REFERENCE REVIEWED 1 NOV 2026

Drill: budget 80% spent, 10 days left

Forecast the month, size what is left of the error budget, and decide what to do about next week's launches. About 30 minutes. Built on the error budget policy guide.

01

How to run this drill

- One facilitator and two to six participants, ideally including someone who owns the product roadmap. Print the page: give out the first sheet, keep the facilitator notes on the second.
- Give the group 15 minutes for the three questions, with a calculator or the error budget calculator, then go through the notes.
- No scores and no record of who said what. Finish by asking what your own error budget policy would say on day 21.

02

Handout

Your service has a 99.9% availability SLO, measured over a 30-day calendar month that resets on the 1st. It is the morning of day 21. The dashboard shows 80% of the month's error budget spent. Traffic is about 500 requests a second. Two features are due to launch next week. Your error budget policy says changes freeze only when the budget is exhausted.

1. If the rest of the month goes like the first 20 days, does the service meet its SLO?
2. How much failure is left, in failed requests and in minutes of full outage?
3. What would you do this week, and who needs to agree?

FACILITATOR NOTES

1. **No.** The average burn so far is above 1:

   `80% × 30 d ÷ 20 d = 1.2×, ending the month at 120% of budget; it runs out 5d from now, at the end of day 25`

2. **What is left:**

   `20% × 1,296,000 = 259,200 failed requests, or 20% × 43m 12s = 8m 38s of full outage`
   To finish within budget the last 10 days must average at most 0.6×, 50% of the pace so far. One bad deploy could spend the rest.

3. **Look before deciding.** Was the 80% one big incident or a steady leak? A single outage that is fixed changes the forecast; a steady leak does not. Find the largest consumers of budget.

4. **Use the policy, and its gaps.** On a graduated policy, 75% or more spent means extra review for risky changes and reliability work first. A policy that acts only at 100% gives no guidance here; agree a response with the product owner rather than improvising one.

5. **Launches.** Options: delay one, launch behind a flag to a small share, or launch with a tested rollback and a person watching. Whoever owns the policy decides, not the on-call engineer.

> Twist: on a 30-day rolling window there is no “10 days left”: errors from 30 days ago drop out daily, so the question becomes how fast the budget is being spent now. Ask which kind of window your own SLOs use.

Sources: Google, Site Reliability Engineering, 1st ed., ch. 28 “Accelerating SREs to On-Call and Beyond” (disaster role playing, “Wheel of Misfortune”, in which a game master plays out an outage live while one or two on-callers ask questions and act; this drill adapts that format as a discussion from a handout rather than following it); The Site Reliability Workbook, ch. 2 “Implementing SLOs” (rolling and calendar windows, error budget policies). The scenario is invented for practice.