Professional-Cloud-DevOps-Engineer Braindumps Real Exam Updated on Jan 20, 2024 with 166 Questions [Q47-Q66

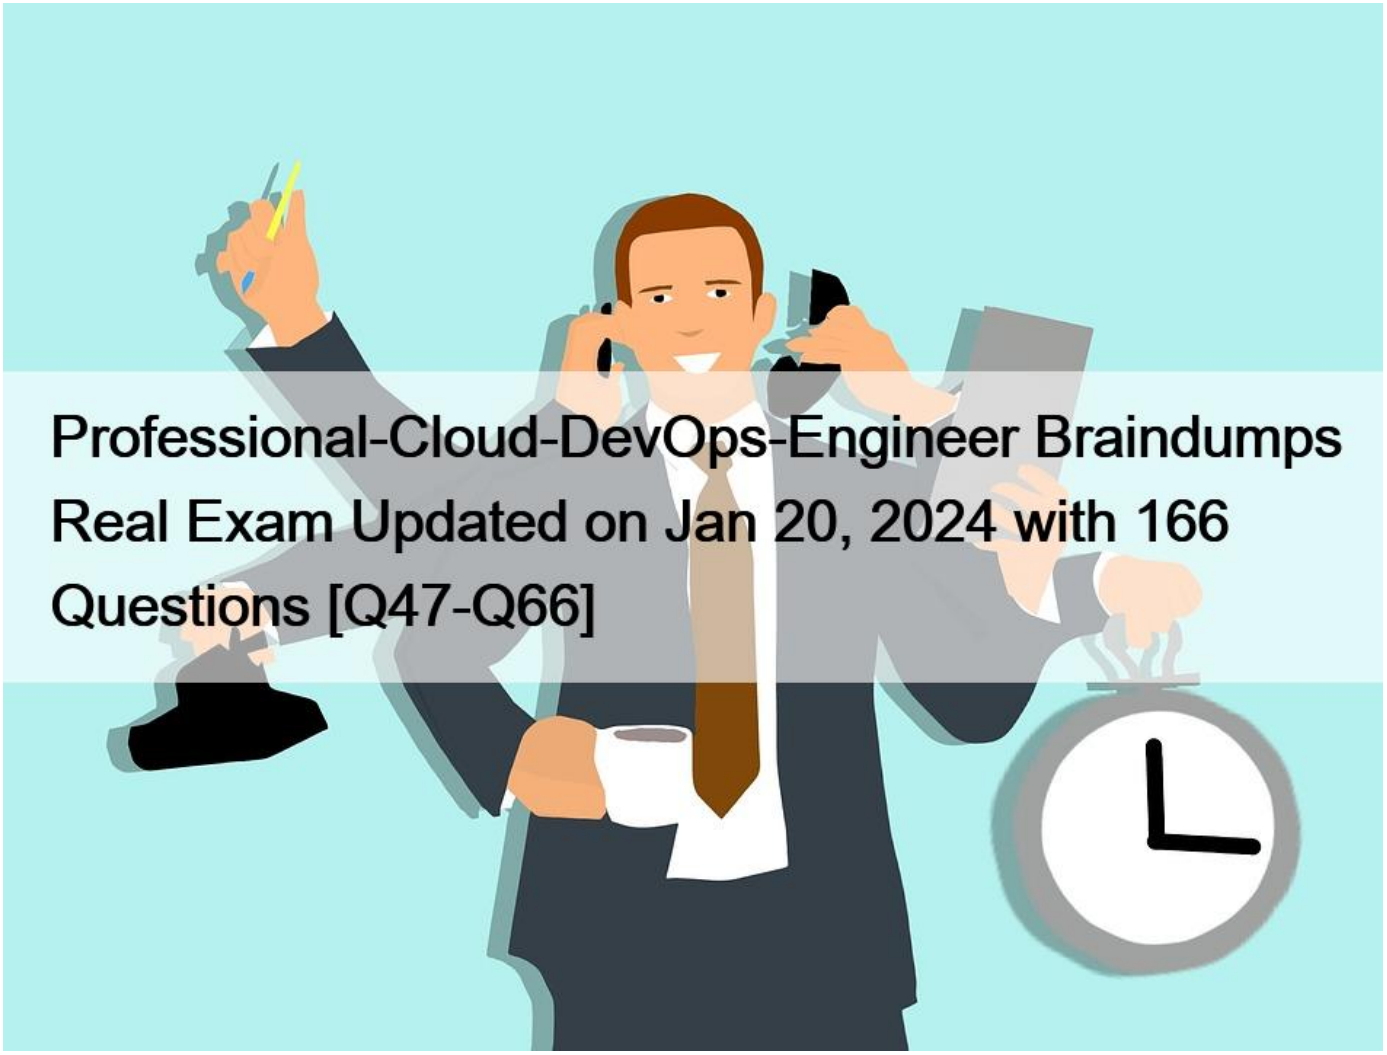

Professional-Cloud-DevOps-Engineer Braindumps Real Exam Updated on Jan 20, 2024 with 166 Questions
Latest Professional-Cloud-DevOps-Engineer PDF Dumps & Real Tests Free Updated Today

To prepare for the exam, candidates are advised to take relevant training courses, read the official study guide, and practice using the Google Cloud Platform. They should also have hands-on experience working with DevOps tools and technologies, such as Docker, Kubernetes, Jenkins, and Terraform. With the right preparation, candidates can pass the Google Professional-Cloud-DevOps-Engineer exam and join the elite group of Cloud DevOps experts who are in high demand in the IT industry.

NEW QUESTION 47

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?
* Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.
* Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.
* Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.
* Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.
https://cloud.google.com/kubernetes-engine/docs/tutorials/autoscaling-metrics The Google Cloud HTTP Load Balancer (GCLB) provides metrics on the number of requests and the response latency for each backend service. These metrics can be used as custom metrics for the horizontal pod autoscaler (HPA) to scale the deployment based on the load. This is the correct solution to use an appropriate SLI for scaling.

NEW QUESTION 48

You are managing an application that runs in Compute Engine The application uses a custom HTTP server to expose an API that is accessed by other applications through an internal TCP/UDP load balancer A firewall rule allows access to the API port from 0.0.0-0/0. You need to configure Cloud Logging to log each IP address that accesses the API by using the fewest number of steps What should you do Bret?
* Enable Packet Mirroring on the VPC
* Install the Ops Agent on the Compute Engine instances.
* Enable logging on the firewall rule
* Enable VPC Flow Logs on the subnet
Explanation

