Deploying on friday afternoons without looking back with Flag

Progressive deployments

DECATHLON Digital

DEVSUMMIT
<MOVE_
FORWARD/>

Maxime Véroone

Staff Engineer @ Decathlon

<Friday prod push with Flagger/>

Agenda

DEVSUMMIT <MOVE_ FORWARD/>

1. Why ?

How we came to need Canary deployments @ Decathlon Ecommerce

2. What ?

Introducing Flagger, your new kubernetes best friend

3. How ?

Anatomy of a Canary deployment with Flagger

4. Learnings

Day 2 challenges & solutions.

DEVSUMMIT <MOVE_ FORWARD/>

2024 : an ambitious rebranding

- New logo
- New color palette
- Renewed brands selection
- New tagline
- New uniforms
- ...
- **New websites ?**

2024: Ambitious Rebranding

DEVSUMMIT <MOVE_ FORWARD/>

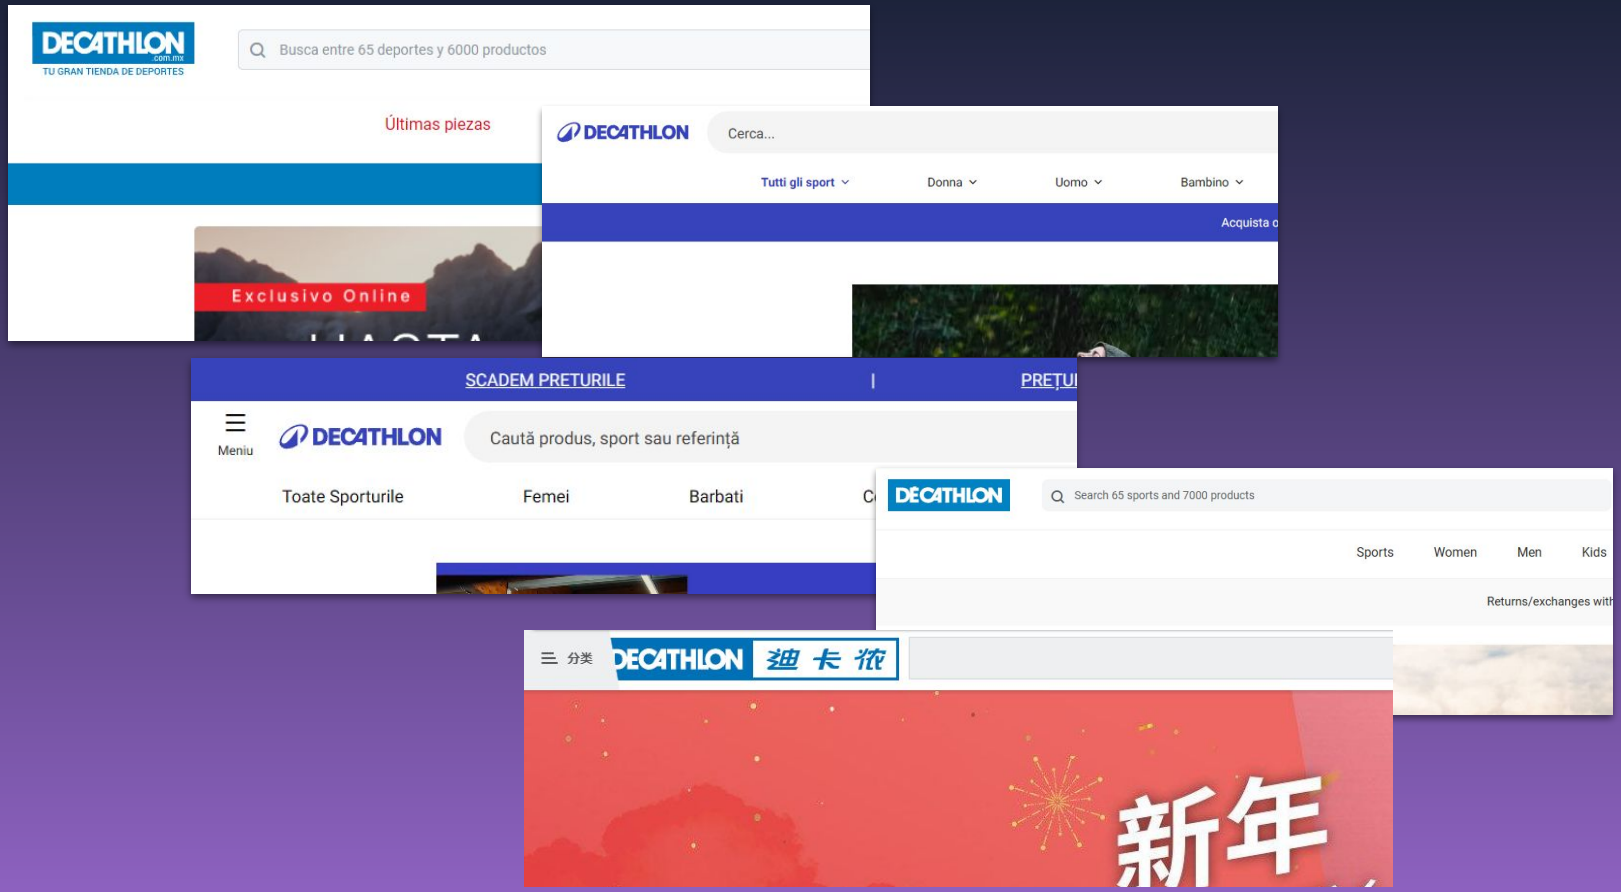

Multi-tenancy

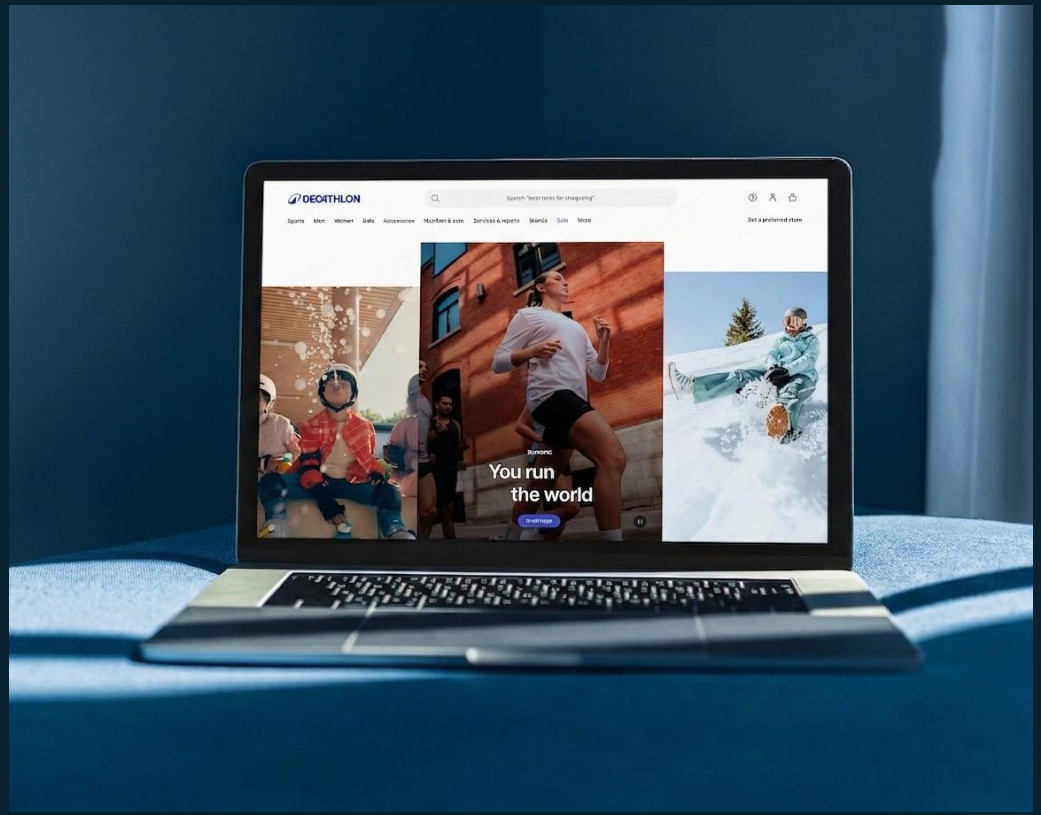

Multi-tenancy

Less workloads, less spread ...

... and no more guinea pigs !

What's so bad about an ecommerce incident ?

DEVSUMMIT <MOVE_FORWARD/>

40 stores

Business equivalent of e-commerce @ DKT

"- It was just an hour of downtime..."

"- We only sent Click & Collect 1h orders to stores after 55 minutes... "

"- Yesterday, we stopped sending orders to warehouses for 4 hours then firehosed 20 thousands in a few minutes..."

What can we do about it ?

DEVSUMMIT <MOVE_ FORWARD/>

Test, test and re-test !

which comes with downsides :

➡️ Impairs velocity

➡️ Flakiness exponential impact

➡️ Load testing is **HARD**

➡️ Some things are just untestable

DEVSUMMIT <MOVE_ FORWARD/>

What can we do about it ?

Just like our beloved products : we beta-test* websites with customers.

We test in prod.
Without hurting the business.
Even on friday evenings.

* cocreation.decathlon.fr

Introducing ...

DEVSUMMIT <MOVE_ FORWARD/>

First release : **2018**
GA : **2022**
Latest : **1.43***

Part of **Flux CD**
(CNCF Graduated)

