Circuit Breaker and Retries

(Production Field Manuals)

Exponential Backoff, Jitter, Kubernetes Operations, and FinOps Impact

Luca Sepe

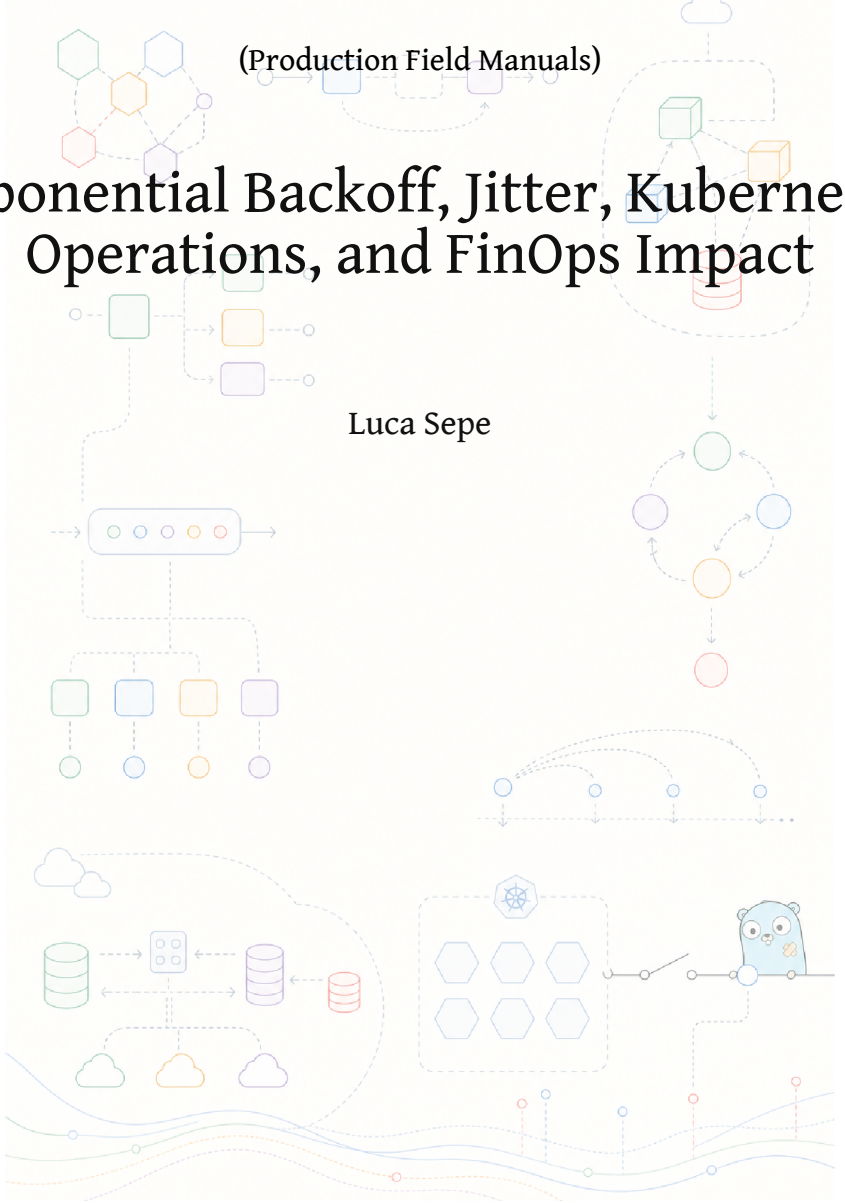

Index

How to Use This Book

The Series

Production Field Manuals is a series of practical guides for production engineering patterns.

Each book explains one topic across four levels: business impact, architecture, implementation and operations.

The goal is to close the gap between people who approve technical direction and people who operate the code under real failure.

This is not a catalog of patterns. It is a practical guide for making one pattern work in production.

This Book

This volume covers **Circuit Breaker and Retries with Exponential Backoff**.

The pattern matters because distributed systems fail partially. A dependency can become slow, overloaded, rate-limited, or unreachable while the rest of the platform is still healthy.

The wrong retry policy can turn a small downstream issue into a cluster-wide incident. The right combination of timeouts, bounded retries, jitter, and circuit breaking limits blast radius and gives the dependency room to recover.

What You Will Get

By the end, you should be able to:

- explain the pattern to leadership in business terms
- design a bounded retry policy from a latency budget
- decide when a circuit breaker should open
- avoid retrying unsafe operations
- expose the right metrics
- operate the pattern in Kubernetes
- run a hands-on lab
- defend the design in a Staff-level review
- use the companion Go implementation as a production starting point

Companion Source Code

The companion source code is public:

```
github.com/lucasepe/circuit-breaker-and-retries
```

The book keeps code snippets intentionally short. The repository contains the complete implementation, tests, benchmarks, and examples.

The companion implementation is written in Go so the book can stay concrete about concurrency, timeouts, contexts, tests, and production ergonomics. After you understand the pattern and its trade-offs, you can implement the same ideas in your preferred programming language.

