

Szkolenie: Google Cloud Site Reliability Engineering Measuring and Managing Reliability



Cel szkolenia:

This one-week course equips students to quantify the reliability of the applications they build in Google Cloud Platform, to assess the risks to these applications' reliability, and to use this data to drive decision-making when prioritizing engineering work.

Course objectives:

This course teaches the theory of Service Level Objectives (SLOs), a principled way of describing and measuring the desired reliability of a service. Upon completion, learners should be able to apply these principles to develop the first SLOs for services they are familiar with in their own organizations. Learners will also learn how to use Service Level Indicators (SLIs) to quantify reliability and Error Budgets to drive business decisions around engineering for greater reliability. The learner will understand the components of a meaningful SLI and walk through the process of developing SLIs and SLOs for an example service.

Audience:

This class is primarily intended for the following participants: DevOps specialists Software developers Product managers and application owner's IT business decision makers.

Plan szkolenia:

- Introduction to SRE
 - This module is intended to bring you up to speed on the concepts underpinning SRE, CRE, and SLOs. If you're already familiar with these concepts, you may still find new information and perspectives in this module, but it is not necessary to complete it.
- Targeting Reliability
 - In this module we're going to talk about how you measure the desired reliability of a service. We will address what to consider when setting SLOs for your application within your organization. We'll look at the three principles we use to measure the desired reliability of a service: figuring out what you want to promise and to whom, figuring out the metrics you care about that make your service reliability "good", and finally, deciding how much reliability is good enough.
- Operating for Reliability
 - In this module, we'll start by introducing a mechanism for quantifying unreliability using something called an error budget. We'll show how error budgets help you decide when to focus on making a service more reliable. And then we'll learn about some of the

engineering and operational improvements that can help you do that.

- Choosing a Good SLI
 - In this module we will start off by taking a look at some characteristics of monitoring metrics that can make them useful as SLIs and contrast these against other metrics that are less useful. Because the choice of where to measure an SLI is a key variable, we'll cover the five main ways you can measure an SLI and compare their pros and cons.
- Developing SLOs and SLIs
 - In this module, we'll start off with an overview of our four step process for developing SLOs and SLIs for a user journey. We'll introduce the fictional company that created our example mobile game, the infrastructure that we'll be working with, and the simple user journey we'll be applying the four step process to.
- Quantifying Risks to SLOs
 - In this module we'll be taking a critical look at the availability risks for our example service. We want to answer the question: "are our SLO targets and error budgets realistic?"
- Consequences of SLO Misses
 - In this module, we'll cover best practices for documenting your SLOs, the rationale behind a formal error budget policy and how best to create one and finally, we'll look at an example error budget policy in order to understand the trade-offs and incentives that play out during negotiations when trying to write an error budget policy.

Poziom trudności



Certyfikaty:

The participants will obtain certificates signed by Google Cloud.

Prowadzący:

Authorized Google Cloud Trainer.