

# Logging, Monitoring and Observability in Google Cloud

**Course 4323– 24 Hours**

## Overview

Learn how to monitor, troubleshoot, and improve your infrastructure and application performance. Guided by the principles of Site Reliability Engineering (SRE), this course features a combination of lectures, demos, hands-on labs, and real-world case studies. In this course, you'll gain experience with full-stack monitoring, real-time log management and analysis, debugging code in production, and profiling CPU and memory usage.

## On Completion, Delegates will be able to

- Plan and implement a well-architected logging and monitoring infrastructure
- Define service level indicators (SLIs) and service level objectives (SLOs)
- Create effective monitoring dashboards and alerts
- Monitor, troubleshoot, and improve Google Cloud infrastructure
- Analyze and export Google Cloud audit logs
- Find production code defects, identify bottlenecks, and improve performance
- Optimize monitoring costs

## Who Should Attend

- Cloud architects, administrators, and SysOps personnel
- Cloud developers and DevOps personnel

## Prerequisites

To get the most out of this course, participants should have:

- Google Cloud Fundamentals: Core Infrastructure or equivalent experience
- Basic scripting or coding familiarity
- Proficiency with command-line tools and Linux operating system environments



## Course Contents

---

### **Introduction to Monitoring in Google Cloud**

In this module, we will take some time to do a high-level overview of the various products which comprise Google Cloud's logging, monitoring, and observability suite.

### **Avoiding Customer Pain**

In this module, we discuss several Site Reliability Engineering (SRE) concepts and how we can use them to help avoid customer pain. In this context, a customer is any consumer of a cloud-based system.

### **Alerting Policies**

Alerting gives timely awareness to problems in your cloud applications so you can resolve the problems quickly. In this module, you will learn how to develop alerting strategies, define alerting policies, add notification channels, identify types of alerts and common uses for each, construct and alert on resource groups, and manage alerting policies programmatically.

### **Monitoring Critical Systems**

Monitoring is all about keeping track of exactly what's happening with the resources we've spun up inside of Google's Cloud. In this module, we'll take a look at options and best practices as they relate to monitoring project architectures. We'll differentiate the core Cloud IAM roles needed to decide who can do what as it relates to monitoring. Just like architecture, this is another crucial early step. We will examine some of the Google created default dashboards, and see how to use them appropriately. We will create charts and use them to build custom dashboards to show resource consumption and application load. And, finally, we will define uptime checks to track liveliness and latency.

### **Configuring Google Cloud Services for Observability**

In the next part of our Metrics discussion, let's take a little time to examine the art of Configuring Google Cloud Services for Observability. In this module, we're going to spend a little time learning how to integrate logging and monitoring agents into Compute Engine VMs and images using Agents, enable and utilize Kubernetes Monitoring, extend and clarify Kubernetes monitoring with Prometheus, and expose custom metrics through code, and with the help of OpenCensus.

## **Advanced Logging and Analysis**

In this module, we will examine some of Google Cloud's advanced logging and analysis capabilities. Specifically, in this module you will learn to identify and choose among resource tagging approaches, define log sinks, create monitoring metrics based on log entries, link application errors to Logging and other operation tools using Error Reporting, and export logs to BigQuery for long term storage and SQL based analysis.

## **Monitoring Network Security and Audit Logs**

In this module, we will examine two key topics: Monitoring as it relates to the VPC network, and how to use Google's Cloud Audit logs. You will learn to collect and analyze VPC Flow, Firewall Rule, and Cloud NAT logs, enable Packet Mirroring, explain the capabilities of the Network Intelligence Center, and use Cloud Audit logs to answer the question, "Who, did what, and when?" We will also cover best practices for Audit Logging.

## **Managing Incidents**

Up to this point in our course, we've mostly focused on ways to inspect and monitor the status of our systems running in Google Cloud. But no matter how solid your planning, design, architecture, and preventive maintenance strategies are, things will go wrong. When they do go wrong, how you manage those incidents will have a huge impact on user perception. In this module, you will learn how to handle incidents using a systematic process.

## **Investigating Application Performance Issues**

When deploying applications to Google Cloud, the Application Performance Management products (Cloud Trace, Cloud Debugger, and Cloud Profiler) provide a suite of tools to give insight into how your code and services are functioning, and to help troubleshoot where needed.

## **Optimizing the Costs of Monitoring**

In our final module we discuss optimizing the costs for Google Cloud's operations suite. Specifically, you will learn to analyze resource utilization costs for operations related components within Google Cloud, and implement best practices for controlling the cost of operations within Google Cloud.