

Google.Premium.Professional Cloud DevOps Engineer.by.VCEplus.52q

Number: Cloud Certified
Passing Score: 800
Time Limit: 120 min
File Version: 2.2



Exam Code: Professional Cloud DevOps Engineer

Exam Name: Professional Cloud DevOps Engineer

Certification Provider: Google

Website: www.vceplus.com - www.vceplus.co

Free Exam: <https://vceplus.com/exam-professional-cloud-devops-engineer/>

Questions & Answers Exam Engine is rigorously checked before being put up for sale. We make sure there is nothing irrelevant in Professional Cloud DevOps Engineer exam products and you get latest questions. We strive to deliver the best Professional Cloud DevOps Engineer exam product for top grades in your first attempt.

Website: <https://vceplus.com> - www.vceplus.co

VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

Exam A

QUESTION 1

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 applications. Use Stackdriver Logging to find dependent applications that are performing poorly.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 2

You created a Stackdriver chart for CPU utilization in a dashboard within your workspace project. You want to share the chart with your Site Reliability Engineering (SRE) team only. You want to ensure you follow the principle of least privilege. What should you do?

- A. Share the workspace Project ID with the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- B. Share the workspace Project ID with the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
- C. Click “Share chart by URL” and provide the URL to the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- D. Click “Share chart by URL” and provide the URL to the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 3

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 4

You have a set of applications running on a Google Kubernetes Engine (GKE) cluster, and you are using Stackdriver Kubernetes Engine Monitoring. You are bringing a new containerized application required by your company into production. This application is written by a third party and cannot be modified or reconfigured. The application writes its log information to `/var/log/app_messages.log`, and you want to send these log entries to Stackdriver Logging. What should you do?

- A. Use the default Stackdriver Kubernetes Engine Monitoring agent configuration.
- B. Deploy a Fluentd daemonset to GKE. Then create a customized input and output configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- C. Install Kubernetes on Google Compute Engine (GCE) and redeploy your applications. Then customize the built-in Stackdriver Logging configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- D. Write a script to tail the log file within the pod and write entries to standard output. Run the script as a sidecar container with the application's pod. Configure a shared volume between the containers to allow the script to have read access to `/var/log` in the application container.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/customizing-stackdriver-logs-fluentd>

QUESTION 5

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 | grep fluentd`.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://groups.google.com/g/google-stackdriver-discussion/c/FXehB9a-5Vk?pli=1>

QUESTION 6

You use a multiple step Cloud Build pipeline to build and deploy your application to Google Kubernetes Engine (GKE). You want to integrate with a third-party monitoring platform by performing a HTTP POST of the build information to a webhook. You want to minimize the development effort. What should you do?

- A. Add logic to each Cloud Build step to HTTP POST the build information to a webhook.
- B. Add a new step at the end of the pipeline in Cloud Build to HTTP POST the build information to a webhook.
- C. Use Stackdriver Logging to create a logs-based metric from the Cloud Build logs. Create an Alert with a Webhook notification type.
- D. Create a Cloud Pub/Sub push subscription to the Cloud Build `cloud-builds` PubSub topic to HTTP POST the build information to a webhook.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 7

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/automated-canary-analysis-kubernetes-engine-spinnaker>

QUESTION 8

You support a high-traffic web application and want to ensure that the home page loads in a timely manner. As a first step, you decide to implement a Service Level Indicator (SLI) to represent home page request latency with an acceptable page load time set to 100 ms. What is the Google-recommended way of calculating this SLI?

- A. Bucketize the request latencies into ranges, and then compute the percentile at 100 ms.
- B. Bucketize the request latencies into ranges, and then compute the median and 90th percentiles.
- C. Count the number of home page requests that load in under 100 ms, and then divide by the total number of home page requests.
- D. Count the number of home page request that load in under 100 ms, and then divide by the total number of all web application requests.

Correct Answer: C

Section: (none)

Explanation



Explanation/Reference:

Reference: <https://sre.google/workbook/implementing-slos/>

QUESTION 9

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 two.)

- 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 server. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.