How to Read It

If you are a CTO or engineering leader, start with: Chapter 1, Chapter 7 and Chapter 9.

If you are an architect, focus on: Chapter 2, Chapter 7 and Chapter 10.

If you are implementing the pattern, read: Chapter 3, Chapter 6 and the companion repository.

If you are operating the system, read: Chapter 4, Chapter 6, Chapter 8 and Chapter 10.

What Not to Do

Do not copy configuration values blindly.

Every timeout, retry count, breaker threshold, and HPA behavior must fit the service's latency budget, dependency SLO, traffic level, and failure model.

Use the defaults in this book as starting points, then validate them with hands-on drills.

Production Principle

Retries are useful only when they are bounded.

Circuit breakers are useful only when they are observable.

Neither pattern is a substitute for clear ownership, capacity planning, idempotency, and operational discipline.

Architectural Deep-Dive

State Management and Flow

The circuit breaker is a finite state machine around a remote call.

It has three standard states:

State	Behavior	Purpose
`Closed`	Requests are allowed. Failures are counted.	Normal operation.
`Open`	Requests fail immediately without touching the dependency.	Stop failure amplification.
`HalfOpen`	A small number of probe requests are allowed.	Test recovery without flooding the dependency.

The usual flow:

1. The caller starts in `Closed`.
2. Each failed call increments a failure counter or updates a rolling error-rate window.
3. When the threshold is exceeded, the breaker moves to `Open`.
4. In `Open`, calls fail fast until the cooldown expires.
5. After cooldown, the breaker moves to `HalfOpen`.
6. A small number of probes are allowed.
7. If probes succeed, the breaker returns to `Closed`.
8. If probes fail, the breaker returns to `Open`.

Retries sit inside the call path, but they must respect the circuit breaker and the request deadline.

A safe order is:

```
incoming request
  -> derive context deadline
  -> ask circuit breaker for permission
  -> execute attempt with per-attempt timeout
  -> classify error
  -> retry only if retryable, idempotent, and deadline allows
  -> record success or failure into breaker
```

Under stress, the important behavior is not that every request eventually succeeds. The important behavior is that the caller stops making the outage worse.

Failure Classification

Do not treat every error equally.

Retryable failures usually include:

- connection reset
- temporary DNS failure
- HTTP `408`, `429`, `500`, `502`, `503`, `504`
- gRPC `Unavailable`, `DeadlineExceeded`, `ResourceExhausted`

Non-retryable failures usually include:

- invalid input
- authentication failure
- authorization failure
- HTTP `400`, `401`, `403`, `404`, `409` unless domain logic says otherwise
- schema validation errors

The breaker should count dependency health failures, not caller mistakes.

System Impact

Network Latency

Retries increase tail latency. That is the trade-off.

The goal is to improve success rate for transient failures without breaking the caller's latency budget.

If the user-facing endpoint has a 500 ms p95 SLO, a retry policy like this is already expensive:

```
attempt timeout: 150 ms
base delay: 50 ms
max attempts: 3
worst case: 150 + 50 + 150 + 100 + 150 = 600 ms before overhead
```

That policy cannot fit inside a 500 ms request unless the first attempt usually fails fast. Production retry policy must start from the end-to-end deadline, not from a library default.

Consistency

Retries interact directly with data consistency.

For reads:

- retries are usually safe
- stale fallback responses may be acceptable
- eventual consistency is often a valid trade-off

For writes:

- retries require idempotency keys or natural idempotency
- duplicate side effects must be prevented server-side
- the caller must handle ”unknown outcome” after timeout

Strong consistency paths need stricter limits. A payment capture, inventory reservation, or account mutation should not rely on blind client retries. Use an idempotency key, durable command record, or queue-backed workflow.

Throughput

Circuit breakers improve cluster throughput during failure by reducing useless work.

Without this pattern:

- workers block on slow I/O
- connection pools fill
- goroutines accumulate
- CPU rises from retries, logging, TLS handshakes, and serialization
- HPA scales out the caller tier
- downstream receives even more pressure

With this pattern:

- open circuit calls are rejected locally
- retry volume is bounded

- jitter spreads load over time
- downstream can recover
- caller capacity remains available for healthy paths

The system-level win is not free. If thresholds are too aggressive, the breaker can reduce availability by opening during small error bursts. If thresholds are too loose, it opens too late and does not protect the cluster.

Architectural Defaults

Good starting defaults:

Parameter	Starting Point	Notes
Per-attempt timeout	50 to 300 ms for internal calls	Must fit inside the caller deadline.
Max attempts	2 to 3	More attempts usually hurt p99 latency.
Base backoff	25 to 100 ms	Depends on service SLO.
Max backoff	500 ms to 2 s	Cap prevents unbounded waits.
Jitter	Full jitter or equal jitter	Required at scale.
Open threshold	5 to 20 consecutive failures, or error-rate window	Error-rate windows are better at high QPS.
Open cooldown	5 to 30 s	Long enough for dependency recovery.
Half-open probes	1 to 10 concurrent probes	Keep small to avoid recovery floods.