The best option for configuring Cloud Logging to log each IP address that accesses the API by using the fewest number of steps is to enable logging on the firewall rule. A firewall rule is a rule that controls the traffic to and from your Compute Engine instances. You can enable logging on a firewall rule to capture information about the traffic that matches the rule, such as source and destination IP addresses, protocols, ports, and actions. You can use Cloud Logging to view and export the firewall logs to other destinations, such as BigQuery, for further analysis.

NEW QUESTION 49

You are reviewing your deployment pipeline in Google Cloud Deploy You must reduce toil in the pipeline and you want to minimize the amount of time it takes to complete an end-to-end deployment What should you do?

Choose 2 answers
* Create a trigger to notify the required team to complete the next step when manual intervention is required
* Divide the automation steps into smaller tasks
* Use a script to automate the creation of the deployment pipeline in Google Cloud Deploy
* Add more engineers to finish the manual steps.
* Automate promotion approvals from the development environment to the test environment
Explanation

The best options for reducing toil in the pipeline and minimizing the amount of time it takes to complete an end-to-end deployment are to create a trigger to notify the required team to complete the next step when manual intervention is required and to automate promotion approvals from the development environment to the test environment. A trigger is a resource that initiates a deployment when an event occurs, such as a code change, a schedule, or a manual request. You can create a trigger to notify the required team to complete the next step when manual intervention is required by using Cloud Build or Cloud Functions. This way, you can reduce the waiting time and human errors in the pipeline. A promotion approval is a process that allows you to approve or reject a deployment

from one environment to another, such as from development to test. You can automate promotion approvals from the development environment to the test environment by using Google Cloud Deploy or Cloud Build. This way, you can speed up the deployment process and avoid manual steps.

