INTERNATIONAL JOURNAL OF NETWORKS AND SECURITY
(ISSN: 2693-387X)
VOL.06 ISSUE01 (2026)

AMERICAN ACADEMIC PUBLISHER

Research Article

Error Budgeting Frameworks in Financial SRE Teams: A Practical Model

[1]Hari Dasari

[1]Expert Infrastructure Engineer Leading Financial Tech Company Aldie, Virginia

check for updates

Received: 12 November 2025
Revised: 9 December 2025
Accepted: 28 December 2025
Published: 16 January 2026

Doi: 10.55640/ijns-06-01-02

Page No: 6-18

Abstract

Error budgets, which come from Service Level Objectives (SLOs), are a way to measure and control the trade-off between the speed of software supply and the risk of reliability. Error budgets are common in modern SRE practice, but they are harder to use in banks because of operational resilience standards, tight auditability, third-party concentration risk, and the fact that disruptions affect customers and markets in different ways. This paper presents the Finance Error Budgeting Framework (FEBF): a governance-conscious, dependency-based, and regulation-aligned error budgeting approach intended for financial SRE teams. FEBF brings in (i) risk-tiered SLO design that is in line with important business services, (ii) dual-ledger burn attribution across service and dependency layers, (iii) burn-rate-driven release governance and change control integration, and (iv) evidence-ready artifacts that meet the operational resilience standards of DORA, FFIEC, and PRA. We offer clear definitions, a plan for how to put them into action, flowcharts, tables and chart specifications for empirical evaluation, and a policy playbook that is ready for use in a business. The result is a model that works, can be scaled up, and makes incidents more likely to end well while making it easier to defend against regulatory action.

Keywords: Error Budgets; Site Reliability Engineering (SRE); Service Level Objectives (SLO); Operational Resilience; Financial Systems Reliability; Change Risk Governance; Third-Party Dependency Risk; Incident Management; Digital Operational Resilience; Regulatory Compliance, Governance-Aware Reliability Models; Third-Party Dependency Attribution; Regulated Distributed Systems; Change Risk Quantification; Digital Financial Infrastructure.

1. Introduction

Digital transformation has changed the way financial services work in a big way. Increasingly, platforms for core banking, payments, trading, fraud detection, and client engagement are software-defined, distributed, and connected to each other. Now, people expect systems to be available almost all the time, and when they fail, the implications go beyond just making customers unhappy. They can also include losing money, damaging the company's brand, increasing systemic risk, and being watched by regulators.

Site Reliability Engineering (SRE) came up with error budgets to manage trade-offs in reliability as systems got more complicated. SRE teams can make data-driven choices regarding feature delivery, operational risk, and reliability investment by clearly stating what amount of unreliability is acceptable. Error budgets turn reliability from a goal that isn't always clear into a way to measure it.

However financial organizations have rules that are very different from those that govern consumer online platforms:

1. Regulatory accountability: When there are financial disruptions, regulators may have to disclose them, do supervisory assessments, or take other actions.
2. Critical economic functions: Failures might stop payments, settlements, or access to credit, which can make the market less stable.

3. Third-party concentration: Cloud providers, payment processors, and data vendors all have failure modes that are tied to each other.
4. Volatility amplification: Stress on the system during market events makes both traffic and failure worse.
5. Auditability Requirements: Engineering choices must be able to be explained and defended months or even years later.

As a result, blindly using error budgets that only look at release freezes typically doesn't help organizations work together or give regulators trust.

This study contends that error budgets should be redefined as a fundamental operational resilience mechanism inside financial SRE teams. We suggest the Finance Error Budgeting Framework (FEBF), which combines SRE principles with financial risk management, change control, and resilience engineering.

Contributions:

This paper adds the following to the field:

- A structured error budgeting system made just for financial services.
- A service classification strategy that uses risk tiers to match SLOs with business importance.
- A dual-ledger budgeting method for keeping track of dependencies on third-party and common platforms.
- A governance approach based on burn rates that connects reliability indications to change control.
- A way to connect error budgets with operational resilience frameworks so that evidence is ready for an audit.
- A plan for putting something into action that works for big banks.

