RELIABILITY UNLEASHED

STRATEGIC FRAMEWORK

From Chaos to Confidence

Site Reliability Engineering for Business Leaders

Executive Track | ~60 Minutes | Strategic Overview

THE BUSINESS CASE FOR RELIABILITY

$400K

Average cost per hour of downtime

79%

Customers leave after bad experience

5x

Cost to acquire new vs retain existing

Reliability isn't just an engineering problem.

It's a business differentiator and competitive advantage.

WHAT IS SITE RELIABILITY ENGINEERING?

> "SRE is what happens when you ask a software engineer to design an operations function."
>
> — Ben Treynor, VP Engineering, Google

TRADITIONAL OPS

- → Manual, reactive
- → Scaling = more people
- → Unmeasured reliability
- → Dev vs Ops tension

SRE APPROACH

- → Automated, proactive
- → Scaling = better software
- → Data-driven SLOs
- → Shared responsibility

THE RESEARCH: DORA METRICS

10+ years of research, 39,000+ participants

DEPLOY FREQ

Low Monthly

High Daily

Elite On-demand

LEAD TIME

Low Months

High Days

Elite <1 hour

FAILURE RATE

Low >30%

High 5-15%

Elite<5%

RECOVERY

Low Weeks

High <1 day

Elite <1 hour

Key insight: Elite performers ship faster AND more reliably

SLOS: MAKING RELIABILITY MEASURABLE

SLI

Service Level Indicator

What we measure

SLO

Service Level Objective

Target we aim for

ERROR BUDGET

Acceptable unreliability

Investment capacity

>50% budget

Ship features

25-50%

Prioritize reliability

<25%

Feature freeze

THE COST OF NINES

AVAILABILITY	DOWNTIME/YEAR	RELATIVE COST	USE CASE
99% (2 nines)	3.65 days	$	Internal tools
99.9% (3 nines)	8.76 hours	$$	Business apps
99.95%	4.38 hours	$$$	Customer-facing
99.99% (4 nines)	52.6 min	$$$$	Core platform
99.999% (5 nines)	5.26 min	$$$$$	Life-critical

Each additional nine roughly 10x the cost. Choose wisely.

LEARNING FROM INDUSTRY LEADERS

GOOGLE

Invented SRE, error budgets, 50% toil cap

NETFLIX

Chaos engineering, test in production

AMAZON

Cell-based architecture, blast radius

META

SEV culture, move fast safely

STRIPE

99.999% uptime, defensive design

SPOTIFY

Golden paths, platform engineering

HIGH-RELIABILITY ORGANIZATIONS

Lessons from Aviation, Nuclear, Healthcare, NASA

1	**Preoccupation with Failure**	— Never ignore small failures
2	**Reluctance to Simplify**	— Embrace complexity, don't oversimplify
3	**Sensitivity to Operations**	— Real-time situational awareness
4	**Commitment to Resilience**	— Detect, contain, recover quickly
5	**Deference to Expertise**	— Let experts decide, regardless of rank

THE OBSERVABILITY INVESTMENT

> *"If you can't monitor a service, you can't be reliable."*
>
> — *Google SRE Book*

METRICS

What's happening?

LOGS

What happened?

TRACES

Why did it happen?

ROI: Faster detection → Faster resolution → Less downtime

INCIDENT MANAGEMENT: THE ROI

WITHOUT PROCESS

- Chaos during incidents
- Unclear ownership
- Same issues recur
- Blame culture

MTTR: 4+ hours

WITH SRE PROCESS

- Structured response
- Clear roles (IC, Comms)
- Blameless postmortems
- Learning culture

MTTR: <1 hour

CULTURE: THE HIDDEN MULTIPLIER

Ron Westrum's Organizational Culture Types

PATHOLOGICAL

- Information is power
- Messengers are shot
- Failure leads to blame
- New ideas are crushed

BUREAUCRATIC

- Information is controlled
- Messengers are tolerated
- Failure leads to justice
- New ideas create problems

GENERATIVE

- Information is shared
- Messengers are trained
- Failure leads to inquiry
- New ideas are welcomed

DORA research: Generative culture predicts software delivery performance

CLOUD STRATEGY: RELIABILITY IMPLICATIONS

ON-PREMISES

- → Full control, full responsibility
- → Capital expenditure model
- → Limited geographic distribution
- → Predictable costs at scale

PUBLIC CLOUD

- → Shared responsibility model
- → Operational expenditure
- → Global distribution possible
- → Variable costs, auto-scaling

KEY DECISION **Match SLO to platform capability**

RISK FACTOR **Vendor lock-in vs portability**

AI/ML: NEW RELIABILITY CHALLENGES

TRADITIONAL SOFTWARE

→ Deterministic behavior

→ Clear failure modes

→ Static once deployed

→ Easy to test

AI/ML SYSTEMS

→ Non-deterministic output

→ Subtle degradation

→ Model drift over time

→ Harder to validate

New metrics needed: Model accuracy, prediction confidence, data quality, inference latency

AGENTIC OPERATIONS: THE FUTURE

1
DETECT
AI anomaly detection

2
DIAGNOSE
Automated analysis

3
REMEDIATE
Execute runbooks

4
LEARN
Improve models

70%
Auto-resolution

<15 min
MTTR target

24/7
Autonomous

PLATFORM ENGINEERING: DEVELOPER PRODUCTIVITY

"A golden path is a paved road to a well-architected production deployment."

— Spotify Engineering

WITHOUT PLATFORM

Each team reinvents deployment, monitoring, security. Weeks to production.

WITH PLATFORM

Self-service, paved paths, built-in best practices. Hours to production.

Make the right thing the easy thing

INVESTMENT PRIORITIES BY MATURITY

1	2	3	4	5
FOUNDATION	**MEASUREMENT**	**AUTOMATION**	**PLATFORM**	**INTELLIGENCE**
Monitoring, on-call	SLOs, error budgets	CI/CD, remediation	Golden paths	AI/ML, agentic

Each phase builds on the previous. Don't skip steps.

MEASURING SRE ROI

AVAILABILITY GAINS

- Reduced downtime cost
- Fewer customer impacts
- Less revenue at risk

VELOCITY GAINS

- Faster time to market
- More deployments/day
- Shorter lead times

EFFICIENCY GAINS

- Less manual toil
- Reduced on-call burden
- Better resource use

PEOPLE GAINS

- Lower attrition
- Higher engagement
- Better recruitment

STRATEGIC ANTI-PATTERNS TO AVOID

✕ RELIABILITY AS AFTERTHOUGHT

"We'll make it reliable after we ship" - technical debt compounds

✕ TOOL-FIRST THINKING

"Let's buy Kubernetes" - tools don't solve culture/process problems

✕ OVER-ENGINEERING SLOS

"We need 99.999%" - costs escalate, value doesn't

✕ BLAME CULTURE

"Find who caused this" - kills learning, hides problems

KEY TAKEAWAYS FOR LEADERS

1. **Reliability = Business Feature** — impacts revenue & retention
2. **Measure What Matters** — SLOs & DORA enable decisions
3. **Culture is the Multiplier** — predicts delivery performance
4. **Invest Progressively** — foundation → automation → intelligence
5. **Future is Agentic** — autonomous ops reduce cost

RELIABILITY UNLEASHED

QUESTIONS?

34 Detailed One-Pagers Available

Technical deep-dives covering all topics discussed today

Essential Reading

Google SRE Book

Accelerate (DORA)

Next Steps

Assess current maturity

Define SLOs

Build roadmap