
Site Reliability Engineering Leadership: Service Targets, Error Budgets and Incident Review Course

31 May – 4 June 2027
Cape town





Site Reliability Engineering Leadership: Service Targets, Error Budgets and Incident Review Course

Ref: 447_25688 Date: 31 May – 4 June 2027 Location: Cape town Fees: 4800 GBP

Programme Summary

Site reliability engineering leadership is the subject of this one-week training course for heads of engineering, platform and operations, ending in a service reliability plan for a portfolio of live services. Many technology functions chase uptime without agreed targets, so release pressure and reliability work compete without a rule to settle them. The programme moves from service targets through diagnosis of budget use, toil and on-call load to policy design, and then to application on real services. Mars Training Centre UK delivers this programme on site reliability engineering leadership for technology leaders from mixed organisations.

The Business Case

Digital services now carry revenue, customer trust and operational continuity in banking, retail, telecommunications and government. When reliability is argued case by case, outages repeat, engineers burn out on interrupt-driven work and executives cannot see what reliability costs or buys. Leaders who set measurable targets, agree what happens when they are missed and learn openly from incidents turn reliability into a managed decision rather than a source of conflict.

- Service Availability: Improves as targets are set from user need and tracked against an agreed budget.
- Engineering Capacity: Is released as toil is measured and removed rather than absorbed.
- Incident Recurrence: Falls as reviews lead to tracked actions with named owners.

Principles and Evidence Base

- The Site Reliability Engineering book published by Google: Its principles on service level indicators, objectives and toil frame the target setting and diagnosis in Sessions 1 and 2.
- The Site Reliability Engineering workbook and its example error budget policy: It supplies the policy structure that participants adapt in Session 3.
- Blameless postmortem practice described in the same body of work: It shapes the incident review format used in Session 4.

Intended Participants

- Engineering heads who answer for the availability of customer-facing services.
- Platform and infrastructure leaders who decide where reliability effort is spent.
- Operations managers who run on-call rotations and incident response.
- Technology directors who report service health and risk to executives.

Typical participants come from banking, retail, telecommunications and public digital services.



Learning Outcomes

1. Participants will be able to construct service level indicators and objectives that reflect user and business need.
2. Participants will be able to evaluate error budget use and diagnose its main causes.
3. Participants will be able to prioritise toil and on-call load for removal.
4. Participants will be able to justify an error budget policy and its release rules to executives.
5. Participants will be able to organise incident command and communication for a major outage.
6. Participants will be able to assess postmortem findings and rank the actions that follow.

Programme Structure

Day 1: Site Reliability Engineering Leadership Principles and Service Targets

How are reliability targets set so that they reflect what users and the business need?

- Service Level Indicator Selection Sheet
- Service Level Objective Setting Method
- Reliability Versus Feature Velocity Trade-Off Map
- Service Criticality Tiering Table
- Reliability Ownership Charter

Day 2: Diagnosing Error Budget Burn, Toil and On-Call Load

Where does reliability work drain engineering time, and where is the budget being spent?

- Error Budget Burn Rate Reading
- Toil Inventory and Measurement Log
- On-Call Load and Fatigue Review
- Monitoring Signal Quality Audit
- Alert Noise Reduction Review

Day 3: Designing Error Budget Policy and Release Decisions

What happens when a service uses up its error budget?

- Error Budget Policy Draft
- Release Gate and Freeze Rule Set
- Capacity and Resilience Planning Review
- Reliability Investment Case Template
- Executive Reliability Report Layout



Day 4: Leading Incident Response and Postmortem Learning

How does a leader run an incident and turn it into lasting change?

- Incident Command Role Card
- Severity Classification Table
- Communication Update Template
- Blameless Postmortem Format
- Action Item Tracking Register

Day 5: Applying Site Reliability Engineering Leadership to One Service Portfolio

Is the reliability plan ready to take to the technology leadership team?

- Applying Session 1 Targets to the Chosen Services
- Testing Session 2 Burn Readings on Past Periods
- Applying the Session 3 Policy to a Release Scenario
- Rehearsing the Session 4 Incident Roles and Review
- Completing the Service Reliability Plan

Methodology

The programme uses a structured exercise method built on anonymised service data, outage timelines and draft policies from banking and retail platforms. Participants set targets, read budget use, draft rules and rehearse incident roles in small groups, then defend their choices to peers acting as executives and product owners. Facilitator feedback is given on every work product before it enters the final plan.

- Suggested activity: Choosing indicators and a target for three services from a short description of user journeys.
- Suggested activity: Reading a burn chart and deciding what the policy should require at each stage.
- Suggested activity: Running a timed incident simulation and writing the first draft of its review.

Capstone Output

Each participant completes a service reliability plan for a portfolio of live services. It holds service level indicators and objectives, an error budget policy with release rules, a toil and on-call reduction list, an incident roles and communication guide, and a postmortem action register. The plan is used at work to agree targets with product teams, brief executives and track reliability work.



Frequently Asked Questions

Is site reliability engineering leadership suitable for managers who do not write code every day?

Yes. Site reliability engineering leadership is taught for people who direct engineering and operations, and it assumes familiarity with running digital services rather than hands-on programming.

Which sources does site reliability engineering leadership rest on in this programme?

Site reliability engineering leadership here rests on the Site Reliability Engineering book and workbook published by Google, covering service level objectives, error budget policy, toil and blameless postmortems.

How is the service reliability plan from site reliability engineering leadership used at work?

You use the service reliability plan to agree targets with product teams, brief executives on risk and track the actions that raise reliability.

How does site reliability engineering leadership differ from a programme on IT service management?

Site reliability engineering leadership centres on measurable targets, error budgets and engineering time, while IT service management centres on process frameworks for requests, changes and support.