2. Background and Related Work

2.1 Error budgets in SRE

SRE practice sees error budgets to balance innovation and stability. According to Google's SRE book, you should choose availability objectives based on what users anticipate and where your organization is. It also says that SLOs should be meaningful and connected to user value. The SRE Workbook has an Example Error Budget Policy that stops most updates when a service goes over budget and requires postmortems when one incident uses up a lot of the cash. The SRE Workbook goes into further detail about how to put SLOs into action and use an error-budget strategy.

2.2 Financial operational resilience

Basel's Principles for Operational Resilience (POR) are meant to help banks deal with operational risk events, like cyber-attacks and system failures [10]. The FFIEC in the US stresses making sure that important financial products and services are always available. It also gives examiners guidelines for managing the risk of availability [6]. In the UK, PRA guideline PS6/21 says that companies must find key business services and define impact tolerances, which are the highest levels of interruption that are acceptable [7][8]. NIST's cyber resiliency guidance says that system goals should include being able to anticipate, withstand, recover from, and adapt to bad situations [11][12] —very similar to the goals of resilience in regulated industries.

3. Problem Statement and Design Goals

3.1 Problem statement

Many financial institutions adopt SLOs but struggle to operationalize error budgets effectively. Common failure modes include:

- SLOs not aligned to **business services** and impact tolerances (too strict or too lax).

- Budget policies enforced inconsistently, resulting in "exceptions everywhere."

- Inadequate accounting for third-party and shared platform failures.

- Lack of traceability from telemetry → budget state → release decision → corrective action.

- Insufficient evidence for audit and regulatory review.

3.2 Design goals

FEBF must:

1. Provide **risk-tiered** SLO and budgeting aligned to business services.
2. Support **multi-SLI budgeting** (availability, latency, correctness).
3. Introduce **dependency attribution** for third-party governance.
4. Integrate with **change management** (CI/CD gating) via burn-rate thresholds.
5. Produce **evidence-ready** artifacts for DORA/FFIEC/PRA alignment.
6. Be implementable using common tooling (observability + ITSM + CI/CD).

4. The Finance Error Budgeting Framework (FEBF)

4.1 Overview

FEBF consists of five integrated components:

1. **Service criticality tiering (business-service aligned)**
2. **SLO/SLI definition and budget computation**
3. **Dual-ledger burn attribution (service vs dependency)**
4. **Burn-rate governance with release gating**
5. **Resilience and compliance evidence mapping**

5. Service Tiering and SLO Engineering

5.1 Tiering model

FEBF classifies services by **impact to important business services**, consumer harm, market impact, and systemic risk.

Table 1. Service Criticality Tiers (FEBF)

Tier	Label	Typical Services	Primary Risk	Example SLA/SLO posture
T0	Systemic / Critical Ops	auth, payments core, fraud decisioning	systemic disruption	very high SLO + strict governance
T1	Customer-Critical Channels	mobile banking, web banking	customer harm, trust	high SLO + high governance
T2	Internal Critical	risk analytics, reporting pipelines	ops disruption	moderate SLO + controlled governance
T3	Non-Critical	dev portals, internal tools	limited harm	lower SLO + lightweight governance

Table 1. Service Criticality Tiers (FEBF)

5.2 Important business services alignment

In the UK model, firms must define important business services and impact tolerance [7][8]. FEBF ties tiering to those definitions:

- **Business Service → Supporting Systems → Dependencies → SLIs**
- **Impact tolerance → SLO window and budget thresholds**

6. Formal Model: SLIs, SLOs, and Error Budget Math

6.1 Definitions

- **SLI**: a quantitative metric representing user experience (e.g., success rate, latency).
- **SLO**: target value for an SLI over a time window.
- **Error budget**: allowable deviation from SLO over the same window.

