



Site Reliability Engineering (SRE) and Observability: SLOs, Error Budgets and Incident Response



Site Reliability Engineering (SRE) and Observability: SLOs, Error Budgets and Incident Response

Introduction:

Production software services fail users in ways dashboards miss: pages fire on noise, releases break checkout, and nobody can agree whether the service is reliable enough to ship faster. Site Reliability Engineering replaces opinion with measured targets, so teams set SLIs and SLOs, spend error budgets deliberately and cut repetitive operational toil. This Core Concept course trains engineers to instrument services with metrics, logs and traces, alert on user-facing symptoms, run incidents and blameless postmortems, and release changes safely. Participants build a Service Reliability Plan for a service they operate.

Course Objectives:

- Define SLIs and SLOs for a user-facing service and write an error budget policy that governs release decisions
- Instrument a service with metrics, logs and traces using OpenTelemetry and query them to explain abnormal behaviour
- Design symptom-based alerts on SLO burn rate and organise a sustainable on-call rotation
- Lead a production incident with defined roles and write a blameless postmortem with tracked action items
- Measure operational toil and select automation, capacity and canary release practices that reduce reliability risk
- Produce a Service Reliability Plan ready for review by engineering and product leads

Target Audience:

- DevOps engineers who build and maintain deployment pipelines and production environments
 - Platform engineers who provide shared infrastructure, tooling and runtime services to product teams
 - Operations and infrastructure engineers who monitor services and respond to outages
 - Software engineers who carry on-call duty for the services they develop
 - Technical team leads who coordinate releases and incident handling for a product area
-

Course Outline:

Day 1: SRE Principles and Service Reliability Baseline

- SRE Origins, Responsibilities and the Relationship between SRE and DevOps
- Embracing Risk: Availability Targets and the Cost of Additional Nines
- Toil Definition and Operational Work Inventory for a Service Team
- Simplicity Principles for Production Systems and Change Management
- Service Reliability Baseline Assessment of a Participant Service

Day 2: SLIs, SLOs, Error Budgets and Observability Signals

- SLI Selection by Service Type: Request, Pipeline and Storage Services
- SLO Document Structure and Target Setting from Historical Data
- Error Budget Policy: Release Freezes, Exceptions and Escalation
- OpenTelemetry Signals: Traces, Metrics and Logs with Shared Context
- Monitoring Distributed Systems: Four Golden Signals Dashboard Design

Day 3: Instrumentation, Alerting and On-Call Practice

- Lab: OpenTelemetry Instrumentation and Collector Pipeline Configuration
- Lab: Distributed Trace Analysis to Locate Latency in a Request Path
- Alerting on SLOs with Multi-Window Burn Rate Rules
- On-Call Rotation Design, Handover Notes and Pager Load Review
- Incident Command Roles, Severity Levels and Communication Channels

Day 4: Postmortems, Capacity, Releases and Resilience Testing

- Blameless Postmortem Template, Timeline Reconstruction and Action Tracking
- Capacity Planning with Demand Forecasts and Load Testing Results
- Overload Handling: Load Shedding, Retries with Backoff and Graceful Degradation
- Canarying Releases and Progressive Delivery with Automated Rollback
- Chaos Engineering Experiments and Container Orchestration Health Checks

Day 5: Lab Scenarios and the Service Reliability Plan

- Lab Scenario: Latency Regression Traced from Alert to Root Cause
 - Lab Scenario: Failed Canary Release and Error Budget Decision
 - SRE Engagement Model and Team Lifecycle Options for the Organisation
 - Service Reliability Plan Drafting: SLO Document, Alerts and Runbooks
 - Service Reliability Plan Peer Review and Presentation
-

Skills You Will Gain:

- Service Level Objective Design
- Error Budget Management
- Telemetry Instrumentation
- Distributed Tracing Analysis
- Burn Rate Alerting
- Incident Command
- Toil Reduction
- Capacity Planning

Why Attend This Course:

- Return with a Service Reliability Plan containing an SLO document, error budget policy, alert rules and postmortem template for a service you operate
- Cut noisy pages by replacing threshold alerts with alerts tied to user-facing reliability
- Gain hands-on practice instrumenting, tracing and debugging a service in a lab environment
- Compare on-call, release and incident practices with engineers from other sectors and organisations

Conclusion:

Reliable services come from agreed targets, useful telemetry and disciplined response rather than heroics. The course moves from SRE principles and a reliability baseline, through SLIs, SLOs, error budgets and observability signals, to hands-on instrumentation, burn rate alerting and on-call practice, and then to postmortems, capacity, overload handling, canary releases and chaos experiments. The final day produces a Service Reliability Plan that gives each team measured targets and repeatable procedures for running its production services.