NEW QUESTION 50

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?
* Use a partitioned rolling update.
* Use Node taints with NoExecute.
* Use a replica set in the deployment specification.
* Use a stateful set with parallel pod management policy.
https://medium.com/velotio-perspectives/exploring-upgrade-strategies-for-stateful-sets-in-kubernetes-c02b8286f251

NEW QUESTION 51

You support an application that stores product information in cached memory. For every cache miss, an entry is logged in Stackdriver Logging. You want to visualize how often a cache miss happens over time. What should you do?
* Link Stackdriver Logging as a source in Google Data Studio. Filler (he logs on the cache misses.
* Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
* Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
* Configure BigOuery as a sink for Stackdriver Logging. Create a scheduled query to filter the cache miss logs and write them to a separate table
https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric

NEW QUESTION 52

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?
* Share the workspace Project ID with the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
* Share the workspace Project ID with the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
* Click “Share chart by URL” and provide the URL to the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
* Click “Share chart by URL” and provide the URL to the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

NEW QUESTION 53

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
* Operations Lead

* Engineering Lead
* Communications Lead
* Customer Impact Assessor
* External Customer Communications Lead

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).” The Operations Lead is responsible for managing the operational aspects of the incident, such as deploying fixes, rolling back changes, or restoring backups. The External Customer Communications Lead is not a standard role in incident response, but it could be delegated by the Communications Lead if needed.

NEW QUESTION 54

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

* Your application servers’ logs
* Instrumentation coded directly in the client
* Metrics exported from the application servers
* GKE health checks for your application servers
* A synthetic client that periodically sends simulated user requests

NEW QUESTION 55

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?

* Implement Jenkins on local workstations.
* Implement Jenkins on Kubernetes on-premises
* Implement Jenkins on Google Cloud Functions.
* Implement Jenkins on Compute Engine virtual machines.

NEW QUESTION 56

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?

* Configure the VPC as a Shared VPC Host project.
* Configure your network services on the Standard Tier.
* Configure your Kubernetes duster as a Private Cluster.
* Configure a Google Cloud HTTP Load Balancer as Ingress.

The Standard Tier network service offers lower network costs than the Premium Tier. This is the correct option to reduce the network cost for the application3.

NEW QUESTION 57

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

NEW QUESTION 58

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

NEW QUESTION 59

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?
* Reroute the user traffic from the affected region to other regions that don’t report issues.
* Use Stackdriver Monitoring to check for a spike in CPU or memory usage for the affected region.
* Add an extra node pool that consists of high memory and high CPU machine type instances to the cluster.
* Use Stackdriver Logging to filter on the clusters in the affected region, and inspect error messages in the logs.
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 60

Your organization wants to increase the availability target of an application from 99 9% to 99 99% for an investment of $2 000 The application’s current revenue is S1,000,000 You need to determine whether the increase in availability is worth the investment for a single year of usage What should you do?
* Calculate the value of improved availability to be $900, and determine that the increase in availability is not worth the investment
* Calculate the value of improved availability to be $1 000 and determine that the increase in availability is not worth the investment
* Calculate the value of improved availability to be $1 000 and determine that the increase in availability is worth the investment
* Calculate the value of improved availability to be $9,000. and determine that the increase in availability is worth the investment
The best option for determining whether the increase in availability is worth the investment for a single year of usage is to calculate the value of improved availability to be $900, and determine that the increase in availability is not worth the investment. To calculate the value of improved availability, we can use the following formula:

Value of improved availability = Revenue * (New availability – Current availability) Plugging in the given numbers, we get:

Value of improved availability = $1,000,000 * (0.9999 – 0.999) = $900

Since the value of improved availability is less than the investment of $2,000, we can conclude that the increase in availability is not worth the investment.

NEW QUESTION 61

You are configuring connectivity across Google Kubernetes Engine (GKE) clusters in different VPCs You notice that the nodes in Cluster A are unable to access the nodes in Cluster B You suspect that the workload access issue is due to the network configuration You need to troubleshoot the issue but do not have execute access to workloads and nodes You want to identify the layer at which the network connectivity is broken What should you do?
* Use Network Connectivity Center to perform a Connectivity Test from Cluster A to Cluster
* Use a debug container to run the traceroute command from Cluster A to Cluster B and from Cluster B to Cluster A Identify the common failure point
* Install a toolbox container on the node in Cluster A Confirm that the routes to Cluster B are configured appropriately
* Enable VPC Flow Logs in both VPCs and monitor packet drops
Explanation