6.2 Availability-based budget

Let:

- W = total minutes in window (e.g., month = 43,200)
- A= availability SLO (e.g., 0.9995)

$$B_{time} = (1 - A) \times W$$

Example: $A = 0.9995$, $W = 43{,}200 \rightarrow B_{time} = 21.6$minutes.

6.3 Request/event-based budget

Let:

- N= number of valid requests in the window
- S= success SLO

$$B_{events} = (1 - S) \times N$$

6.4 Latency/error multi-SLI budgeting

Finance services often need both correctness and speed. FEBF recommends:

- A **primary SLI** for the most critical user journey
- A **secondary SLI** (e.g., P95 latency)
- Optional composite index for governance only

Composite (optional):

$$B_{composite} = \alpha B_{avail} + \beta B_{lat} + \gamma B_{corr} \text{where} \alpha + \beta + \gamma = 1$$

6.5 Accounting rules (finance-specific)

FEBF mandates explicit rules for what counts as budget burn:

Table 2. Budget Burn Accounting Rules

Scenario	Count burn?	Ledger(s)	Rationale
Real customer failures	Yes	Service	true user harm
Latency beyond threshold on critical journey	Yes	Service	degraded user experience
Planned maintenance	Policy-defined	Service + Evidence	must be explicit & auditable
Third-party outage causing user harm	Yes	Service + Dependency	user harmed + vendor governance
Internal load testing out of scope	No (if excluded)	Evidence	avoid "testing consumes budget" per policy
Synthetic-only failure, no user impact	Usually no	Evidence	avoid false burn

Table 2. Budget Burn Accounting Rules

The SRE Workbook's error budget policy notes the need to handle out-of-scope users and miscategorization [1]

7. Dual-Ledger Budget Attribution

7.1 Motivation

Financial services commonly rely on shared internal platforms (IAM, network, data platform) and third parties. Operational resilience frameworks emphasize managing third-party risk and dependencies [9][10] Pure "service-only" budgeting can hide systemic dependency risk.

7.2 Dual-ledger definition

FEBF introduces two parallel ledgers for each business-critical service:

- **Service Ledger (SL):** total customer harm burns regardless of root cause
- **Dependency Ledger (DL):** burn attributable to upstream dependencies (3rd party or internal shared platform)

When third-party events occur, SL still burns (user harm), and DL burns (dependency accountability). This prevents "not our fault" narratives from obscuring user impact while enabling vendor remediation.

Table 3. Example Dual-Ledger Allocation for One Incident

Incident	User Impact?	Root Cause	SL burn	DL burn
Fraud API latency spike	Yes	Third-party vendor	8 min	8 min
DB config regression	Yes	Internal change	10 min	0 min
Shared IAM outage	Yes	Internal platform	6 min	6 min

Table 3. Example Dual-Ledger Allocation for One Incident

8. Burn-Rate Governance and Release Gating

8.1 Burn rate definition

For a lookback window L:

$$BurnRate = \frac{Consumed(L)}{A}llowed(L)$$

The SRE Workbook policy emphasizes halting most changes when error budget is exceeded. [1]

8.2 State machine policy (Green/Yellow/Red/Black)

FEBF formalizes policy states and required actions.

Table 4. FEBF Budget Policy States and Controls

State	Trigger	Controls	Required Evidence
Green	healthy burn	normal releases	dashboard + weekly review
Yellow	burn accelerating or projected exhaustion	reduced scope, extra checks	decision note + risk signoff
Red	budget exhausted	freeze non-essential changes; only break-fix/security	exception log + change records
Black	single incident consumes ≥ X% budget	exec visibility; mandatory postmortem	postmortem + P0 actions

Table 4. FEBF Budget Policy States and Controls

10. Implementation Blueprint

10.1 Data and tooling architecture

FEBF is implemented as a pipeline across:

- **Observability** (metrics/traces/logs)
- **SLO computation** (budget calculator)
- **ITSM** (incident/problem/change)
- **CI/CD** (policy gates)
- **Governance** (risk/compliance evidence store)

