

Professional-Cloud-DevOps-Engineer PDF Demo & Professional-Cloud-DevOps-Engineer Pruefungssimulationen



Übrigens, Sie können die vollständige Version der DeutschPrüfung Professional-Cloud-DevOps-Engineer Prüfungsfragen aus dem Cloud-Speicher herunterladen: <https://drive.google.com/open?id=1BOIwuAGc8LAcihEup-kFzfYJAXLXjkVq>

Die Google Professional-Cloud-DevOps-Engineer Zertifizierungsprüfung ist sehr populär in IT-Industrie. Es spielt eine übergreifende Bedeutung für die Leute, die ihre Arbeitsstelle erhöhen wollen. Und es ist auch die Wahl, die Leute klar sehen können. Außerdem dadurch können Sie Ihre Fähigkeit verbessern und mehr verwendbare Technik beherrschen. Damit können Sie Ihre Arbeit besser fertigen und auch anderen Ihre Fähigkeit zeigen.

Die Prüfung konzentriert sich auf verschiedene Aspekte von DevOps, einschließlich kontinuierlicher Lieferung, Infrastruktur als Code, Überwachung und Protokollierung sowie Sicherheit. Kandidaten, die die Prüfung bestehen, werden als Experten anerkannt, die DevOps-Workflows auf der Google Cloud Plattform implementieren und verwalten können. Sie sind auch in der Lage, hoch skalierbare und zuverlässige Cloud-basierte Anwendungen zu entwerfen und umzusetzen.

>> Professional-Cloud-DevOps-Engineer PDF Demo <<

Hilfsreiche Prüfungsunterlagen verwirklicht Ihren Wunsch nach der Zertifikat der Google Cloud Certified - Professional Cloud DevOps Engineer Exam

Wie kann man Erfolge erlangen. Es gibt nur eine Ankürzung, nämlich: die Lernhilfe zur Google Professional-Cloud-DevOps-EngineerZertifizierungsprüfung von DeutschPrüfung zu benutzen. Das ist unser Vorschlag für jeden Kandidaten. Wir hoffen, dass Sie Ihren Traum erfüllen können.

Google Cloud Certified - Professional Cloud DevOps Engineer Exam Professional-Cloud-DevOps-Engineer Prüfungsfragen mit Lösungen (Q194-Q199):

194. Frage

You need to define SLOs for a high-traffic web application. Customers are currently happy with the application performance and availability. Based on current measurement, the 90th percentile Of latency is 160 ms and the 95th percentile of latency is 300 ms over a 28-day window. What latency SLO should you publish?

- A. 90th percentile - 150 ms
95th percentile - 290 ms
- B. 90th percentile - 160 ms
95th percentile - 300 ms
- C. 90th percentile - 300 ms

- 95th percentile - 450 ms
- D. 90th percentile - 190 ms
- 95th percentile - 330 ms

Antwort: B

Begründung:

a latency SLO is a service level objective that specifies a target level of responsiveness for a web application¹. A latency SLO can be expressed as a percentile of latency over a time window, such as the 90th percentile of latency over 28 days². A percentile of latency is the maximum amount of time that a given percentage of requests take to complete. For example, the 90th percentile of latency is the maximum amount of time that 90% of requests take to complete³.

To define a latency SLO, you need to consider the following factors²⁴:

The expectations and satisfaction of your customers. You want to set a latency SLO that reflects the level of performance that your customers are happy with and willing to pay for.

The current and historical measurements of your latency. You want to set a latency SLO that is based on data and realistic for your web application.

The trade-offs and costs of improving your latency. You want to set a latency SLO that balances the benefits of faster response times with the costs of engineering work, infrastructure, and complexity.

Based on these factors, the best option for defining a latency SLO for your web application is option B. Option B sets the latency SLO to match the current measurement of your latency, which means that you are meeting the expectations and satisfaction of your customers. Option B also sets a realistic and achievable target for your web application, which means that you do not need to invest extra resources or effort to improve your latency. Option B also aligns with the best practice of setting conservative SLOs, which means that you have some buffer or margin for error in case your latency fluctuates or degrades⁵.

195. Frage

You are the on-call Site Reliability Engineer for a microservice that is deployed to a Google Kubernetes Engine (GKE) Autopilot cluster. Your company runs an online store that publishes order messages to Pub/Sub and a microservice receives these messages and updates stock information in the warehousing system. A sales event caused an increase in orders, and the stock information is not being updated quickly enough. This is causing a large number of orders to be accepted for products that are out of stock. You check the metrics for the microservice and compare them to typical levels.