The best option for troubleshooting the issue without having execute access to workloads and nodes is to use Network Connectivity Center to perform a Connectivity Test from Cluster A to Cluster B. Network Connectivity Center is a service that allows you to create, manage, and monitor network connectivity across Google Cloud, hybrid, and multi-cloud environments. You can use Network Connectivity Center to perform a Connectivity Test, which is a feature that allows you to test the reachability and latency between two endpoints, such as GKE clusters, VM instances, or IP addresses. By using Network Connectivity Center to perform a Connectivity Test from Cluster A to Cluster B, you can identify the layer at which the network connectivity is broken, such as the firewall, routing, or load balancing.

NEW QUESTION 62

You are configuring a Cl pipeline. The build step for your Cl pipeline integration testing requires access to APIs inside your private VPC network. Your security team requires that you do not expose API traffic publicly. You need to implement a solution that minimizes management overhead. What should you do?
* Use Cloud Build private pools to connect to the private VPC.
* Use Spinnaker for Google Cloud to connect to the private VPC.
* Use Cloud Build as a pipeline runner. Configure Internal HTTP(S) Load Balancing for API access.
* Use Cloud Build as a pipeline runner. Configure External HTTP(S) Load Balancing with a Google Cloud Armor policy for API access.
Cloud Build is a service that executes your builds on Google Cloud Platform infrastructure1. Cloud Build can be used as a pipeline runner for your CI pipeline, which is a process that automates the integration and testing of your code2. Cloud Build private pools are private, dedicated pools of workers that offer greater customization over the build environment, including the ability to access resources in a private VPC network3. A VPC network is a virtual network that provides connectivity for your Google Cloud resources and services. By using Cloud Build private pools, you can implement a solution that minimizes management overhead, as Cloud Build private pools are hosted and fully-managed by Cloud Build and scale up and down to zero, with no infrastructure to set up, upgrade, or scale3. You can also implement a solution that meets your security requirement, as Cloud Build private pools use network peering to connect into your private VPC network and do not expose API traffic publicly.

NEW QUESTION 63

Your company operates in a highly regulated domain. Your security team requires that only trusted container images can be deployed to Google Kubernetes Engine (GKE). You need to implement a solution that meets the requirements of the security team, while minimizing management overhead. What should you do?

* Grant the roles/artifactregistry. writer role to the Cloud Build service account. Confirm that no employee has Artifact Registry write permission.
* Use Cloud Run to write and deploy a custom validator Enable an Eventarc trigger to perform validations when new images are uploaded.
* Configure Kritis to run in your GKE clusters to enforce deploy-time security policies.
* Configure Binary Authorization in your GKE clusters to enforce deploy-time security policies

NEW QUESTION 64

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?
* Focus on responding to internal stakeholders at least every 30 minutes. Commit to “next update” times.
* Provide periodic updates to all stakeholders in a timely manner. Commit to a “next update” time in all communications.
* Delegate the responding to internal stakeholder emails to another member of the Incident Response Team.

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

NEW QUESTION 65

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?
* Configure the build system with protected branches that require pull request approval.
* Use an Admission Controller to verify that incoming requests originate from approved sources.
* Leverage Kubernetes Role-Based Access Control (RBAC) to restrict access to only approved users.
* Enable binary authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.
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 66

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?
* 90th percentile – 100ms

95th percentile – 250ms
* 90th percentile – 120ms

95th percentile – 275ms
* 90th percentile – 150ms

95th percentile – 300ms
* 90th percentile – 250ms

95th percentile – 400ms
https://sre.google/sre-book/service-level-objectives/

Google Professional-Cloud-DevOps-Engineer is a certification exam offered by Google Cloud that tests the skills and knowledge of professionals in the field of cloud-based DevOps. Google Cloud Certified - Professional Cloud DevOps Engineer Exam certification demonstrates a candidate's ability to design, implement, and manage DevOps practices on the Google Cloud platform. It requires a solid understanding of software development, deployment, and cloud infrastructure management.

Professional-Cloud-DevOps-Engineer Dumps With 100% Verified Q&As - Pass Guarantee or Full Refund: https://www.vceprep.com/Professional-Cloud-DevOps-Engineer-latest-vce-prep.html]