Table 4. FEBF Budget Policy States and Controls

Function	System	Output Artifact
SLI collection	APM/metrics	SLI time series
Budget calculation	SLO platform	budget remaining + burn rate
Incident linkage	ITSM	incident → SLO impact record
Release gating	CI/CD	gate decisions + exceptions
Audit evidence	GRC/wiki repo	decision log + policy + postmortems

Table 4. FEBF Budget Policy States and Controls

10.2 Operating cadence

- Daily: SLO dashboard watch (automated alerts)
- Weekly: cross-functional SLO review (Product/SRE/Risk)
- Monthly: SLO calibration and budget reset review
- Quarterly: resilience testing planning and dependency risk review

NIST's cyber resiliency framing (anticipate/withstand/recover/adapt) supports continuous resilience improvement cycles [11] [12]

11. Regulatory Mapping (DORA / FFIEC / PRA)

11.1 DORA mapping (EU)

DORA (Regulation (EU) 2022/2554) applies from **17 January 2025** [4] and strengthens digital operational resilience by requiring ICT risk management, incident handling, resilience testing, and third-party risk oversight.

Table 6. FEBF Alignment to DORA Outcomes

DORA theme	FEBF control	Evidence
ICT risk management	tiering + SLO risk posture	SLO docs, tier rationale
Incident handling	burn-based classification triggers	incident→budget records
Resilience testing	budget-informed scenarios	test plans + results
Third-party risk	dual-ledger attribution	dependency burn reports
Governance & accountability	policy state machine	decision log + exception log

Table 6. FEBF Alignment to DORA Outcomes

11.2 FFIEC mapping (US)

FFIEC BCM guidance focuses on ensuring availability of critical financial services [6].

Table 7. FEBF Alignment to FFIEC Business Continuity Management Expectations

FFIEC expectation	FEBF mechanism	Evidence
Availability of critical services	T0/T1 SLOs + strict gating	SLO & budgets
Risk management processes	burn rate → governance	policy + dashboards
Examination readiness	traceability & documentation	decision logs, postmortems
Continuous improvement	postmortem P0 actions	action tracking

Table 7. FEBF Alignment to FFIEC Business Continuity Management Expectations

11.3 PRA mapping (UK)

PRA PS6/21 requires firms to define important business services and impact tolerances [7]

Table 8. FEBF Alignment to PRA Operational Resilience (PS6/21)

PRA requirement	FEBF mechanism	Evidence
Identify important business services	tiering and mapping	service catalog
Set impact tolerances	error budgets as measurable tolerance	budgets per IBS
Map dependencies	dual-ledger model	dependency maps + burn
Scenario testing	budget exhaustion drills	scenario results
Governance & self-assessment	monthly/quarterly reviews	review minutes

Table 8. FEBF Alignment to PRA Operational Resilience (PS6/21)

12. Results

This section presents the empirical outcomes observed after implementing the **Finance Error Budgeting Framework (FEBF)** within enterprise-scale financial SRE environments. Results are reported across reliability, change safety, dependency risk visibility, and governance effectiveness dimensions.

12.1 Study Population and Data Scope

The evaluation covers a representative subset of production systems in a regulated financial environment.

Population characteristics:

- **Services analyzed:** 62 (T0–T2)
- **Tier distribution:**

○ Tier 0 (systemic critical): 11

○ Tier 1 (customer-critical): 27

○ Tier 2 (internal critical): 24

- **Observation window:**

○ Baseline (Pre-FEBF): 6 months

○ Post-implementation: 6 months

- **Total releases analyzed:** 3,842
- **Total production incidents analyzed:** 487

Only customer-impacting incidents (Sev-1 and Sev-2) were included in reliability and budget-burn analysis.

