

# BEYOND THE TOGGLE

Feature Flags + Measurement as Your Migration Safety Net

**Ryan Vila**

**FME Advisory Director**

# Every Migration Is a Risk Event

*Three failure modes that kill replatforms — and the pattern that prevents all three.*



## Big-bang cutover

One moment separates "working" from "broken." No migration is one service: UI, Data pipeline, Service Introduction, etc



## No rollback path

The new system is live. The old system is gone. Recovery means forward-only fixes under pressure.



## Blast radius you can't control

A bug ships to 100% of users, in every region, at the same time.

**Feature flags + measurement** turn a migration from an event into a **controlled experiment**.

# The Strangler Fig Pattern



## 1. Route Traffic

Direct a fraction of traffic to new cloud-based components while legacy remains the host.



## 2. Measure Impact

Monitor real-world performance and stability behind the feature flag.

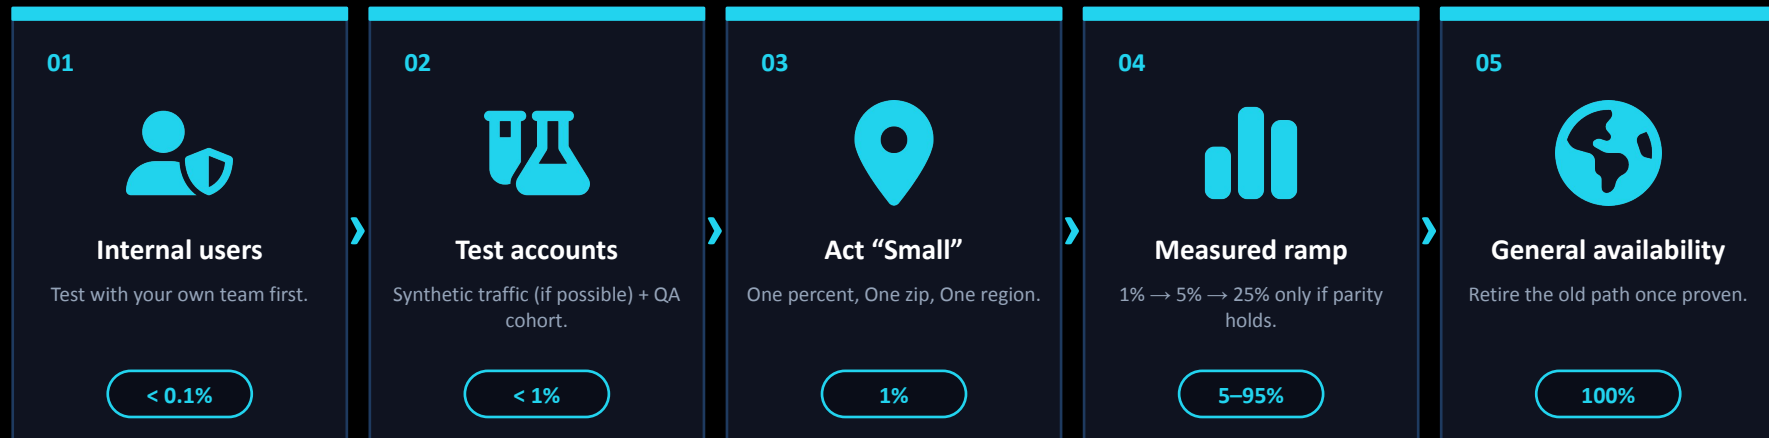


## 3. Retire Debt

Scale traffic to 100% and safely retire legacy services once proven.

# Canary Release

*Progressive rollout — earn every percent of traffic.*



## Measure frequently, early

Parity, error budget, cost, cohort — at 1%, not at 100%.



## Think like a load test

Each step is a controlled experiment. Scale only on green.



## Know your rollback path

Scale back, don't kill. Default treatment = one-click return.

If you can't **scale back or pause**, you don't have a canary — you have a **deployment**.

# What You Actually Watch at 1%

*A canary without metrics is just a slower outage.*



## Parity metrics

Is the new path matching the old path on latency, error rate, and throughput? Not "is it fast enough" — is it as fast as what we're replacing?



## Error budget burn

How much SLO headroom is this canary consuming per percent of traffic? If 1% is burning 10% of budget, do not scale.



## Cost per request

New infrastructure often has different cost curves. Watch unit economics at low traffic before you find them at full traffic.



## Cohort behavior

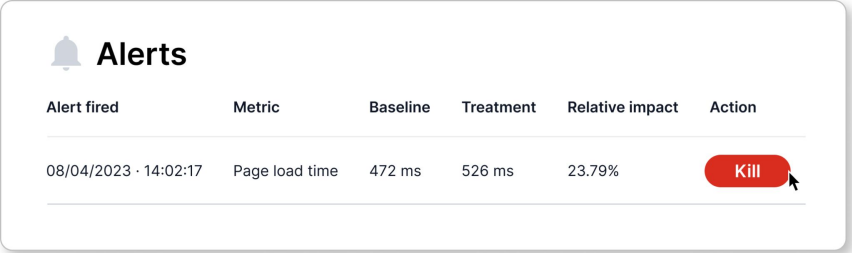
Aggregate numbers hide regressions. Segment by region, customer tier, SDK version — the failure is usually in a slice, not the whole.