Correct Answer: AC

Section: (none)

Explanation**Explanation/Reference:****QUESTION 10**

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. • Deploy the Stackdriver logging agent to the application servers.
 - Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.
- B. • Deploy the Stackdriver logging agent to the application servers.
 - Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.
- C. • Deploy the Stackdriver monitoring agent to the application servers.
 - Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.
- D. • Install the gsutil command line tool on your application servers.
 - 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.
 - Give the developers the IAM Object Viewer access to view the logs in the specified bucket.

Correct Answer: B

Section: (none)

Explanation

**Explanation/Reference:****QUESTION 11**

You support the backend of a mobile phone game that runs on a Google Kubernetes Engine (GKE) cluster. The application is serving HTTP requests from users. You need to implement a solution that will reduce the network cost. What should you do?

- A. Configure the VPC as a Shared VPC Host project.
- B. Configure your network services on the Standard Tier.
- C. Configure your Kubernetes cluster as a Private Cluster.
- D. Configure a Google Cloud HTTP Load Balancer as Ingress.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/prep-kubernetes-engine-for-prod>

QUESTION 12

You encountered a major service outage that affected all users of the service for multiple hours. After several hours of incident management, the service returned to normal, and user access was restored. You need to provide an incident summary to relevant stakeholders following the Site Reliability Engineering recommended practices. What should you do first?

- A. Call individual stakeholders to explain what happened.
- B. Develop a post-mortem to be distributed to stakeholders.
- C. Send the Incident State Document to all the stakeholders.
- D. Require the engineer responsible to write an apology email to all stakeholders.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 13

You are performing a semi-annual capacity planning exercise for your flagship service. You expect a service user growth rate of 10% month-over-month over the next six months. Your service is fully containerized and runs on Google Cloud Platform (GCP), using a Google Kubernetes Engine (GKE) Standard regional cluster on three zones with cluster autoscaler enabled. You currently consume about 30% of your total deployed CPU capacity, and you require resilience against the failure of a zone. You want to ensure that your users experience minimal negative impact as a result of this growth or as a result of zone failure, while avoiding unnecessary costs. How should you prepare to handle the predicted growth?

- A. Verify the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verify your expected resource needs.
- B. Because you are deployed on GKE and are using a cluster autoscaler, your GKE cluster will scale automatically, regardless of growth rate.
- C. Because you are at only 30% utilization, you have significant headroom and you won't need to add any additional capacity for this rate of growth.
- D. Proactively add 60% more node capacity to account for six months of 10% growth rate, and then perform a load test to make sure you have enough capacity.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 14

Your application images are built and pushed to Google Container Registry (GCR). You want to build an automated pipeline that deploys the application when the image is updated while minimizing the development effort. What should you do?

- A. Use Cloud Build to trigger a Spinnaker pipeline.
- B. Use Cloud Pub/Sub to trigger a Spinnaker pipeline.
- C. Use a custom builder in Cloud Build to trigger Jenkins pipeline.
- D. Use Cloud Pub/Sub to trigger a custom deployment service running in Google Kubernetes Engine (GKE).

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 15

Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between the zones. You can fail over from one zone to another, but it causes a 10-minute service disruption for the affected users. You typically experience a database failure once per quarter and can detect it within five minutes. You are cataloging the reliability risks of a new real-time chat feature for your product. You catalog the following information for each risk:

- Mean Time to Detect (MTTD) in minutes
- Mean Time to Repair (MTTR) in minutes
- Mean Time Between Failure (MTBF) in days
- User Impact Percentage

The chat feature requires a new database system that takes twice as long to successfully fail over between zones. You want to account for the risk of the new database failing in one zone. What would be the values for the risk of database failover with the new system?

- A. MTTD: 5
MTTR: 10
MTBF: 90
Impact: 33%
- B. MTTD: 5
MTTR: 20
MTBF: 90
Impact: 33%
- C. MTTD: 5
MTTR: 10
MTBF: 90
Impact: 50%

D. MTTD: 5
MTTR: 20
MTBF: 90
Impact: 50%

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 16

You are managing the production deployment to a set of Google Kubernetes Engine (GKE) clusters. You want to make sure only images which are successfully built by your trusted CI/CD pipeline are deployed to production. What should you do?