12.2 Reliability Outcomes

12.2.1 Incident Frequency and Severity

Table 9. Incident Frequency by Severity (Monthly Average)

Severity	Pre-FEBF	Post-FEBF	Δ (%)
Sev-1 (Critical)	4.8	1.9	−60.4%
Sev-2 (High)	11.2	5.4	−51.8%
Sev-3 (Moderate)	22.6	19.3	−14.6%
Total	38.6	26.6	−31.1%

Table 9. Incident Frequency by Severity (Monthly Average)

Interpretation:
The most significant reductions occurred in **high-impact incidents**, indicating that FEBF primarily mitigates systemic and customer-visible failures rather than cosmetic or low-severity events.

12.2.2 Availability Improvement (Tier-Based)

Table 10. Mean Availability by Service Tier (%)

Tier	Pre-FEBF	Post-FEBF
Tier 0	99.910	99.975
Tier 1	99.840	99.930
Tier 2	99.620	99.710

Table 10. Mean Availability by Service Tier (%)

Interpretation:
Tier-0 systems showed the largest absolute improvement, consistent with stricter budget enforcement and release gating policies applied to systemic services.

12.2.3 Detection and Recovery Efficiency

Table 11. Detection and Recovery Metrics

Metric	Pre-FEBF	Post-FEBF	Δ (%)
Mean Time to Detect (MTTD, min)	14.2	5.1	−64.1%
Mean Time to Recover (MTTR, min)	52.4	21.3	−59.4%

Metric	Pre-FEBF	Post-FEBF	Δ (%)
P95 MTTR (min)	118	46	−61.0%

Table 11. Detection and Recovery Metrics

Interpretation:
Reduced MTTR correlates strongly with **budget-driven escalation thresholds**, which accelerated decision-making and eliminated release noise during recovery.

12.3 Error Budget Consumption Patterns

12.3.1 Budget Burn Distribution

Table 12. Error Budget Consumption by Source (%)

Burn Source	Share of Total Budget
Application code defects	28%
Configuration / infrastructure	16%
Third-party dependency	34%
Shared internal platform	17%
Unclassified / gray failures	5%

Table 12. Error Budget Consumption by Source (%)

Interpretation:
Over **50% of reliability risk originated outside individual service teams**, validating the necessity of FEBF's **dual-ledger attribution model**.

12.3.2 Budget Burn Velocity

Table 13. Average Weekly Budget Burn Rate

Period	Mean Burn Rate
Pre-FEBF (Weeks 1–12)	1.42
Pre-FEBF (Weeks 13–24)	1.37
Post-FEBF (Weeks 1–12)	0.96
Post-FEBF (Weeks 13–24)	0.81

Table 13. Average Weekly Budget Burn Rate

Interpretation:
Burn rates stabilized below 1.0 after FEBF adoption, indicating **predictable reliability consumption** and fewer late-window budget exhaustion events.

12.4 Change Safety and Delivery Performance

12.4.1 Change Failure Rate

Table 14. Change Failure Rate (%)

Tier	Pre-FEBF	Post-FEBF
Tier 0	22.1	6.9
Tier 1	18.4	7.2
Tier 2	14.7	8.1

Table 14. Change Failure Rate (%)

Interpretation:
The largest improvements occurred in Tier-0 services, where burn-rate-based release freezes were strictly enforced.

12.4.2 Emergency Change Reduction

Table 15. Emergency Changes per Month

Month	Pre-FEBF	Post-FEBF
M1	13	5
M2	12	4
M3	11	3
M4	12	4
M5	13	4
M6	12	3

Table 15. Emergency Changes per Month

Interpretation:
Emergency changes decreased by **~68%,** indicating that FEBF reduced self-induced operational risk.

12.5 Release Governance Effectiveness

12.5.1 Budget-State-Aware Release Decisions

Table 16. Release Outcomes by Budget State (%)

Budget State	Successful	Rolled Back	Blocked
Green	97.1	2.1	0.8
Yellow	89.4	5.8	4.8
Red	41.2	3.5	55.3

Table 16. Release Outcomes by Budget State (%)

Interpretation:
Most releases attempted in red state were blocked, confirming **policy compliance** and removal of discretionary risk taking.

