

Professional-Cloud-DevOps-Engineer Dumps

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

<https://www.certleader.com/Professional-Cloud-DevOps-Engineer-dumps.html>



NEW QUESTION 1

You have an application running in Google Kubernetes Engine. The application invokes multiple services per request but responds too slowly. You need to identify which downstream service or services are causing the delay. What should you do?

- A. Analyze VPC flow logs along the path of the request.
- B. Investigate the Liveness and Readiness probes for each service.
- C. Create a Dataflow pipeline to analyze service metrics in real time.
- D. Use a distributed tracing framework such as OpenTelemetry or Stackdriver Trace.

Answer: C

NEW QUESTION 2

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 3

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 4

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.

Answer: B

NEW QUESTION 5

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.

Answer: B

Explanation:

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

NEW QUESTION 6

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

Answer: A

Explanation:

https://cloud.google.com/monitoring/api/metrics_gcp#gcp-appengine

NEW QUESTION 7

You support a popular mobile game application deployed on Google Kubernetes Engine (GKE) across several Google Cloud regions. Each region has multiple Kubernetes clusters. You receive a report that none of the users in a specific region can connect to the application. You want to resolve the incident while following Site Reliability Engineering practices. What should you do first?

- A. Reroute the user traffic from the affected region to other regions that don't report issues.
- B. Use Stackdriver Monitoring to check for a spike in CPU or memory usage for the affected region.
- C. Add an extra node pool that consists of high memory and high CPU machine type instances to the cluster.
- D. Use Stackdriver Logging to filter on the clusters in the affected region, and inspect error messages in the logs.

Answer: A

Explanation:

Google always aims to first stop the impact of an incident, and then find the root cause (unless the root cause just happens to be identified early on).

NEW QUESTION 8

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 GK
- C. Then create a customized input and output configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- D. Install Kubernetes on Google Compute Engine (GCE) and redeploy your application
- E. Then customize the built-in Stackdriver Logging configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- F. Write a script to tail the log file within the pod and write entries to standard output
- G. Run the script as a sidecar container with the application's pod
- H. Configure a shared volume between the containers to allow the script to have read access to `/var/log` in the application container.

Answer: B

Explanation:

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

Besides the list of default logs that the Logging agent streams by default, you can customize the Logging agent to send additional logs to Logging or to adjust agent settings by adding input configurations. The configuration definitions in these sections apply to the fluent-plugin-google-cloud output plugin only and specify how logs are transformed and ingested into Cloud Logging. <https://cloud.google.com/logging/docs/agent/logging/configuration#configure>

NEW QUESTION 9

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 10

You are ready to deploy a new feature of a web-based application to production. You want to use Google Kubernetes Engine (GKE) to perform a phased rollout to half of the web server pods. What should you do?

- A. Use a partitioned rolling update.
- B. Use Node taints with NoExecute.
- C. Use a replica set in the deployment specification.
- D. Use a stateful set with parallel pod management policy.

Answer: A

Explanation:

<https://medium.com/velotio-perspectives/exploring-upgrade-strategies-for-stateful-sets-in-kubernetes-c02b8286f>

NEW QUESTION 10

You are using Stackdriver to monitor applications hosted on Google Cloud Platform (GCP). You recently deployed a new application, but its logs are not appearing on the Stackdriver dashboard.

You need to troubleshoot the issue. What should you do?

- A. Confirm that the Stackdriver agent has been installed in the hosting virtual machine.
- B. Confirm that your account has the proper permissions to use the Stackdriver dashboard.
- C. Confirm that port 25 has been opened in the firewall to allow messages through to Stackdriver.
- D. Confirm that the application is using the required client library and the service account key has proper permissions.

Answer: A

Explanation:

<https://cloud.google.com/monitoring/agent/monitoring/troubleshooting#checklist>

NEW QUESTION 12

You need to define Service Level Objectives (SLOs) for a high-traffic multi-region web application. Customers expect the application to always be available and have fast response times. Customers are currently happy with the application performance and availability. Based on current measurement, you observe that the 90th percentile of latency is 120ms and the 95th percentile of latency is 275ms over a 28-day window. What latency SLO would you recommend to the team to publish?

- A. 90th percentile – 100ms 95th percentile – 250ms
- B. 90th percentile – 120ms 95th percentile – 275ms
- C. 90th percentile – 150ms 95th percentile – 300ms
- D. 90th percentile – 250ms 95th percentile – 400ms

Answer: C

Explanation:

<https://sre.google/sre-book/service-level-objectives/>

NEW QUESTION 13

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

Answer: A

Explanation:

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

NEW QUESTION 16

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.

Answer: D

Explanation:

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

NEW QUESTION 20

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.

Answer: A

NEW QUESTION 23

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 26

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.

Answer: D

Explanation:

"After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline"- means something wrong on the recent change not with the image registry.

NEW QUESTION 31

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 34

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.

