SITE RELIABILITY ENGINEERING

The SRE Frameworks Field Guide

A practical reference to all six core SRE frameworks — with worked examples, decision trees, comparison matrices, an error-budget calculator, a postmortem template, and a maturity self-assessment you can use Monday morning.

- 6 Framework Deep-Dives
- SLI/SLO Worked Examples
- Error Budget Calculator
- Decision Trees & Matrices
- Maturity Self-Assessment

Global Skill Development Council
Vendor-neutral · Internationally recognized credentialing body

Trusted by 75,000+ professionals worldwide
www.gsdcouncil.org

HOW TO USE THIS GUIDE

Six frameworks, one reliable system

SRE is not one big idea — it's six interlocking frameworks that turn "make it reliable" into measurable, repeatable practice. This field guide gives each a deep-dive plus the practical artifacts to apply it: worked examples, decision trees, matrices, and templates.

#	Framework	What it answers
1	**SLI / SLO / SLA**	How reliable are we, and what do we promise?
2	**Error Budgets**	How much unreliability can we afford to spend?
3	**Toil Reduction**	What manual work should we automate away?
4	**Observability & Golden Signals**	How do we see what the system is doing?
5	**Incident Management**	How do we respond when things break?
6	**Blameless Postmortems**	How do we learn so it doesn't recur?

Read it cover to cover

The frameworks build on each other — measurement first, learning last.

Or jump to a tool

Need the calculator or template today? Each lives on its own page.

Then go deeper

Every framework maps to a module in the GSDC SRE certification.

Framework content reflects established SRE practice (Google SRE book & workbook) and current 2025–2026 reliability engineering guidance.

FRAMEWORK DEEP-DIVE · 1 OF 6

Service levels: SLI, SLO, SLA

1

SLI / SLO / SLA

THE MEASUREMENT FRAMEWORK

WHAT

A layered way to express reliability in numbers, not adjectives — measure it, target it, and promise it.

WHY

Turns "make it reliable" into a shared contract between engineering and the business.

Term	Definition	Example
SLI — Indicator	A quantitative measure of what users experience	% of requests returning success
SLO — Objective	A target value for an SLI over a time window	99.9% success over 30 days
SLA — Agreement	A contract with consequences for missing an SLO	Credits if availability < 99.9%

The golden rule of good SLIs

- ✓ **Measure the user's experience**, not internal noise — track impact, not every metric.
- ✓ **Keep it few** — the most common mistake is tracking too many SLIs.
- ✓ **Count what matters** — for availability, only server errors (5xx) should burn budget, not client errors (4xx).
- ✓ **Use percentiles for latency** — p50 for typical, p99 (or p99.9) for the tail.

SLOs sit slightly below 100% on purpose — the gap is your error budget (Framework 2).

WORKED EXAMPLES

SLI & SLO, step by step

Scenario: you run an online **checkout API**. The business says it must be "fast and reliable." Reliability needs math — here's how to turn that into SLIs and SLOs.

1 · Availability SLI

```
availability = successful_requests ÷ total_requests
```

Definition: "Percentage of HTTP requests returning 2xx or expected 4xx." Only unexpected 5xx server errors count against you. **SLO: 99.9% successful over 30 days.**

2 · Latency SLI

```
latency_sli = requests_under_threshold ÷ total_requests
```

Definition: "Percentage of requests served faster than 300 ms." Track at multiple percentiles. **SLO: 99% of requests < 300 ms (p99) over 30 days.**

3 · Putting it together

SLI	Measurement	SLO target
Availability	Successful ÷ total requests	**99.9% / 30 days**
Latency (p99)	Requests under 300 ms	**99.0% / 30 days**
Freshness	Data updated within target window	**99.5% / 30 days**

Pro tip — don't guess your targets

Run the service for 2–4 weeks and measure actual performance first, then set SLOs slightly above the floor your users tolerate. Targets that are too tight create needless toil; too loose, and users hurt before anyone reacts.

FRAMEWORK DEEP-DIVE · 2 OF 6

Error budgets: reliability as currency

2

Error Budgets

THE DECISION FRAMEWORK

WHAT

The gap between your SLO and 100% — the unreliability you're allowed to "spend" on velocity.

WHY

Ends the dev-vs-ops tug-of-war: ship fast while budget is healthy, slow down when it runs low.

```
error budget = 100% − SLO
```

The four-level error budget policy

Budget remaining	Action
> 50%	Ship normally — take healthy risks.
25–50%	Increased review on changes.
10–25%	Feature freeze — reliability work only.
< 10%	Emergency mode — all hands on reliability.

This makes release decisions automatic and objective: teams stop arguing and rely on the budget math.