*With our contributions inside

Reminder : Kubernetes Operator Concept

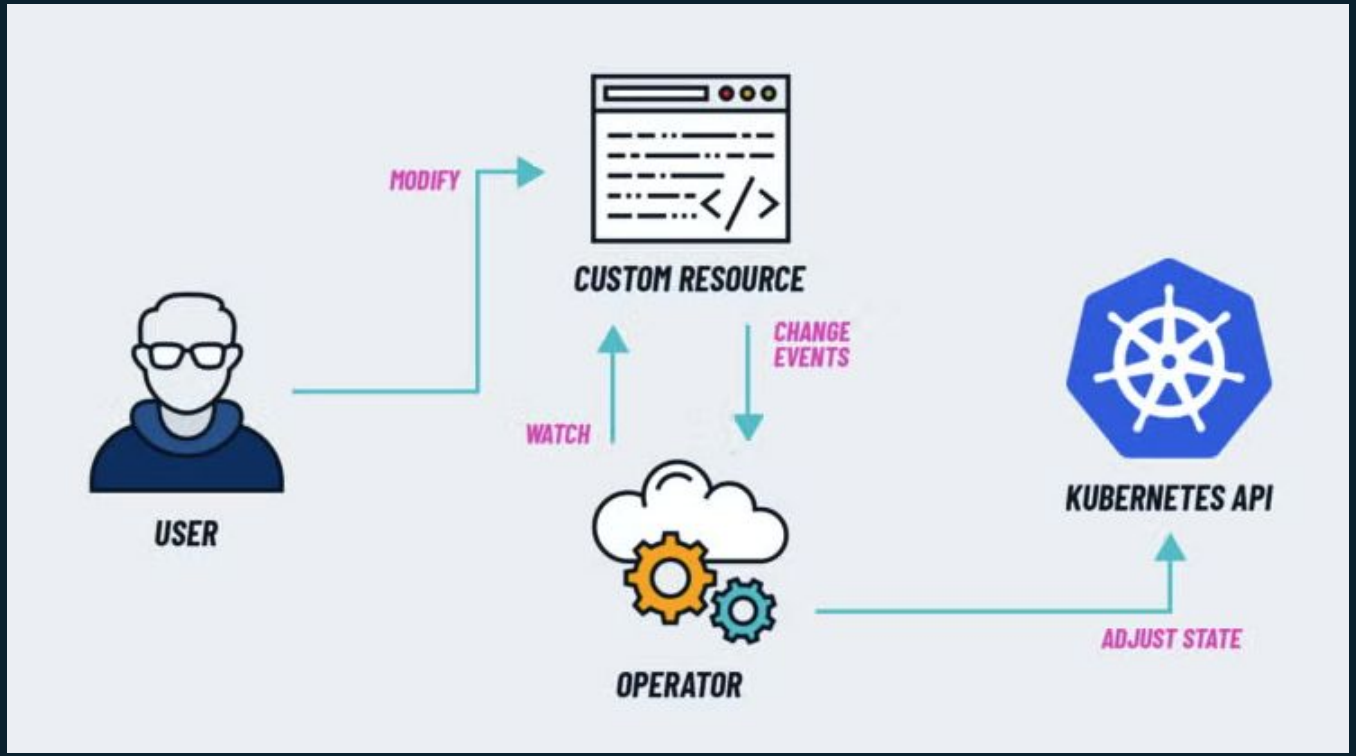

How Flagger works

Flagger concept

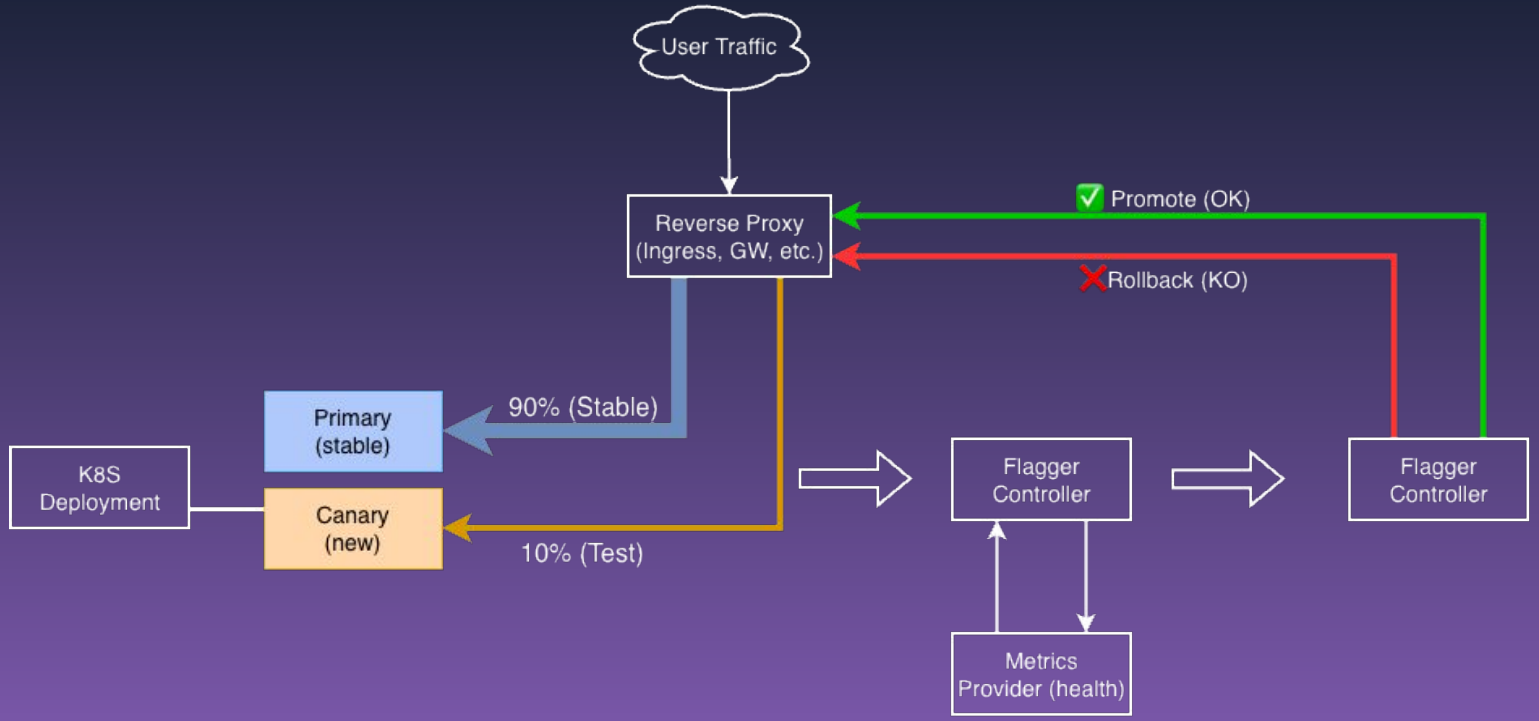

Flagger Strategies

- **Shadowing** : Mirroring real traffic without user impact
- **Blue/Green** : Instant switchover of 100% traffic
- **A/B Testing** : Targeting specific user populations
- ***Canary*** : Progressive deployment with increasing traffic steps