- A. Enable Cloud Security Scanner on the clusters.
- B. Enable Vulnerability Analysis on the Container Registry.
- C. Set up the Kubernetes Engine clusters as private clusters.
- D. Set up the Kubernetes Engine clusters with Binary Authorization.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://codelabs.developers.google.com/codelabs/cloud-builder-gke-continuous-deploy/index.html#1>

QUESTION 17

You support an e-commerce application that runs on a large Google Kubernetes Engine (GKE) cluster deployed on-premises and on Google Cloud Platform. The application consists of microservices that run in containers. You want to identify containers that are using the most CPU and memory. What should you do?

- A. Use Stackdriver Kubernetes Engine Monitoring.
- B. Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.
- C. Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.
- D. Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

Correct Answer: C

Section: (none)

Explanation**Explanation/Reference:**

Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

QUESTION 18

Your company experiences bugs, outages, and slowness in its production systems. Developers use the production environment for new feature development and bug fixes. Configuration and experiments are done in the production environment, causing outages for users. Testers use the production environment for load testing, which often slows the production systems. You need to redesign the environment to reduce the number of bugs and outages in production and to enable testers to load test new features. What should you do?

- A. Create an automated testing script in production to detect failures as soon as they occur.
- B. Create a development environment with smaller server capacity and give access only to developers and testers.
- C. Secure the production environment to ensure that developers can't change it and set up one controlled update per year.
- D. Create a development environment for writing code and a test environment for configurations, experiments, and load testing.

Correct Answer: A

Section: (none)

Explanation**Explanation/Reference:****QUESTION 19**

You support an application running on App Engine. The application is used globally and accessed from various device types. You want to know the number of connections. You are using Stackdriver Monitoring for App Engine. What metric should you use?

- A. flex/connections/current
- B. tcp_ssl_proxy/new_connections
- C. tcp_ssl_proxy/open_connections
- D. flex/instance/connections/current

Correct Answer: D

Section: (none)

Explanation**Explanation/Reference:**

Reference: https://cloud.google.com/monitoring/api/metrics_gcp

QUESTION 20

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: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 21

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 22

You are on-call for an infrastructure service that has a large number of dependent systems. You receive an alert indicating that the service is failing to serve most of its requests and all of its dependent systems with hundreds of thousands of users are affected. As part of your Site Reliability Engineering (SRE) incident management protocol, you declare yourself Incident Commander (IC) and pull in two experienced people from your team as Operations Lead (OL) and Communications Lead (CL). What should you do next?

- A. Look for ways to mitigate user impact and deploy the mitigations to production.
- B. Contact the affected service owners and update them on the status of the incident.
- C. Establish a communication channel where incident responders and leads can communicate with each other.
- D. Start a postmortem, add incident information, circulate the draft internally, and ask internal stakeholders for input.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 23

You are developing a strategy for monitoring your Google Cloud Platform (GCP) projects in production using Stackdriver Workspaces. One of the requirements is to be able to quickly identify and react to production environment issues without false alerts from development and staging projects. You want to ensure that you adhere to the principle of least privilege when providing relevant team members with access to Stackdriver Workspaces. What should you do?

- A. Grant relevant team members read access to all GCP production projects. Create Stackdriver workspaces inside each project.
- B. Grant relevant team members the Project Viewer IAM role on all GCP production projects. Create Stackdriver workspaces inside each project.
- C. Choose an existing GCP production project to host the monitoring workspace. Attach the production projects to this workspace. Grant relevant team members read access to the Stackdriver Workspace.
- D. Create a new GCP monitoring project and create a Stackdriver Workspace inside it. Attach the production projects to this workspace. Grant relevant team members read access to the Stackdriver Workspace.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 24

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 from Stackdriver to BigQuery.
 2. From BigQuery, 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.
- D.
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.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

QUESTION 25

You need to run a business-critical workload on a fixed set of Compute Engine instances for several months. The workload is stable with the exact amount of resources allocated to it. You want to lower the costs for this workload without any performance implications. What should you do?

- A. Purchase Committed Use Discounts.
- B. Migrate the instances to a Managed Instance Group.
- C. Convert the instances to preemptible virtual machines.
- D. Create an Unmanaged Instance Group for the instances used to run the workload.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 26

You are part of an organization that follows SRE practices and principles. You are taking over the management of a new service from the Development Team, and you conduct a Production Readiness Review (PRR). After the PRR analysis phase, you determine that the service cannot currently meet its Service Level Objectives (SLOs). You want to ensure that the service can meet its SLOs in production. What should you do next?

