



# Google

## Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

## About Exambible

### *Your Partner of IT Exam*

## Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

#### 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

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 3

You encounter a large number of outages in the production systems you support. You receive alerts for all the outages that wake you up at night. The alerts are due to unhealthy systems that are automatically restarted within a minute. You want to set up a process that would prevent staff burnout while following Site Reliability Engineering practices. What should you do?

- A. Eliminate unactionable alerts.
- B. Create an incident report for each of the alerts.
- C. Distribute the alerts to engineers in different time zones.
- D. Redefine the related Service Level Objective so that the error budget is not exhausted.

**Answer:** A

#### Explanation:

Eliminate bad monitoring : Unactionable alerts (i.e., spam) <https://cloud.google.com/blog/products/management-tools/meeting-reliability-challenges-with-sre-principles>

agree with kyubiblaze about having to remove unactionable items aka spam: "good monitoring alerts on actionable problems" @ <https://cloud.google.com/blog/products/management-tools/meeting-reliability-challenges-with-sre-principles>

#### NEW QUESTION 4

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

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

**Answer:** B

#### NEW QUESTION 5

You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. All PII entries begin with the text userinfo. You want to capture these log entries in a secure location for later review and prevent them from leaking to Stackdriver Logging. What should you do?

- A. Create a basic log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.
- B. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, and then copy the entries to a Cloud Storage bucket.
- C. Create an advanced log filter matching userinfo, configure a log export in the Stackdriver console with Cloud Storage as a sink, and then configure a log exclusion with userinfo as a filter.
- D. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, create an advanced log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 6

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

- A. flex/connections/current
- B. tcp\_ssl\_proxy/new\_connections
- C. tcp\_ssl\_proxy/open\_connections

D. flex/instance/connections/current

**Answer:** A

**Explanation:**

[https://cloud.google.com/monitoring/api/metrics\\_gcp#gcp-appengine](https://cloud.google.com/monitoring/api/metrics_gcp#gcp-appengine)

#### NEW QUESTION 7

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

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

**Answer:** A

**Explanation:**

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

#### NEW QUESTION 8

You 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 9

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 have a pool of application servers running on Compute Engine. You need to provide a secure solution that requires the least amount of configuration and allows developers to easily access application logs for troubleshooting. How would you implement the solution on GCP?

- A. • Deploy the Stackdriver logging agent to the application servers. • Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.
- B. • Deploy the Stackdriver logging agent to the application servers. • Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.
- C. • Deploy the Stackdriver monitoring agent to the application servers. • Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.
- D. • Install the gsutil command line tool on your application servers. • Write a script using gsutil to upload your application log to a Cloud Storage bucket, and then schedule it to run via cron every 5 minutes. • Give the developers IAM Object Viewer access to view the logs in the specified bucket.

**Answer:** A

**Explanation:**

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

#### NEW QUESTION 10

You are working with a government agency that requires you to archive application logs for seven years. You need to configure Stackdriver to export and store the logs while minimizing costs of storage. What should you do?

- A. Create a Cloud Storage bucket and develop your application to send logs directly to the bucket.
- B. Develop an App Engine application that pulls the logs from Stackdriver and saves them in BigQuery.
- C. Create an export in Stackdriver and configure Cloud Pub/Sub to store logs in permanent storage for seven years.
- D. Create a sink in Stackdriver, name it, create a bucket on Cloud Storage for storing archived logs, and then select the bucket as the log export destination.

**Answer:** D

**Explanation:**

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

**NEW QUESTION 13**

You need to run a business-critical workload on a fixed set of Compute Engine instances for several months. The workload is stable with the exact amount of resources allocated to it. You want to lower the costs for this workload without any performance implications. What should you do?

- A. Purchase Committed Use Discounts.
- B. Migrate the instances to a Managed Instance Group.
- C. Convert the instances to preemptible virtual machines.
- D. Create an Unmanaged Instance Group for the instances used to run the workload.

**Answer:** A

**NEW QUESTION 14**

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 16**

You have a CI/CD pipeline that uses Cloud Build to build new Docker images and push them to Docker Hub. You use Git for code versioning. After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline. You need to resolve the issue following Site Reliability Engineering practices. What should you do?

- A. Disable the CI pipeline and revert to manually building and pushing the artifacts.
- B. Change the CI pipeline to push the artifacts to Container Registry instead of Docker Hub.
- C. Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.
- D. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

**Answer:** D

**Explanation:**

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

**NEW QUESTION 19**

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. You want to prevent these fields from being written in new log entries as quickly as possible. What should you do?

- A. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
- B. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.
- C. Wait for the application developers to patch the application, and then verify that the log entries are no longer exposing PII.
- D. Stage log entries to Cloud Storage, and then trigger a Cloud Function to remove the fields and write the entries to Stackdriver via the Stackdriver Logging API.