# Measuring versus Monitoring

- Tightly couple rollouts with data
- Detect problems with minimal customer exposure
- Empower engineers with causation vs correlation
- Triage without wasting time

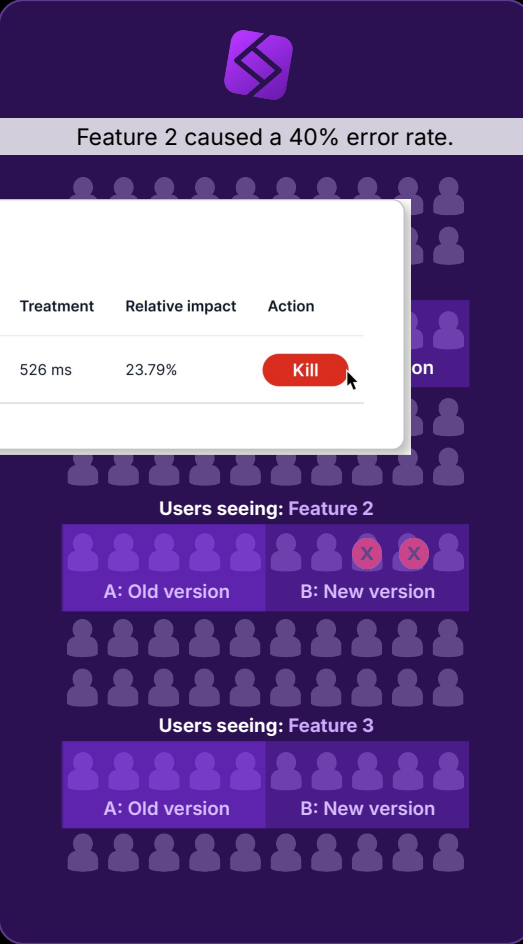
### Detecting problems in production

2% of users saw an error  
Unable to determine causality



Alert fired	Metric	Baseline	Treatment	Relative impact	Action
08/04/2023 · 14:02:17	Page load time	472 ms	526 ms	23.79%	<a href="#">Kill</a>

The background shows a grid of user icons, with two icons in the middle row marked with a red 'X', indicating a small number of users affected by the error.



Feature 2 caused a 40% error rate.

Users seeing: Feature 2

A: Old version      B: New version

Users seeing: Feature 3

A: Old version      B: New version

The diagram illustrates a feature rollout strategy. It shows two groups of users: 'Users seeing: Feature 2' and 'Users seeing: Feature 3'. Each group is split into 'A: Old version' and 'B: New version'. In the 'Feature 2' group, the 'B: New version' group has a 40% error rate, indicated by a red 'X' on the icon. In the 'Feature 3' group, both 'A: Old version' and 'B: New version' groups are shown without error indicators.

# Case Study: Rebuilding a Production Data Pipeline Behind a Flag

*How Split replaced its own S3 inbound integration with zero customer impact.*



## PROBLEM

- Single Spark batch job, sequential processing
- Scaling cost up 50%+ with customer demand
- Status reports missing data customers needed




## APPROACH

- New streaming pipeline built in parallel
- Flag controls which SQS queue each message enters
- Per-customer targeting to onboard gradually
- Default treatment = one-click rollback



## RESULT

- 80% lower resource cost
- 10 hours → 15 minutes for the same 50GB load
- Full quarter of migration, zero customer impact
- Improved status observability



# Flags Don't Save You From These

*Four ways teams defeat their own safety net.*



## The flag that never turns off

Two years later it's still there, branching every request, owned by nobody, removed by no one. Every stale flag is a landmine for the next engineer.



## The rollback nobody tested

You have a kill switch. You've never pulled it. The first time you try it will be at 3am during an incident, and the runbook is a Confluence page from 2023.



## Config drift across environments

The flag is 10% in staging, 50% in prod, 0% in DR. When prod fails over, what ships?



## Ownership ambiguity

Whose flag is this? Who decides when it graduates? Who removes it? If the answer is "the team that left last quarter," it's already broken.

**A flag controls a change and is a change itself. Own it, test the rollback, clean it up.**

# Start Monday.

*One migration. Three questions. If you can't answer all three in a sentence, that's your first project.*

**01**

## **What's behind a flag?**

Name one migration in flight. Is the new path gated? If not, why?

**02**

## **What's your kill switch?**

If the canary breaks, what flag change reverts it? Who can pull it at 3am?

**03**

## **What's your rollback SLA?**

From "we have a problem" to "we're back on known-good" — how many minutes?

# THANK YOU



Download The Beyond the Toggle EBook