For high-QPS services, prefer rolling-window error rate over consecutive failures. Consecutive failure thresholds are simple but noisy.

FinOps and ROI Cost Model

Why Cost Modeling Belongs Here

Circuit breakers and bounded retries are reliability patterns, but they also control cloud spend.

During dependency incidents, waste usually comes from:

- caller pods doing useless retries
- HPA scaling callers because retries increase CPU
- connection pools and TLS handshakes increasing
- logs and metrics volume exploding
- downstream systems being overprovisioned for failure peaks
- engineers spending hours across multiple teams diagnosing secondary symptoms

The savings are not only steady-state compute. The larger value is avoiding cost spikes and revenue-impacting cascading outages.

Basic Variables

Use these variables for a quick model.

Variable	Meaning
`R`	incoming requests per second to caller
`D`	percentage of requests calling the dependency
`A`	max attempts per dependency call
`P`	number of caller pods
`C`	average CPU cost per pod per hour
`H`	incident duration in hours
`S`	HPA scale multiplier during incident
`L`	log and metrics cost per million events
`E`	extra error events per failed attempt

Without control, worst-case attempted dependency calls can approach:

```
R * D * A * S
```

With circuit breaking, once the breaker is open, attempted dependency calls move closer to:

```
R * D * probe_rate
```

The exact numbers depend on traffic, but the direction matters: open circuits convert remote waste into local rejections.

Retry Amplification Example

Assume:

Input	Value
Incoming traffic	10,000 RPS
Dependency call ratio	30 percent
Max attempts	3
HPA incident multiplier	2x
Dependency outage duration	1 hour

Without breaker containment:

```
10,000 * 0.30 * 3 * 2 = 18,000 attempted dependency calls/sec
```

Baseline healthy dependency volume:

```
10,000 * 0.30 = 3,000 calls/sec
```

Amplification:

```
18,000 / 3,000 = 6x
```

That means the dependency receives six times its normal request volume while unhealthy.

With circuit breaking, the dependency may receive only a small probe rate during the open period. If 100 caller pods each allow one half-open probe every 10 seconds:

```
100 / 10 = 10 probe calls/sec
```

That is the difference between recovery and collapse.

Compute Cost Example

Assume:

Input	Value
Normal caller pods	80
Incident caller pods	200
Cost per pod-hour	$0.12
Incident duration	3 hours

Extra compute cost:

```
(200 - 80) * 0.12 * 3 = $43.20
```

This number may look small. It is not the main cost.

The larger costs are:

- revenue loss during degraded checkout or login
- engineer time during incident response
- database or cache overprovisioning
- observability ingestion spike
- customer support load
- SLO burn and error-budget depletion

Compute waste is the visible cost. Reliability damage is the expensive cost.

Observability Cost Example

Assume:

Input	Value
Extra failed attempts	50 million
Log lines per failed attempt	2
Cost per million log events	$0.50

Extra log cost:

```
50 * 2 * 0.50 = $50
```

Again, the dollar amount may not be huge for one incident, but the signal-to-noise damage is severe. During incidents, noisy logs slow diagnosis.

Downtime Avoidance Model

Leadership usually understands avoided outage cost better than avoided CPU.

Use:

```
avoided_loss = revenue_per_hour * avoided_outage_hours * affected_revenue_percentage
```

Example:

Input	Value
Revenue per hour	$80,000
Avoided outage	0.5 hours
Affected path	40 percent

```
80,000 * 0.5 * 0.40 = $16,000
```

If the pattern avoids even one 30-minute cascading outage, it can justify far more engineering time than the implementation costs.

Engineering Time Model

Incident response cost:

```
engineer_cost = engineers * hours * loaded_hourly_rate
```

Example:

Input	Value
Engineers involved	8
Hours per engineer	4
Loaded hourly rate	$120

```
8 * 4 * 120 = $3,840
```

This excludes opportunity cost. The same engineers are not shipping product while debugging cascading failure.

FinOps Dashboard Inputs

Track:

- caller pod count
- caller CPU during dependency incidents
- retry rate
- rejected call rate
- dependency QPS
- log ingestion rate
- metric series cardinality
- HPA desired replicas
- incident duration
- SLO burn rate

The FinOps story is strongest when you show before/after incident curves.

Leadership Summary

Use this language:

> Circuit breakers and bounded retries reduce the financial impact of dependency failures. They prevent caller services from wasting compute on doomed calls, reduce autoscaling spikes, lower observability noise, and shorten incident diagnosis. The main ROI is not steady-state CPU savings; it is avoided cascading downtime and reduced incident blast radius.

Purchase Value Checklist

This ebook is worth buying if the reader can leave with:

- a cost model they can present to leadership
- a lab to reproduce the failure mode
- a tested Go implementation
- alert and runbook templates
- production templates
- design review material

That is the difference between content and an operational asset.