

KillTest

質量更高 服務更好



練習題

<https://www.killtest.net>

一年免費更新服務

**Exam : Professional Cloud
DevOps Engineer**

**Title : Google Cloud Certified -
Professional Cloud DevOps
Engineer Exam**

Version : DEMO

1.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

2.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 ms95th percentile - 290 ms
- B. 90th percentile - 160 ms95th percentile - 300 ms
- C. 90th percentile - 190 ms95th percentile - 330 ms
- D. 90th percentile - 300 ms95th percentile - 450 ms

Answer: B

Explanation:

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 factors⁴:

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.

3.You are creating a CI/CD pipeline in Cloud Build to build an application container image The application code is stored in GitHub Your company requires that production image builds are only run against the main branch and that the change control team approves all pushes to the main branch You

want the image build to be as automated as possible.

What should you do? Choose 2 answers

- A. Create a trigger on the Cloud Build job Set the repository event setting to Pull request'
- B. Add the owners file to the Included files filter on the trigger
- C. Create a trigger on the Cloud Build job Set the repository event setting to Push to a branch
- D. Configure a branch protection rule for the main branch on the repository
- E. Enable the Approval option on the trigger

Answer: C,D

Explanation:

The best options for creating a CI/CD pipeline in Cloud Build to build an application container image and ensuring that production image builds are only run against the main branch and that the change control team approves all pushes to the main branch are to create a trigger on the Cloud Build job, set the repository event setting to Push to a branch, and configure a branch protection rule for the main branch on the repository. A trigger is a resource that starts a build when an event occurs, such as a code change. By creating a trigger on the Cloud Build job and setting the repository event setting to Push to a branch, you can ensure that the image build is only run when code is pushed to a specific branch, such as the main branch. A branch protection rule is a rule that enforces certain policies on a branch, such as requiring reviews, status checks, or approvals before merging code. By configuring a branch protection rule for the main branch on the repository, you can ensure that the change control team approves all pushes to the main branch.

4.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 – 100ms95th percentile – 250ms
- B. 90th percentile – 120ms95th percentile – 275ms
- C. 90th percentile – 150ms95th percentile – 300ms
- D. 90th percentile – 250ms95th percentile – 400ms

Answer: C

Explanation:

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

5.Your company runs services by using multiple globally distributed Google Kubernetes Engine (GKE) clusters Your operations team has set up workload monitoring that uses Prometheus-based tooling for metrics alerts: and generating dashboards This setup does not provide a method to view metrics globally across all clusters You need to implement a scalable solution to support global Prometheus querying and minimize management overhead.

What should you do?

- A. Configure Prometheus cross-service federation for centralized data access
- B. Configure workload metrics within Cloud Operations for GKE
- C. Configure Prometheus hierarchical federation for centralized data access

D. Configure Google Cloud Managed Service for Prometheus

Answer: D

Explanation:

The best option for implementing a scalable solution to support global Prometheus querying and minimize management overhead is to use Google Cloud Managed Service for Prometheus. Google Cloud Managed Service for Prometheus is a fully managed service that allows you to collect, query, and visualize metrics from your GKE clusters using Prometheus-based tooling. You can use Google Cloud Managed Service for Prometheus to query metrics across multiple clusters and regions using a global view. You can also use Google Cloud Managed Service for Prometheus to integrate with other Google Cloud services, such as Cloud Monitoring, Cloud Logging, and BigQuery. By using Google Cloud Managed Service for Prometheus, you can avoid managing and scaling your own Prometheus servers and focus on your application performance.