

Google

Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam



NEW QUESTION 1

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

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

Answer: B

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

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

Answer: B

NEW QUESTION 4

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

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

Answer: A

Explanation:

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

NEW QUESTION 5

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 6

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 7

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 8

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 9

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

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

Answer: C

Explanation:

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

NEW QUESTION 10

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 10

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 12

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 17

You are running an application in a virtual machine (VM) using a custom Debian image. The image has the Stackdriver Logging agent installed. The VM has the cloud-platform scope. The application is logging information via syslog. You want to use Stackdriver Logging in the Google Cloud Platform Console to visualize the logs. You notice that syslog is not showing up in the "All logs" dropdown list of the Logs Viewer. What is the first thing you should do?

- A. Look for the agent's test log entry in the Logs Viewer.
- B. Install the most recent version of the Stackdriver agent.
- C. Verify the VM service account access scope includes the monitoring.write scope.
- D. SSH to the VM and execute the following commands on your VM: ps ax | grep fluentd

Answer: D

Explanation:

https://cloud.google.com/compute/docs/access/service-accounts#associating_a_service_account_to_an_instance

NEW QUESTION 22

You are on-call for an infrastructure service that has a large number of dependent systems. You receive an alert indicating that the service is failing to serve most of its requests and all of its dependent systems with hundreds of thousands of users are affected. As part of your Site Reliability Engineering (SRE) incident management protocol, you declare yourself Incident Commander (IC) and pull in two experienced people from your team as Operations Lead (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 24

You support an application that stores product information in cached memory. For every cache miss, an entry is logged in Stackdriver Logging. You want to visualize how often a cache miss happens over time. What should you do?

- A. Link Stackdriver Logging as a source in Google Data Studi
- B. Filter (he logs on the cache misses.
- C. Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
- D. Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
- E. Configure BigQuery as a sink for Stackdriver Loggin
- F. Create a scheduled query to filter the cache miss logs and write them to a separate table

Answer: C

Explanation:

<https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric>

NEW QUESTION 28

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 30

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 34

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 35

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 36

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

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

Answer: A

Explanation:

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

NEW QUESTION 41

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 44

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 49

You support a web application that is hosted on Compute Engine. The application provides a booking service for thousands of users. Shortly after the release of a new feature, your monitoring dashboard shows that all users are experiencing latency at login. You want to mitigate the impact of the incident on the users of your service. What should you do first?

- A. Roll back the recent release.
- B. Review the Stackdriver monitoring.
- C. Upsize the virtual machines running the login services.
- D. Deploy a new release to see whether it fixes the problem.

Answer: C

Explanation:

Rollback to previous stable version. Then you need to find what is causing the issue.

NEW QUESTION 53

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 58

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 59

Your development team has created a new version of their service's API. You need to deploy the new versions of the API with the least disruption to third-party developers and end users of third-party installed applications. What should you do?

- A. Introduce the new version of the API. Announce deprecation of the old version of the AP
- B. Deprecate the old version of the API. Contact remaining users of the old API. Provide best effort support to users of the old AP
- C. Turn down the old version of the API.
- D. Announce deprecation of the old version of the AP
- E. Introduce the new version of the API. Contact remaining users on the old AP
- F. Deprecate the old version of the AP
- G. Turn down the old version of the API. Provide best effort support to users of the old API.
- H. Announce deprecation of the old version of the AP
- I. Contact remaining users on the old API. Introduce the new version of the AP
- J. Deprecate the old version of the API. Provide best effort support to users of the old AP
- K. Turn down the old version of the API.
- L. Introduce the new version of the AP
- M. Contact remaining users of the old API. Announce deprecation of the old version of the AP
- N. Deprecate the old version of the API. Turn down the old version of the API. Provide best effort support to users of the old API.

Answer: A

NEW QUESTION 62

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 65

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 66

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 71

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 74

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

- A. Grant relevant team members read access to all GCP production project
- B. Create Stackdriver workspaces inside each project.
- C. Grant relevant team members the Project Viewer IAM role on all GCP production project
- D. Create Slackdriver workspaces inside each project.
- E. Choose an existing GCP production project to host the monitoring workspac
- F. Attach the production projects to this workspac
- G. Grant relevant team members read access to the Stackdriver Workspace.
- H. Create a new GCP monitoring project, and create a Stackdriver Workspace inside i
- I. Attach the production projects to this workspac
- J. Grant relevant team members read access to the Stackdriver Workspace.

Answer: D

Explanation:

"A Project can host many Projects and appear in many Projects, but it can only be used as the scoping project once. We recommend that you create a new Project for the purpose of having multiple Projects in the same scope."

NEW QUESTION 76

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 match the application's observed reliability.
- D. Implement and measure additional Service Level Indicators (SLIs) fro 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 it's 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 78

.....

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)