Answer: B

NEW QUESTION 36

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 41

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 43

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 too
- C. Ensure that your team members add their SMS/phone numbers to the external tool.
- D. Ensure that your team members set their SMS/phone numbers in their Stackdriver Profile
- E. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list.
- F. Configure a Slack notification for each alerting policy
- G. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are received
- H. Ensure that your team members add their SMS/phone numbers to the external integration.

Answer: C

Explanation:

https://cloud.google.com/monitoring/support/notification-options#creating_channels To configure SMS notifications, do the following:

In the SMS section, click Add new and follow the instructions. Click Save. When you set up your alerting policy, select the SMS notification type and choose a verified phone number from the list.

NEW QUESTION 46

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
- B. Commit to "next update" times.
- C. Provide periodic updates to all stakeholders in a timely manner
- D. Commit to a "next update" time in all communications.
- E. Delegate the responding to internal stakeholder emails to another member of the Incident Response Team
- F. Focus on providing responses directly to customers.
- G. Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communication
- H. Focus on providing responses directly to customers.

Answer: B

Explanation:

When disaster strikes, the person who declares the incident typically steps into the IC role and directs the high-level state of the incident. The IC concentrates on the 3Cs and does the following: Commands and coordinates the incident response, delegating roles as needed. By default, the IC assumes all roles that have not been delegated yet. Communicates effectively. Stays in control of the incident response. Works with other responders to resolve the incident. <https://sre.google/workbook/incident-response/>

NEW QUESTION 49

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 50

You are creating and assigning action items in a postmortem for an outage. The outage is over, but you need to address the root causes. You want to ensure that your team handles the action items quickly and efficiently. How should you assign owners and collaborators to action items?

- A. Assign one owner for each action item and any necessary collaborators.

- B. Assign multiple owners for each item to guarantee that the team addresses items quickly
- C. Assign collaborators but no individual owners to the items to keep the postmortem blameless.
- D. Assign the team lead as the owner for all action items because they are in charge of the SRE team.

Answer: A

Explanation:

<https://devops.com/when-it-disaster-strikes-part-3-conducting-a-blameless-post-mortem/>

NEW QUESTION 55

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

Answer: B

Explanation:

<https://www.atlassian.com/incident-management/kpis/common-metrics> <https://linkedin.github.io/school-of-sre/>

NEW QUESTION 60

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 64

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

Answer: B

Explanation:

<https://cloud.google.com/architecture/continuous-delivery-toolchain-spinnaker-cloud> <https://spinnaker.io/guides/user/pipeline/triggers/pubsub/>

NEW QUESTION 65

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 68

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 70

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.

Answer: D

Explanation:

<https://cloud.google.com/binary-authorization/docs/overview>

NEW QUESTION 71

You manage several production systems that run on Compute Engine in the same Google Cloud Platform (GCP) project. Each system has its own set of dedicated Compute Engine instances. You want to know how much it costs to run each of the systems. What should you do?

- A. In the Google Cloud Platform Console, use the Cost Breakdown section to visualize the costs per system.
- B. Assign all instances a label specific to the system they run
- C. Configure BigQuery billing export and query costs per label.
- D. Enrich all instances with metadata specific to the system they run
- E. Configure Stackdriver Logging to export to BigQuery, and query costs based on the metadata.
- F. Name each virtual machine (VM) after the system it runs
- G. Set up a usage report export to a Cloud Storage bucket
- H. Configure the bucket as a source in BigQuery to query costs based on VM name.

Answer: B

Explanation:

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

NEW QUESTION 74

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 2 answers

- 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

Answer: BE

NEW QUESTION 77

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.

Answer: A

Explanation:

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

NEW QUESTION 78

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 to METRIC_KIND_UNSPECIFIED and a 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 79

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 84

You support a user-facing web application. When analyzing the application's error budget over the previous six months, you notice that the application has never consumed more than 5% of its error budget in any given time window. You hold a Service Level Objective (SLO) review with business stakeholders and confirm that the SLO is set appropriately. You want your application's SLO to more closely reflect its observed reliability. What steps can you take to further that goal while balancing velocity, reliability, and business needs? (Choose two.)

- A. Add more serving capacity to all of your application's zones.
- B. Have more frequent or potentially risky application releases.
- C. Tighten the SLO to match the application's observed reliability.
- D. Implement and measure additional Service Level Indicators (SLIs) from the application.
- E. Announce planned downtime to consume more error budget, and ensure that users are not depending on a tighter SLO.

Answer: DE

Explanation:

<https://sre.google/sre-book/service-level-objectives/>
You want the application's SLO to more closely reflect its observed reliability. The key here is error budget never goes over 5%. This means they can have additional downtime and still stay within their budget.

NEW QUESTION 88

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Professional-Cloud-DevOps-Engineer Exam with Our Prep Materials Via below:

<https://www.certleader.com/Professional-Cloud-DevOps-Engineer-dumps.html>