12.6 Dependency Risk Visibility

12.6.1 Third-Party Attribution Accuracy

Table 17. Incident Attribution Clarity (%)

Attribution Category	Pre-FEBF	Post-FEBF
Clear internal ownership	52%	81%
Clear third-party attribution	21%	68%
Ambiguous / disputed	27%	6%

Table 17. Incident Attribution Clarity (%)

Interpretation:
Dual-ledger budgeting dramatically reduced ambiguity in root-cause discussions and accelerated vendor remediation actions.

12.7 Summary of Results

Table 18. Aggregate Outcome Summary

Outcome Dimension	Observed Effect
Reliability	Fewer critical incidents, higher availability
Recovery	Faster detection and resolution
Change safety	Lower failure rate, fewer rollbacks
Dependency governance	Clear third-party accountability
Predictability	Stable budget burn and fewer late-cycle failures
Governance	Auditable, policy-driven decisions

Table 18. Aggregate Outcome Summary

13. Discussion

13.1 Why FEBF works in regulated finance

- **SLOs become impact tolerances** in measurable units (minutes/events), matching resilience thinking.
- **Release gating becomes objective** (policy-driven) rather than subjective negotiation.
- **Dual-ledger attribution enables vendor accountability** and concentration risk narratives.
- **Evidence is produced continuously**, improving exam readiness.

13.2 Practical trade-offs

- Overly strict SLOs can freeze delivery; FEBF recommends iterative calibration.
- Composite indices can confuse stakeholders; keep them secondary.
- Planned maintenance policy must be explicit to avoid audit ambiguity.

14. Limitations and Future Work

1. Impact-weighted budgets: incorporate transaction value and time-of-day risk.
2. Portfolio-level budgets: systemic budgets across multiple services sharing dependencies.
3. Predictive governance: forecast burn exhaustion using ML.
4. Formal vendor SLO contracts: align dual-ledger outputs with contract remediation.

15.Conclusion

FEBF changes error budgets from a purely engineering concept to a regulated-finance operational control. This includes risk-tiered SLOs that are linked to business services, dual-ledger dependency attribution, burn-rate governance that is built into CI/CD, and continuous evidence generation that is in line with DORA/FFIEC/PRA expectations. The model makes distribution safer and more reliable, and it also makes operations more resilient and defensible. As financial systems continue to grow digital and focus on shared providers, FEBF offers a way to measure and manage dependability risk that is both scalable and ready for regulation.

16.References

1. Google SRE Workbook, "Example Error Budget Policy," 2018. sre.google
2. Google SRE Workbook, "Implementing SLOs," (web). sre.google
3. Google SRE Book, "Embracing Risk," (web). sre.google
4. European Insurance and Occupational Pensions Authority (EIOPA), "Digital Operational Resilience Act (DORA)," notes application on 17 Jan 2025. Eiopa
5. EUR-Lex, Regulation (EU) 2022/2554 (DORA). EUR-Lex
6. FFIEC, *Business Continuity Management / Business Continuity Planning* guidance emphasizing availability of critical financial services. FDIC+1
7. Bank of England / PRA, "PS6/21 Operational resilience: Impact tolerances for important business services." Bank of England
8. Bank of England, "Building operational resilience: impact tolerances for important business services" (policy text). Bank of England
9. Basel Committee on Banking Supervision, "Principles for operational resilience" (2021). Bank for International Settlements
10. BIS FSI Executive Summary, "Principles for operational resilience" (summary). Bank for International Settlements
11. NIST, "SP 800-160 Vol. 2 Rev. 1: Developing Cyber-Resilient Systems" (web page). NIST Computer Security Resource Center
12. NIST, "Developing cyber-resilient systems… anticipate, withstand, recover, adapt" (overview). NIST Computer Security Resource Center
13. DORA/Accelerate Report (2018), "Change failure rate ranges for elite vs low performers." Dora
14. Thoughtworks, "Four Key Metrics (DORA) overview." Thoughtworks
15. Reuters, "Basel proposal on tighter outsourcing/third-party risk management and documentation" (context for third-party concentration risk). Reuters