- A. Adjust the SLO targets to be achievable by the service so you can bring it into production.
- B. Notify the development team that they will have to provide production support for the service.
- C. Identify recommended reliability improvements to the service to be completed before handover.
- D. Bring the service into production with no SLOs and build them when you have collected operational data.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 27

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/automated-canary-analysis-kubernetes-engine-spinnaker>

QUESTION 28

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 to a versioned Cloud Storage basket 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 day, 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.
- D. • 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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 29

You support a high-traffic web application with a microservice architecture. The home page of the application displays multiple widgets containing content such as the current weather, stock prices, and news headlines. The main serving thread makes a call to a dedicated microservice for each widget and then lays out the homepage for the user. The microservices occasionally fail; when that happens, the serving thread serves the homepage with some missing content. Users of the application are unhappy if this degraded mode occurs too frequently, but they would rather have some content served instead of no content at all. You want to set a Service Level Objective (SLO) to ensure that the user experience does not degrade too much. What Service Level Indicator (SLI) should you use to measure this?

- A. A quality SLI: the ratio of non-degraded responses to total responses.
- B. An availability SLI: the ratio of healthy microservices to the total number of microservices.
- C. A freshness SLI: the proportion of widgets that have been updated within the last 10 minutes.
- D. A latency SLI: the ratio of microservice calls that complete in under 100 ms to the total number of microservice calls.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/stackdriver/docs/solutions/slo-monitoring>

QUESTION 30

You support a multi-region web service running on Google Kubernetes Engine (GKE) behind a Global HTTP/S Cloud Load Balancer (CLB). For legacy reasons, user requests first go through a third-party Content Delivery Network (CDN), which then routes traffic to the CLB. You have already implemented an availability Service Level Indicator (SLI) at the CLB level. However, you want to increase coverage in case of a potential load balancer misconfiguration, CDN failure, or

other global networking catastrophe. Where should you measure this new SLI? (Choose two.)

- A. Your application servers' logs.
- B. Instrumentation coded directly in the client.
- C. Metrics exported from the application servers.
- D. GKE health checks for your application servers.
- E. A synthetic client that periodically sends simulated user requests.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 31

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:

Section: (none)

Explanation

Explanation/Reference:

QUESTION 32

You support a production service that runs on a single Compute Engine instance. You regularly need to spend time on recreating the service by deleting the crashing instance and creating a new instance based on the relevant image. You want to reduce the time spent performing manual operations while following Site Reliability Engineering principles. What should you do?

- A. File a bug with the development team so they can find the root cause of the crashing instance.
- B. Create a Managed instance Group with a single instance and use health checks to determine the system status.
- C. Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.
- D. Create a Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it was crashed.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 33

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 time. Instruct developers to not store secrets at rest.
- B. Store secrets in a separate configuration file on Git. Provide select developers with access to the configuration file.
- C. Store secrets in Cloud Storage encrypted with a key from Cloud KMS. Provide the CI/CD pipeline with access to Cloud KMS via IAM.
- D. Encrypt the secrets and store them in the source code repository. Store a decryption key in a separate repository and grant your pipeline access to it.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 34

Your company follows Site Reliability Engineering practices. You are the person in charge of Communications for a large, ongoing incident affecting your customer-facing applications. There is still no estimated time for a resolution of the outage. You are receiving emails from internal stakeholders who want updates on the outage, as well as emails from customers who want to know what is happening. You want to efficiently provide updates to everyone affected by the outage. What should you do?

- A. Focus on responding to internal stakeholders at least every 30 minutes. Commit to “next update” times.
- B. Provide periodic updates to all stakeholders in a timely manner. Commit to a “next update” time in all communications.
- C. Delegate the responding to internal stakeholder emails to another member of the Incident Response Team. Focus on providing responses directly to customers.

D. Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communications. Focus on providing responses directly to customers.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 35

Your team uses Cloud Build for all CI/CD 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 36

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 Studio. Filter the logs on the cache misses.
- B. Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
- C. Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
- D. Configure BigQuery as a sink for Stackdriver Logging. Create a scheduled query to filter the cache miss logs and write them to a separate table.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 37

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 38

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 log 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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 39

You have a CI/CD pipeline that uses Cloud Build to build new Docker images and push them to Docker Hub. You use Git for code versioning. After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline. You need to resolve the issue following Site Reliability Engineering practices. What should you do?

- A. Disable the CI pipeline and revert to manually building and pushing the artifacts.
- B. Change the CI pipeline to push the artifacts to Container Registry instead of Docker Hub.
- C. Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.
- D. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 40

Your company follows Site Reliability Engineering principles. You are writing a postmortem for an incident, triggered by a software change, that severely affected users. You want to prevent severe incidents from happening in the future. What should you do?

- A. Identify engineers responsible for the incident and escalate to their senior management.
- B. Ensure that test cases that catch errors of this type are run successfully before new software releases.
- C. Follow up with the employees who reviewed the changes and prescribe practices they should follow in the future.
- D. Design a policy that will require on-call teams to immediately call engineers and management to discuss a plan of action if an incident occurs.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 41

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 two.)

- 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.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 42

You manage an application that is writing logs to Stackdriver Logging. You need to give some team members the ability to export logs. What should you do?

- A. Grant the team members the IAM role of logging.configWriter on Cloud IAM.
- B. Configure Access Context Manager to allow only these members to export logs.
- C. Create and grant a custom IAM role with the permissions logging.sinks.list and logging.sink.get.
- D. Create an Organizational Policy in Cloud IAM to allow only these members to create log exports.

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

Reference: <https://cloud.google.com/logging/docs/access-control>

QUESTION 43

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/altostrat-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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 44

Your team has recently deployed an NGINX-based application into Google Kubernetes Engine (GKE) and has exposed it to the public via an HTTP Google Cloud Load Balancer (GCLB) ingress. You want to scale the deployment of the application's frontend using an appropriate Service Level Indicator (SLI). What should you do?

- A. Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.
- B. Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.
- C. Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.
- D. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 45

Your company follows Site Reliability Engineering practices. You are the Incident Commander for a new, customer-impacting incident. You need to immediately assign two incident management roles to assist you in an effective incident response. What roles should you assign? (Choose two.)

- A. Operations Lead
- B. Engineering Lead
- C. Communications Lead
- D. Customer Impact Assessor
- E. External Customer Communications Lead

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

QUESTION 46

You support an application running on GCP and want to configure SMS notifications to your team for the most critical alerts in Stackdriver Monitoring. You have already identified the alerting policies you want to configure this for. What should you do?

- A. Download and configure a third-party integration between Stackdriver Monitoring and an SMS gateway. Ensure that your team members add their SMS/phone numbers to the external tool.
- B. Select the Webhook notifications option for each alerting policy, and configure it to use a third-party integration tool. Ensure that your team members add their SMS/phone numbers to the external tool.
- C. Ensure that your team members set their SMS/phone numbers in their Stackdriver Profile. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list.
- D. Configure a Slack notification for each alerting policy. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are received. Ensure that your team members add their SMS/phone numbers to the external integration.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 47

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 Stacked 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 `metricKind` set to `METRIC_KIND_UNSPECIFIED` and a `valueType` set to `INT64`.
 - In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 48

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: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 49

You need to reduce the cost of virtual machines (VM) for your organization. After reviewing different options, you decide to leverage preemptible VM instances. Which application is suitable for preemptible VMs?

- A. A scalable in-memory caching system.
- B. The organization's public-facing website.
- C. A distributed, eventually consistent NoSQL database cluster with sufficient quorum.
- D. A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/preemptible-vms>

QUESTION 50

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 51

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 52

You are running a real-time gaming application on Compute Engine that has a production and testing environment. Each environment has their own Virtual Private Cloud (VPC) network. The application frontend and backend servers are located on different subnets in the environment's VPC. You suspect there is a malicious process communicating intermittently in your production frontend servers. You want to ensure that network traffic is captured for analysis. What should you do?

- A. Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 0.5.
- B. Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 1.0.
- C. Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 0.5. Apply changes in testing before production.
- D. Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 1.0. Apply changes in testing before production.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

