. Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action item
D. Share it on the engineering organization's document portal.
E. Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each perso

F. Share it with the manager only.
G. Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each perso
H. Share it on the engineering organization's document portal.

Answer: B

NEW QUESTION 30
Your team of Infrastructure DevOps Engineers is growing, and you are starting to use Terraform to manage infrastructure. You need a way to implement code versioning and to share code with other team members. What should you do?

A. Store the Terraform code in a version-control syste
B. Establish procedures for pushing new versions and merging with the master.
C. Store the Terraform code in a network shared folder with child folders for each version releas
D. Ensure that everyone works on different files.
E. Store the Terraform code in a Cloud Storage bucket using object versionin
F. Give access to the bucket to every team member so they can download the files.
G. Store the Terraform code in a shared Google Drive folder so it syncs automatically to every team member's compute
H. Organize files with a naming convention that identifies each new version.

Answer: A

Explanation:
https://www.terraform.io/docs/cloud/guides/recommended-practices/part3.3.html

NEW QUESTION 34
Your development team has created a new version of their service's API. You need to deploy the new versions of the API with the least disruption to third-party developers and end users of third-party installed applications. What should you do?

A. Introduce the new version of the API.Announce deprecation of the old version of the AP
B. Deprecate the old version of the API.Contact remaining users of the old API.Provide best effort support to users of the old AP
C. Turn down the old version of the API.
D. Announce deprecation of the old version of the AP
E. Introduce the new version of the API.Contact remaining users on the old AP
F. Deprecate the old version of the AP
G. Turn down the old version of the API.Provide best effort support to users of the old API.
H. Announce deprecation of the old version of the AP
I. Contact remaining users on the old API.Introduce the new version of the AP
J. Deprecate the old version of the API.Provide best effort support to users of the old AP
K. Turn down the old version of the API.
L. Introduce the new version of the AP
M. Contact remaining users of the old API.Announce deprecation of the old version of the AP
N. Deprecate the old version of the API.Turn down the old version of the API.Provide best effort support to users of the old API.

Answer: A

NEW QUESTION 38
You are writing a postmortem for an incident that severely affected users. You want to prevent similar incidents in the future. Which two of the following sections should you include in the postmortem? (Choose two.)

A. An explanation of the root cause of the incident
B. A list of employees responsible for causing the incident
C. A list of action items to prevent a recurrence of the incident
D. Your opinion of the incident's severity compared to past incidents
E. Copies of the design documents for all the services impacted by the incident

Answer: AC

Explanation:
For a postmortem to be truly blameless, it must focus on identifying the contributing causes of the incident without indicting any individual or team for bad or inappropriate behavior.

NEW QUESTION 42
Your application artifacts are being built and deployed via a CI/CD pipeline. You want the CI/CD pipeline to securely access application secrets. You also want to more easily rotate secrets in case of a security breach. What should you do?

A. Prompt developers for secrets at build tim
B. Instruct developers to not store secrets at rest.
C. Store secrets in a separate configuration file on Gi
D. Provide select developers with access to the configuration file.
E. Store secrets in Cloud Storage encrypted with a key from Cloud KM
F. Provide the CI/CD pipeline with access to Cloud KMS via IAM.
G. Encrypt the secrets and store them in the source code repositor
H. Store a decryption key in a separate repository and grant your pipeline access to it

Answer: C

NEW QUESTION 46
......

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/