ERROR BUDGET CALCULATOR

Turn SLOs into minutes & decisions

Downtime allowance by SLO (per 30-day month)

SLO	Allowed unreliability	≈ Downtime / month
99%	1%	7h 18m
99.5%	0.5%	3h 39m
99.9% ("three nines")	0.1%	43.2 min
99.95%	0.05%	21.6 min
99.99% ("four nines")	0.01%	4.3 min

Worked example — did we blow the budget?

1. **Setup:** Checkout API has a 99.9% monthly SLO → budget = 43.2 min of "full-downtime equivalent."
2. **Incident:** 20% of requests failed for 15 minutes.
3. **Budget consumed:** 0.20 × 15 min = **3.0 min** of equivalent downtime.
4. **Share of monthly budget:** 3.0 ÷ 43.2 = **≈ 6.9%** consumed.

Verdict: ~93% of budget remains → > 50% → ship normally. One incident didn't break the bank — but watch the trend.

Burn rate — how fast are you spending?

```
burn rate = (1 − actual SLI) ÷ (1 − SLO)
```

During the incident: (1 − 0.80) ÷ (1 − 0.999) = 0.20 ÷ 0.001 = **200×**. A burn rate of 1.0 means spending exactly as fast as allowed; 200× is a five-alarm spike that should page someone.

Time-based illustration for clarity; production teams typically compute request-based budgets and multi-window burn-rate alerts.

FRAMEWORK DEEP-DIVE · 3 OF 6

Toil reduction

3

Toil Reduction

THE EFFICIENCY FRAMEWORK

WHAT

Toil is manual, repetitive, automatable work that scales with load and has no lasting value.

WHY

Unchecked toil consumes the engineering time you need for reliability improvements.

Decision tree — is it toil, and what do I do?

Q1 · Is the task manual and repetitive?

- NO → Not toil. Leave it; revisit if it changes.
- YES ↓ Continue to Q2.

Q2 · Does the effort scale linearly with service growth?

- NO → Lower priority — document in a runbook.
- YES ↓ Continue to Q3.

Q3 · Is it automatable with reasonable effort?

- YES → Automate it — highest-leverage toil to remove.
- NO → Reduce frequency / re-architect to eliminate the trigger.

FRAMEWORK DEEP-DIVE · 4 OF 6

Observability & the golden signals

4

Observability & Golden Signals

THE VISIBILITY FRAMEWORK

WHAT

Understanding system state from its outputs — metrics, logs, and traces working together.

WHY

Metrics say latency spiked; logs say which service erred; traces say which call cascaded.

The four golden signals

Signal	What it tells you	Watch for
Latency	How long requests take	Rising p99 / tail latency
Traffic	Demand on the system	Spikes or sudden drops
Errors	Rate of failed requests	5xx climbing vs. budget
Saturation	How "full" the system is	CPU, memory, queue depth

The three pillars

Metrics

Aggregated time-series — the "what" and "how much," cheap to store and alert on.

Logs

Event detail — the "why," invaluable once metrics point you to a service.

Traces

Request paths across services — the "where," to find the cascade's origin.

Alert on the golden signals and SLO burn — not raw thresholds — to monitor impact, not noise.

FRAMEWORK DEEP-DIVE · 5 OF 6

Incident management

5

Incident Management

THE RESPONSE FRAMEWORK

WHAT

A structured lifecycle — detect, triage, mitigate, resolve — with clear roles and severity.

WHY

Calm, coordinated response cuts time-to-recover and protects the error budget.

Severity matrix

Sev	Impact	Response
SEV1	Major outage; many users / revenue hit	Page on-call + incident commander now
SEV2	Significant degradation; SLO at risk	Page on-call; engage owners
SEV3	Minor / partial; limited users	Handle in business hours
SEV4	Cosmetic / no user impact	Backlog & track

Escalation decision

Is an SLO actively burning or are users impacted?

YES →
Declare incident, assign a commander, open comms channel.

NO →
Log it, monitor, and address in normal workflow.

FRAMEWORK DEEP-DIVE · 6 OF 6

Blameless postmortems

6

Blameless Postmortems

THE LEARNING FRAMEWORK

WHAT

A structured review after incidents that focuses on systems and process — never individuals.

WHY

Psychological safety produces honest analysis, which produces durable fixes.

Blameless in practice

- ✓ **Ask "why did the system allow this?"** not "why did you do that?"
- ✓ **When code breaks prod,** ask why CI didn't catch it — not who pushed it.
- ✓ **Quantify with the budget** — "this outage consumed 30% of our quarterly error budget."
- ✓ **Rank incidents** by budget consumed to focus the next quarter's reliability work.

Feeds the loop

Action items become backlog work, closing the reliability improvement cycle.

Builds culture

Teams share lessons openly when no one is hunting for blame.