Example deployment

Strategy :
Canary
Steps : 5, 10, 20, 50
interval : 5 min

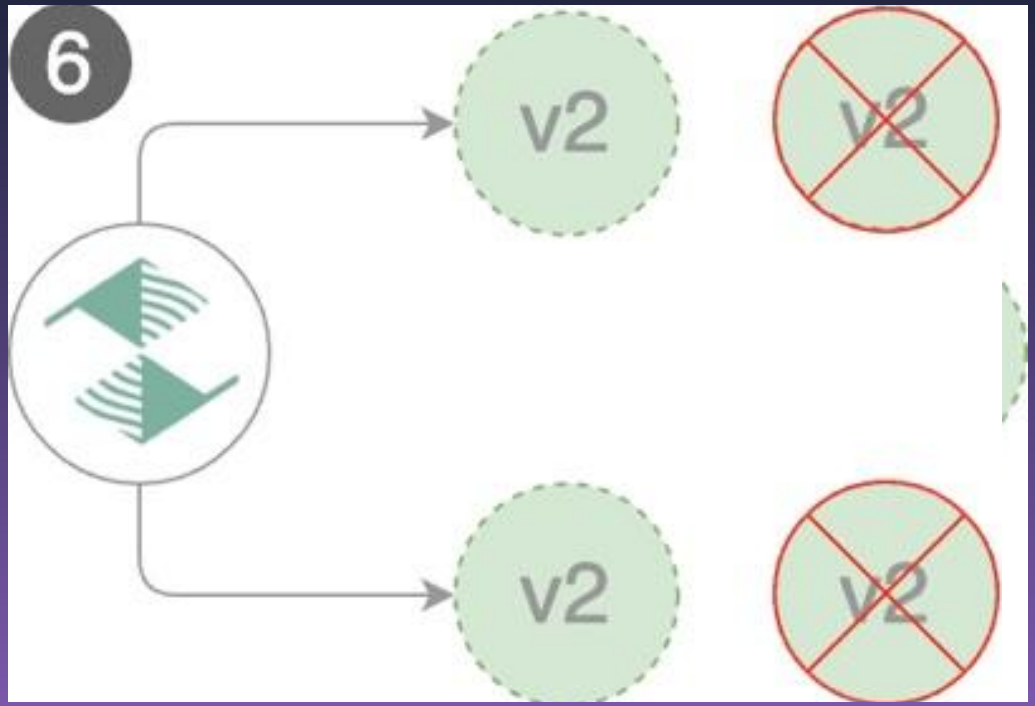

Reality

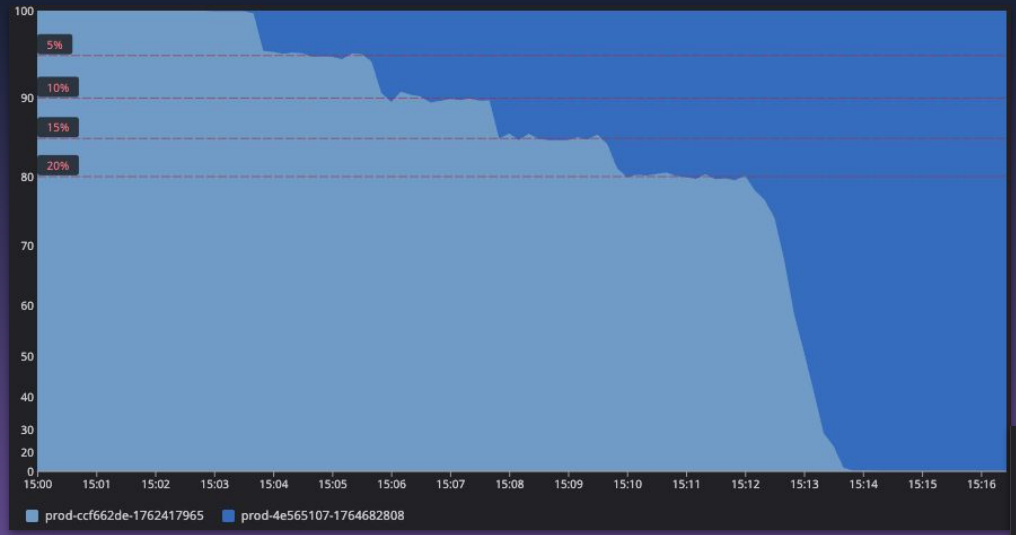

This is how it looks in our Dashboards during a Canary.

```
analysis:
  interval: '2m'
  stepWeight: 5
  maxWeight: 20
```

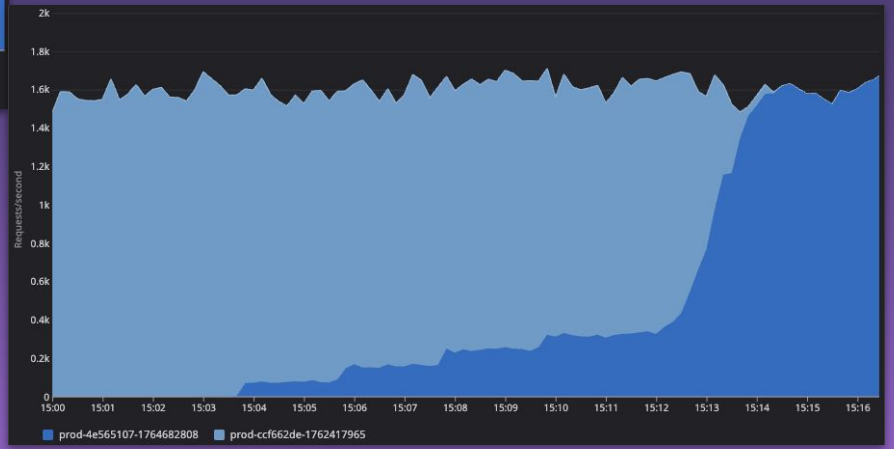

Flagger
is there still a pilot in the plane ?

DEVSUMMIT <MOVE_ FORWARD/>

Webhooks

Usage 1 :
Send notifications between each step of the deployment

Via:

pre-rollout (at start)

rollout (next step)

post-rollout (promotion done)

events (tout)

DEVSUMMIT <MOVE_ FORWARD/>

Webhooks

Usage 2 :
Allow manual control over the deployment cycle

hooks:
confirm-rollout,
confirm-traffic-increase,
rollback,
...