Microservice metrics	Typical state	Current state
Average CPU across all Pods	20% of Pod limit	30% of Pod limit
Average memory across all Pods	10% of Pod limit	10% of Pod limit
Pub/Sub subscription: Average oldest unacknowledged message age	347 milliseconds	8074 milliseconds
Pub/Sub subscription: Average undelivered messages	5 messages	14705 messages
Pub/Sub subscription: Average acknowledgment latency	312 milliseconds	354 milliseconds

You need to ensure that the warehouse system accurately reflects product inventory at the time orders are placed and minimize the impact on customers. What should you do?

- A. Add a virtual queue to the online store that allows typical traffic levels
- **B. Increase the number of Pod replicas**
- C. Increase the Pod CPU and memory limits
- D. Decrease the acknowledgment deadline on the subscription

Antwort: B

Begründung:

The best option for ensuring that the warehouse system accurately reflects product inventory at the time orders are placed and minimizing the impact on customers is to increase the number of Pod replicas. Increasing the number of Pod replicas will increase the scalability and availability of your microservice, which will allow it to handle more Pub/Sub messages and update stock information faster. This way, you can reduce the backlog of undelivered messages and oldest unacknowledged message age, which are causing delays in updating product inventory. You can use Horizontal Pod Autoscaler or Cloud Monitoring metrics-based autoscaling to automatically adjust the number of Pod replicas based on load or custom metrics.

196. Frage

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

Antwort: A

Begründung:

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.

197. Frage

You are monitoring a service that uses n2-standard-2 Compute Engine instances that serve large files. Users have reported that downloads are slow. Your Cloud Monitoring dashboard shows that your VMS are running at peak network throughput. You want to improve the network throughput performance. What should you do?

- **A. Change the machine type for your VMS to n2-standard-8.**
- B. Deploy a Cloud NAT gateway and attach the gateway to the subnet of the VMS.
- C. Add additional network interface controllers (NICs) to your VMS.
- D. Deploy the Ops Agent to export additional monitoring metrics.

Antwort: A

Begründung:

The correct answer is C, Change the machine type for your VMs to n2-standard-8.

According to the Google Cloud documentation, the network throughput performance of a Compute Engine VM depends on its machine type¹. The n2-standard-2 machine type has a maximum egress bandwidth of 4 Gbps, which can be a bottleneck for serving large files. By changing the machine type to n2-standard-8, you can increase the maximum egress bandwidth to 16 Gbps, which can improve the network throughput performance and reduce the download time for users. You also need to enable per VM Tier_1 networking performance, which is a feature that allows VMs to achieve higher network performance than the default settings².

The other options are incorrect because they do not improve the network throughput performance of your VMs. Option A is incorrect because Cloud NAT is a service that allows private IP addresses to access the internet, but it does not increase the network bandwidth or speed³. Option B is incorrect because adding additional network interfaces (NICs) or IP addresses per NIC does not increase ingress or egress bandwidth for a VM¹. Option D is incorrect because deploying the Ops Agent can help you monitor and troubleshoot your VMs, but it does not affect the network throughput performance⁴.

Reference:

Cloud NAT overview, Cloud NAT overview. Network bandwidth, Bandwidth summary. Installing the Ops Agent, Installing the Ops Agent. Configure per VM Tier_1 networking performance, Configure per VM Tier_1 networking performance.

198. Frage

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 - 150ms 95th percentile - 300ms**
- C. 90th percentile - 250ms 95th percentile - 400ms
- D. 90th percentile - 120ms 95th percentile - 275ms

Antwort: B

Begründung:

• • • • •

Professional-Cloud-DevOps-Engineer Pruefungssimulationen: <https://www.deutschpruefung.com/Professional-Cloud-DevOps-Engineer-deutsch-pruefungsfragen.html>

- [illegible]

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,
myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, Disposable vapes

Außerdem sind jetzt einige Teile dieser DeutschPrüfung Professional-Cloud-DevOps-Engineer Prüfungsfragen kostenlos erhältlich:
<https://drive.google.com/open?id=1BOIwuAGc8LAcihEup-kFzfYJAxLXjkVq>