**Answer:** A

**NEW QUESTION 22**

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 to 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 day, confirm that all changes have been captured in the files within the folder structure.
- D. • 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 23**



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 26

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 29

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 32

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 34

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 35

You use Cloud Build to build and deploy your application. You want to securely incorporate database credentials and other application secrets into the build pipeline. You also want to minimize the development effort. What should you do?

- A. Create a Cloud Storage bucket and use the built-in encryption at rest
- B. Store the secrets in the bucket and grant Cloud Build access to the bucket.

- C. Encrypt the secrets and store them in the application repository
- D. Store a decryption key in a separate repository and grant Cloud Build access to the repository.
- E. Use client-side encryption to encrypt the secrets and store them in a Cloud Storage bucket
- F. Store a decryption key in the bucket and grant Cloud Build access to the bucket.
- G. Use Cloud Key Management Service (Cloud KMS) to encrypt the secrets and include them in your Cloud Build deployment configuration
- H. Grant Cloud Build access to the KeyRing.

**Answer:** D

**Explanation:**

<https://cloud.google.com/build/docs/securing-builds/use-encrypted-credentials>

**NEW QUESTION 39**

Your team is designing a new application for deployment both inside and outside Google Cloud Platform (GCP). You need to collect detailed metrics such as system resource utilization. You want to use centralized GCP services while minimizing the amount of work required to set up this collection system. What should you do?

- A. Import the Stackdriver Profiler package, and configure it to relay function timing data to Stackdriver for further analysis.
- B. Import the Stackdriver Debugger package, and configure the application to emit debug messages with timing information.
- C. Instrument the code using a timing library, and publish the metrics via a health check endpoint that is scraped by Stackdriver.
- D. Install an Application Performance Monitoring (APM) tool in both locations, and configure an export to a central data storage location for analysis.

**Answer:** A

**NEW QUESTION 40**

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 43**

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 server
- 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 45**

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 50**

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 51**

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 in testing 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 55**

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 56**

You support a service that recently had an outage. The outage was caused by a new release that exhausted the service memory resources. You rolled back the release successfully to mitigate the impact on users. You are now in charge of the post-mortem for the outage. You want to follow Site Reliability Engineering practices when developing the post-mortem. What should you do?

- A. Focus on developing new features rather than avoiding the outages from recurring.
- B. Focus on identifying the contributing causes of the incident rather than the individual responsible for the cause.
- C. Plan individual meetings with all the engineers involved
- D. Determine who approved and pushed the new release to production.
- E. Use the Git history to find the related code commit
- F. Prevent the engineer who made that commit from working on production services.

**Answer:** B

**NEW QUESTION 61**

You support a web application that runs on App Engine and uses CloudSQL and Cloud Storage for data storage. After a short spike in website traffic, you notice a big increase in latency for all user requests, increase in CPU use, and the number of processes running the application. Initial troubleshooting reveals: After the initial spike in traffic, load levels returned to normal but users still experience high latency. Requests for content from the CloudSQL database and images from Cloud Storage show the same high latency.

No changes were made to the website around the time the latency increased. There is no increase in the number of errors to the users.

You expect another spike in website traffic in the coming days and want to make sure users don't experience latency. What should you do?

- A. Upgrade the GCS buckets to Multi-Regional.
- B. Enable high availability on the CloudSQL instances.
- C. Move the application from App Engine to Compute Engine.
- D. Modify the App Engine configuration to have additional idle instances.

**Answer:** D

**Explanation:**

Scaling App Engine scales the number of instances automatically in response to processing volume. This scaling factors in the automatic\_scaling settings that are provided on a per-version basis in the configuration file. A service with basic scaling is configured by setting the maximum number of instances in the max\_instances parameter of the basic\_scaling setting. The number of live instances scales with the processing volume. You configure the number of instances of each version in that service's configuration file. The number of instances usually corresponds to the size of a dataset being held in memory or the desired throughput for offline work. You can adjust the number of instances of a manually-scaled version very quickly, without stopping instances that are currently running, using the Modules API set\_num\_instances function. <https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed> <https://cloud.google.com/appengine/docs/standard/python/config/appref> max\_idle\_instances Optional. The maximum number of idle instances that App Engine should maintain for this version. Specify a value from 1 to 1000. If not specified, the default value is automatic, which means App Engine will manage the number of idle instances. Keep the following in mind: A high maximum reduces the number of idle instances more gradually when load levels return to normal after a spike. This helps your application maintain steady performance through fluctuations in request load, but also raises the number of idle instances (and consequent running costs) during such periods of heavy load.

**NEW QUESTION 65**



.....

## Relate Links

**100% Pass Your Professional-Cloud-DevOps-Engineer Exam with Exambible Prep Materials**

<https://www.exambible.com/Professional-Cloud-DevOps-Engineer-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>