The postmortem template on the next page gives you a ready structure to run your first one.

READY-TO-USE TEMPLATE

The blameless postmortem template

Copy this structure for any incident that impacted an SLO. Keep the language factual and system-focused.

Incident Postmortem

BLAMELESS TEMPLATE

Title & date	Short incident name and time window
Severity & impact	Sev level; users, services, and revenue affected
Budget consumed	% of error budget spent by this incident
Summary	2–3 sentence plain-language overview
Timeline	Detection → escalation → mitigation → resolution
Root cause	Contributing factors (systems & process, not people)
Detection	How it was found; time-to-detect; gaps
Resolution	What restored service; time-to-recover
What went well	Effective responses worth repeating
Action items	Owners, due dates, preventive fixes
Lessons learned	Durable takeaways shared across teams

COMPARISON MATRIX

SRE vs DevOps — how they differ

Both bridge development and operations, but they're not the same thing. A useful shorthand: **"class SRE implements interface DevOps"** — SRE is a concrete implementation of DevOps principles.

Dimension	DevOps	SRE
Nature	A broad cultural movement	A concrete engineering discipline
Origin	Industry philosophy	Pioneered at Google
View of ops	"You build it, you run it"	Operations as a software problem
Reliability	Implied goal	Measured via SLOs & error budgets
Key metrics	Deploy frequency, lead time	SLO compliance, error budget, MTTR
Release control	Continuous delivery	Error-budget-governed releases
Failure response	Continuous improvement	Blameless postmortems + action items

Complementary, not rival

SRE gives DevOps the metrics and practices to make reliability concrete.

Shared destination

Both want fast, safe delivery — SRE just measures the "safe" part.

SRE MATURITY SELF-ASSESSMENT

How mature is your reliability practice?

LVL 1	LVL 2	LVL 3	LVL 4	LVL 5
Reactive	**Managed**	**Defined**	**Measured**	**Optimizing**
Firefighting; no SLOs; manual ops; blame culture.	Basic monitoring; some runbooks; ad-hoc incidents.	SLIs/SLOs set; blameless postmortems; some automation.	Error budgets govern releases; strong observability.	Automation-first; chaos testing; self-healing.

Score yourself — tick what's true today

- ✓ We have defined SLIs and SLOs for our critical services.
- ✓ Error budgets actually influence our release decisions.
- ✓ We run blameless postmortems for every significant incident.
- ✓ Toil is measured and actively reduced through automation.
- ✓ On-call is humane, with clear runbooks and escalation paths.
- ✓ Observability spans metrics, logs, and traces, correlated.

Boxes ticked	Indicative maturity
0–1	Level 1–2 — start with SLIs/SLOs and blameless reviews.
2–4	Level 3 — formalize error budgets and reduce toil.
5–6	Level 4–5 — push toward automation-first and chaos testing.

★ RELATED

Move your team up a maturity level

Whatever your score, the next level is a skill set away. Enroll in the GSDC SRE certification and lead the change with confidence.

Get Certified →

FROM FIELD GUIDE TO CERTIFICATION

Where you go deeper on each framework

Every framework in this guide maps to a module in the GSDC SRE certification, so you can turn reference knowledge into validated, hands-on skill.

Framework	GSDC module
1 · SLI / SLO / SLA	Module 2 — Service Level Objectives & Error Budgets
2 · Error Budgets	Module 2 — Service Level Objectives & Error Budgets
3 · Toil Reduction	Module 3 — Reducing Toil
4 · Observability & Golden Signals	Module 4 — Monitoring & SLIs
5 · Incident Management	Module 6 — Anti-Fragility & Learning from Failure
6 · Blameless Postmortems	Module 6 — Anti-Fragility & Learning from Failure
Foundations & culture	Modules 1, 5 & 7 — Principles, Automation, Org Impact

The full 7-module curriculum

1 · Principles & Practices | **2 · SLOs & Error Budgets** | **3 · Reducing Toil** | **4 · Monitoring & SLIs** | **5 · Tools & Automation** | **6 · Anti-Fragility** | **7 · Organizational Impact**

Each module pairs concepts with hands-on Learn-by-Doing activities — so you don't just read the frameworks, you apply them.

SALARY GUIDE

What framework fluency pays

Engineers who can apply these frameworks — not just name them — command strong compensation. Pay rises steadily with experience and specialized skills.

$120K–$170K
Typical US average SRE range

$210K+
Lead / Principal SRE upper range

+20–40%
Total comp uplift from bonus & equity

Career stage	Indicative annual range (US)
Entry-level SRE	$80,000 – $110,000
Mid-level SRE	$110,000 – $145,000
Senior SRE	$145,000 – $185,000+
Lead / Principal SRE	$170,000 – $210,000+

Indicative US ranges aggregated from public 2025–2026 sources (Glassdoor, Built In, Payscale). Pay varies by location, employer, experience and skills.