Flagger

Real life production challenges

Datadog & rate-limits

Problem:

- **Datadog** rate-limits its timeseries APIs

Impact :

- HTTP 429 == Canary unhealthy
  => Rollback

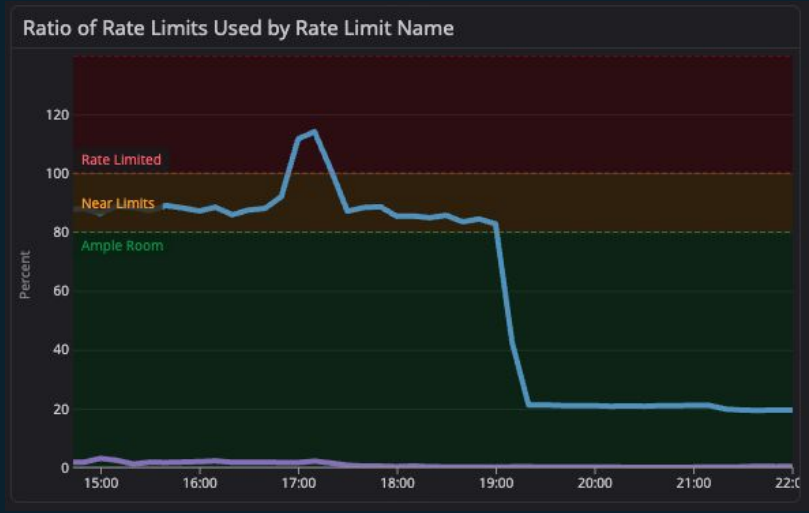

Solutions :

- Don’t use Flagger for KEDA & Flagger…
- Batch query via **DatadogMetric** CRD :
  Available since Flagger 1.43.0 Thanks to our contribution with Johan Lore

Flagger duplication of resources

DEVSUMMIT <MOVE_ FORWARD/>

Flagger automatically creates copies of ***certain*** **resources** to isolate the stable version.

These are suffixed with **-primary**

Including references between these resources in their definition!

deployments(france)[2]

NAME↑	READY	UP-TO-DATE	AVAILABLE	AGE
app-live	0/0	0	0	57d
app-live-primary	2/2	2	2	56d

Flagger duplication of resources

DEVSUMMIT <MOVE_ FORWARD/>

Nothing else gets duplicated !

- *Exemple :* DatadogMetric

Risk : surprises in behavior.

Solution : Duplicate them yourself or modify to aggregate both workloads

Sample size & measurement interval

Measuring several ”nines” of availability requires a lot of data points.
If your interval is too short, your measurement won’t be precise enough

You can’t measure 99.99% availability with 10 requests per interval

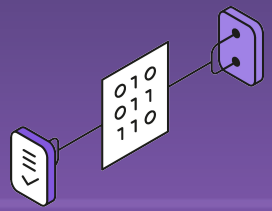

Sample size & measurement interval

Telemetry latency : Time to collect telemetry signals / metrics introduces delays in error detection during analysis.

➡️ **Solution :** step interval > collection-availability delay of telemetry

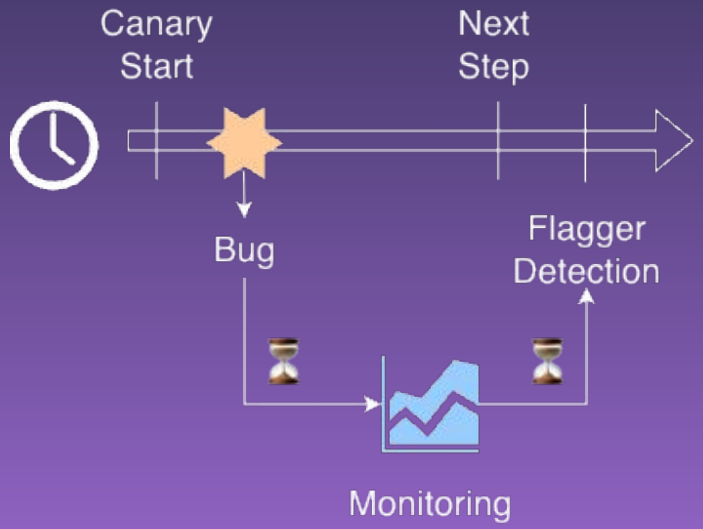

Manual Rollback

DEVSUMMIT <MOVE_ FORWARD/>

Auto-rollback is awesome, but...

