Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

https://www.2passeasy.com/dumps/Professional-Cloud-DevOps-Engineer/

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 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. Buckelize Ihe 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 requests that load in under 100 m
E. and then divide by the total number of all web application requests.

Answer: C

Explanation:
https://sre.google/workbook/implementing-slos/
In the SRE principles book, it's recommended treating the SLI as the ratio of two numbers: the number of good events divided by the total number of events. For example: Number of successful HTTP requests / total HTTP requests (success rate)

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 lo 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 tog exclusion with userinfo as a filter.
D. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, create an advanced log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.

Answer: B

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

NEW QUESTION 6
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 Slackdriver 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 outpu
G. Run the script as a sidecar container with the application's po
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 7
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 8
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 9
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 support a trading application written in Python and hosted on App Engine flexible environment. You want to customize the error information being sent to Stackdriver Error Reporting. What should you do?

A. Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.
B. Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.
C. Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.

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

Answer: D

Explanation:
https://cloud.google.com/error-reporting/docs/formatting-error-messages https://cloud.google.com/error-reporting/docs/reference/libraries#client-libraries-install-python no need to install error reporting library on App Engine Flex.

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 15
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 has crashed.

Answer: B

NEW QUESTION 20
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 Buitd 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 24
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 29
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 (OLJ 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.

Answer: A

Explanation:
https://sre.google/sre-book/managing-incidents/

NEW QUESTION 31
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 36
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 Slackdriver 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 37
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 39
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 40
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 42
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 47
You support a stateless web-based API that is deployed on a single Compute Engine instance in the
europe-west2-a zone . The Service Level Indicator (SLI) for service availability is below the specified Service Level Objective (SLO). A postmortem has revealed that requests to the API regularly time out. The time outs are due to the API having a high number of requests and running out memory. You want to improve service availability. What should you do?

A. Change the specified SLO to match the measured SLI.
B. Move the service to higher-specification compute instances with more memory.
C. Set up additional service instances in other zones and load balance the traffic between all instances.
D. Set up additional service instances in other zones and use them as a failover in case the primary instance is unavailable.

Answer: C

NEW QUESTION 52
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 the production environment. A recent security audit alerted your team that the code pushed to production could contain vulnerabilities and that the existing tooling around virtual machine (VM) vulnerabilities no longer applies to the containerized environment. You need to ensure the security and patch level of all code running through the pipeline. What should you do?

A. Set up Container Analysis to scan and report Common Vulnerabilities and Exposures.
B. Configure the containers in the build pipeline to always update themselves before release.
C. Reconfigure the existing operating system vulnerability software to exist inside the container.
D. Implement static code analysis tooling against the Docker files used to create the containers.

Answer: D

Explanation:
https://cloud.google.com/binary-authorization
Binary Authorization is a deploy-time security control that ensures only trusted container images are deployed on Google Kubernetes Engine (GKE) or Cloud Run. With Binary Authorization, you can require images to be signed by trusted authorities during the development process and then enforce signature validation when deploying. By enforcing validation, you can gain tighter control over your container environment by ensuring only verified images are integrated into the build-and-release process.

NEW QUESTION 56
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 59
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 tea

B. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
C. Share the workspace Project ID with the SRE tea
D. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
E. Click "Share chart by URL" and provide the URL to the SRE tea
F. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
G. Click "Share chart by URL" and provide the URL to the SRE tea
H. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

Answer: C

Explanation:
https://cloud.google.com/monitoring/access-control

NEW QUESTION 62
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.

Answer: C

Explanation:
https://cloud.google.com/kubernetes-engine/docs/tutorials/autoscaling-metrics

NEW QUESTION 65
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 duster 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 70
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 2 answers

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

Answer: BE

NEW QUESTION 71
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 74
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 intesting 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.

Answer: D

NEW QUESTION 75
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 76
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 BigOuer
E. aggregate logs per container, and then analyze CPU and memory consumption.

Answer: A

Explanation:
https://cloud.google.com/anthos/clusters/docs/on-prem/1.7/concepts/logging-and-monitoring

NEW QUESTION 80
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Professional-Cloud-DevOps-Engineer Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Professional-Cloud-DevOps-Engineer Product From:

https://www.2passeasy.com/dumps/Professional-Cloud-DevOps-Engineer/

Money Back Guarantee

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