CAREER ROADMAP

From practitioner to reliability leader

Mastering these six frameworks is the through-line of an SRE career. Each rung deepens how broadly and strategically you apply them.

01 **Junior / Associate SRE**
Apply SLIs, dashboards, and runbooks under guidance.

02 **Site Reliability Engineer**
Own SLOs and error budgets; drive toil reduction and automation.

03 **Senior SRE**
Lead incidents and postmortems; shape reliability architecture.

04 **Staff / Principal SRE**
Set framework standards and resilience strategy across teams.

05 **SRE Manager / Head of Reliability**
Own reliability culture, policy, and maturity enterprise-wide.

Keep advancing

After Foundation, the **GSDC Certified SRE Practitioner** deepens applied expertise across these frameworks — the natural next credential.

AI GOVERNANCE & RELIABILITY ROLES

The frameworks meet AI

These six frameworks extend naturally to AI systems. As teams adopt ML for detection and autonomous remediation, the same SLO, observability, and postmortem discipline keeps AI reliable and governed — powering a new family of roles.

Emerging role	Framework it leans on
AIOps Engineer	Observability + incident management
ML Platform SRE	SLOs, error budgets, toil reduction
AI Reliability Engineer	SLIs/SLOs for latency, availability, drift
AI Governance & Compliance Lead	Postmortems + reliability standards
Observability Engineer (AI)	Golden signals for AI systems

Why the frameworks transfer

- ✓ **SLOs for AI** — define objectives for AI-service latency, availability, and accuracy drift.
- ✓ **Observability first** — golden signals make AI behavior measurable and auditable.
- ✓ **Guardrailed automation** — autonomous remediation needs error-budget and incident discipline.

★ ENROLLMENT OPEN

Apply these frameworks to AI-era systems

AIOps and AI-governance roles build on exactly these six frameworks. Enrollment is open now — join the GSDC SRE certification and get ahead.

Reserve My Spot →

HIRING TRENDS

Framework skills are in demand

Employers increasingly screen for the very frameworks in this guide — SLO thinking, observability, and incident discipline — across a widening set of industries.

Industries hiring SREs

Big Tech · Finance & Banking · Healthcare · E-commerce · Telecom · SaaS & Cloud · Gaming & Streaming · Startups

Trend	What it means for you
SLO-based alerting over thresholds	Teams want engineers who think in objectives.
Autonomous IT / AIOps standard	Framework fluency extends to AI-driven ops.
Open-source observability stacks	Hands-on golden-signals experience is prized.
75% enterprise SRE adoption by 2027	Sustained, rising demand for certified talent.

Credential = visibility

A recognized certification helps recruiters find you faster.

Talk frameworks fluently

Interviewers probe SLOs, error budgets, and postmortems directly.

Adoption and demand context from public industry roadmaps (2025–2026).

EXAM DETAILS & YOUR CREDENTIAL

How you get certified

Exam attribute	Detail
Question type	Multiple-choice (knowledge & scenario-based)
Number of questions	40
Duration	60 minutes
Passing score	Approximately 65–70%
Delivery	Online, remotely proctored, closed book
Attempts	2 attempts (1 free retake included)
Validity	Lifetime — no renewal fees

GSDC
GLOBAL SKILL DEVELOPMENT COUNCIL

ANSI Accredited Member

CERTIFICATE OF ACHIEVEMENT

This is proudly presented to

[Your Name]

for achieving the

Certified SRE Foundation (CSREF)

Certificate is illustrative; final design, wording and seal are issued by GSDC. Exam parameters reflect the standard SRE Foundation format — confirm current specifics at enrollment.

⏰ FINAL 48 HOURS

Last call — your offer expires soon

This is the final window. Enroll in the GSDC SRE certification within 48 hours and turn this field guide into a recognized credential.

Enroll Before It Ends →

YOUR NEXT STEP

From field guide to certified

You've got the six frameworks, the calculator, the matrices, the template, and the maturity check. The reliable next move is to turn reference knowledge into a recognized credential.

How it works

01 **Enroll**
Choose the GSDC SRE certification and start online.

02 **Learn & apply**
Work the 7 modules and hands-on activities on real scenarios.

03 **Certify**
Sit the online proctored exam — with a free retake if needed.

04 **Lead reliability**
Use the frameworks and your credential to raise your team's maturity.

Ready to master the SRE frameworks?

Start the GSDC SRE certification and join 75,000+ professionals who chose a globally recognized credential.

Get Certified →

Global Skill Development Council
www.gsdcouncil.org
accreditations@gsdcouncil.org

Lewes, Delaware, USA
Glattbrugg, Switzerland
Singapore