- Sometimes your *metrics miss* the problems
- Sometimes they appear *later* (memory leak)

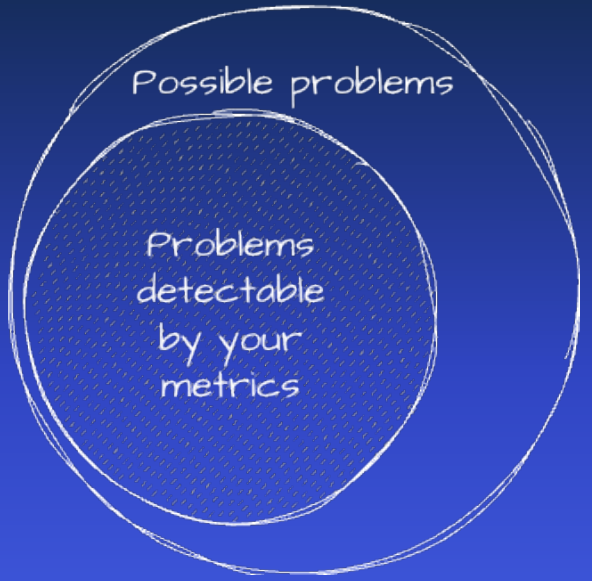

You need a “Rollback Now” button.

- During or After the progression
- Instantaneous, not progressive

Autoscaling at scale with Flagger

Challenge 1 : tuning the autoscaling for both Canary & Primary
You can either :

- Let Flagger **duplicate** your Autoscaler
  - Optionally overriding primary min/max
  - ⚠️ Changes will not be duplicated Until next promotion
- Create **different** autoscalers for each

```
autoscalerRef:
  apiVersion: autoscaling/v2
  kind: HorizontalPodAutoscaler
  name: podinfo
  primaryScalerReplicas:
    minReplicas: 2
    maxReplicas: 5
```

Autoscaling at scale with Flagger

Challenge 2 : By default, promotion is instantaneous.
If primary has scaled down it won't be able to handle nominal traffic.

Solution : Make the Promotion stage of the Canary progressive
(or use KEDA with forecast-based scale-down safeguards)

```
spec:
  analysis:
    interval: 90s
    stepWeightPromotion: 10
```

Autoscaling at scale with Flagger

Challenge 3 : Flagger requires 100% healthy pods before advancing, delaying progression if not met.
Volatile auto-scaling can cause frequent delays.

Solution: Allow a threshold of unhealthy pods:

```
spec:
  analysis:
    interval: 90s
    canaryReadyThreshold: 80
    primaryReadyThreshold: 80
```

The biggest challenge

DEVSUMMIT <MOVE_ FORWARD/>

Conducting change among teams

Working solutions for software engineers without them is bound to lack adoption

- Missing critical issues
- Solution not adapted to real needs

Providing tools is not enough.
What works is building these tools with people who will use them

<SPECIAL THANKS/>

THIS WAS A TEAM EFFORT !

DEVSUMMIT <MOVE_ FORWARD/>

Decathlon **Ecommerce Ops & DevOps** teams who did most of this effort.
I'm just the advertiser here

Decathlon Web Checkout & Product API teams
For being pioneers (mostly guinea pigs)

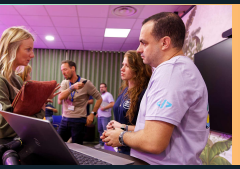

Dev Summit Organization team
for their huge efforts in bringing this event to life !

Special word for **Johan Lore** who was supposed to be with me today
But instead chose to visit Tokyo 🇯🇵 !

Thanks.