

Google

Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam



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 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 3

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

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

Answer: B

Explanation:

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

NEW QUESTION 4

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.

Answer: C

NEW QUESTION 5

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 6

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 7

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

Answer: B

Explanation:

<https://cloud.google.com/blog/products/gcp/available-or-not-that-is-the-question-cre-life-lessons>

NEW QUESTION 8

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 9

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 log
- D. Create an Alert with a Webhook notification type.
- E. Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.

Answer: D

NEW QUESTION 10

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 10

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 15

Your company is developing applications that are deployed on Google Kubernetes Engine (GKE). Each team manages a different application. You need to create the development and production environments for each team, while minimizing costs. Different teams should not be able to access other teams' environments. What should you do?

- A. Create one GCP Project per tea
- B. In each project, create a cluster for Development and one for Productio
- C. Grant the teams IAM access to their respective clusters.
- D. Create one GCP Project per tea
- E. In each project, create a cluster with a Kubernetes namespace for Development and one for Productio
- F. Grant the teams IAM access to their respective clusters.
- G. Create a Development and a Production GKE cluster in separate project
- H. In each cluster, create a Kubernetes namespace per team, and then configure Identity Aware Proxy so that each team can only access its own namespace.
- I. Create a Development and a Production GKE cluster in separate project
- J. In each cluster, create a Kubernetes namespace per team, and then configure Kubernetes Role-based access control (RBAC) so that each team can only access its own namespace.

Answer: D

Explanation:

https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#roles_and_groups

NEW QUESTION 17

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 18

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 23

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

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

Answer: B

NEW QUESTION 26

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 30

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.

Answer: A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#custom_metrics <https://github.com/GoogleCloudPlatform/k8s-stackdriver/blob/master/custom-metrics-stackdriver-adapter/REA> Your application can report a custom metric to Cloud Monitoring. You can configure Kubernetes to respond to these metrics and scale your workload automatically. For example, you can scale your application based on metrics such as queries per second, writes per second, network performance, latency when communicating with a different application, or other metrics that make sense for your workload.
<https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics>

NEW QUESTION 34

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 36

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 40

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 2 answers

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

Answer: AC

Explanation:

<https://sre.google/workbook/incident-response/>

"The main roles in incident response are the Incident Commander (IC), Communications Lead (CL), and Operations or Ops Lead (OL)."

NEW QUESTION 43

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.

Answer: D

NEW QUESTION 45

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 50

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 minute
- B. Commit to "next update" times.
- C. Provide periodic updates to all stakeholders in a timely manne
- 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 Tea
- 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 53

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 (MUD) 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 56

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 58

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 60

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 62

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 person
- 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 person
- H. Share it on the engineering organization's document portal.

Answer: B

NEW QUESTION 64

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 69

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 72

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.

Answer: D

Explanation:

Costs associated with a load balancer are charged to the project containing the load balancer components. Because of these benefits, container-native load balancing is the recommended solution for load balancing through Ingress. When NEGs are used with GKE Ingress, the Ingress controller facilitates the creation of all aspects of the L7 load balancer. This includes creating the virtual IP address, forwarding rules, health checks, firewall rules, and more.

<https://cloud.google.com/architecture/best-practices-for-running-cost-effective-kubernetes-applications-on-gke>

NEW QUESTION 73

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 74

You have migrated an e-commerce application to Google Cloud Platform (GCP). You want to prepare the application for the upcoming busy season. What should you do first to prepare for the busy season?

- A. Load test the application to profile its performance for scaling.
- B. Enable AutoScaling on the production clusters, in case there is growth.
- C. Pre-provision double the compute power used last season, expecting growth.
- D. Create a runbook on inflating the disaster recovery (DR) environment if there is growth.

Answer: A

Explanation:

<https://cloud.google.com/blog/topics/retail/preparing-for-peak-holiday-season-while-wfh>

NEW QUESTION 77

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 78

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

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

Answer: C

Explanation:

https://cloud.google.com/compute/quotas#understanding_quotas <https://cloud.google.com/compute/quotas>

NEW QUESTION 83

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

Answer: C

NEW QUESTION 86

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 90

You are performing a semiannual 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
- C. your GKE cluster will scale automatically, regardless of growth rate.
- D. 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.
- E. 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.

Answer: A

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler>

The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload's CPU or memory consumption

NEW QUESTION 92

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Professional-Cloud-DevOps-Engineer Practice Exam Features:

- * Professional-Cloud-DevOps-Engineer Questions and Answers Updated Frequently
- * Professional-Cloud-DevOps-Engineer Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-DevOps-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Professional-Cloud-DevOps-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-DevOps-Engineer Practice Test Here](https://www.surepassexam.com/Professional-Cloud-DevOps-Engineer-exam-dumps.html)