WOSCx4

TECHNISCHE UNIVERSITÄT BERLIN

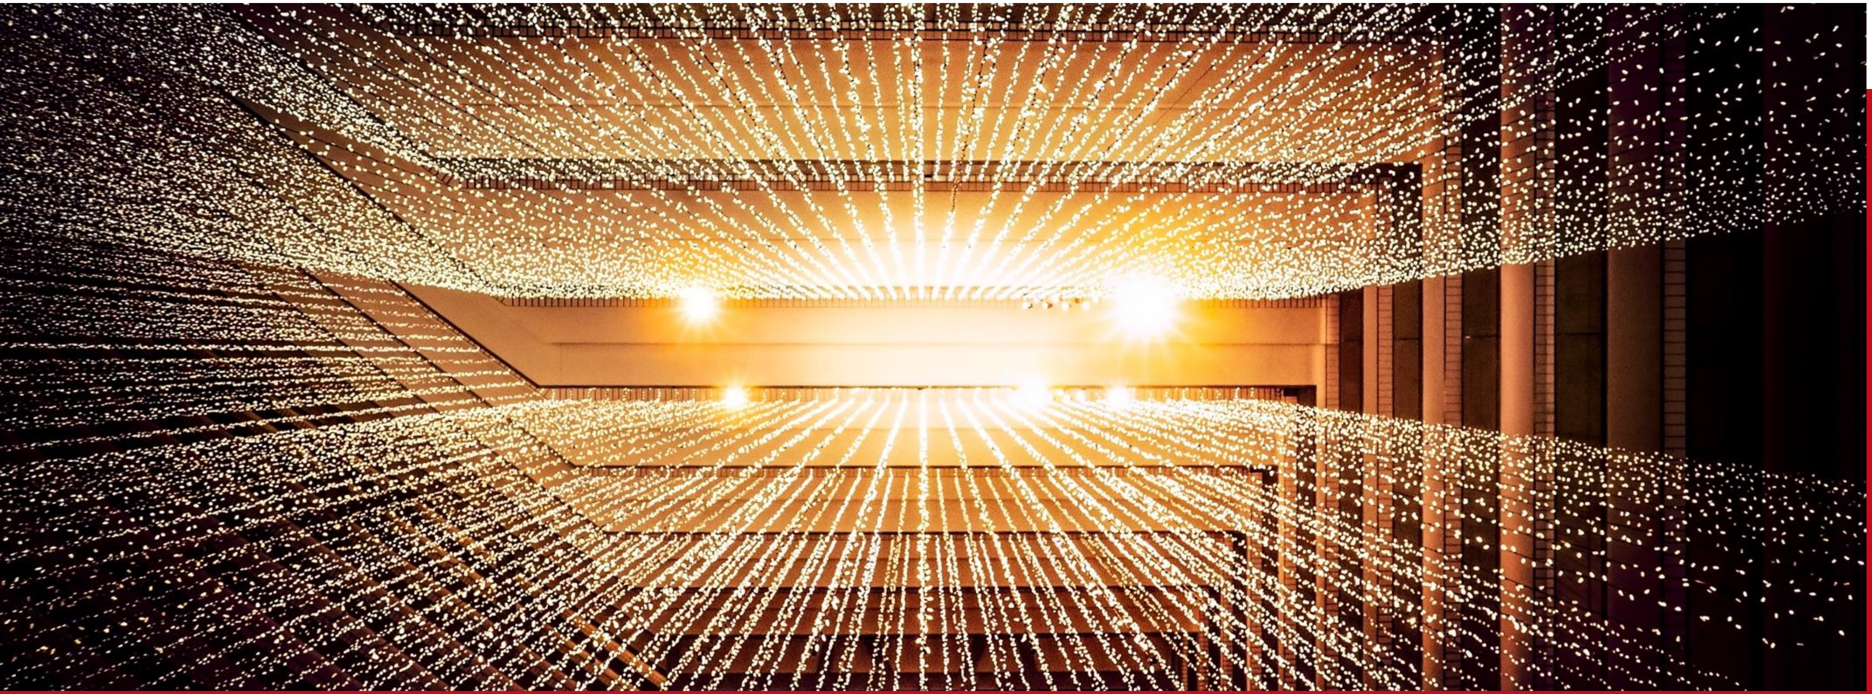

Umbilical Choir:

Automated Live Testing for Edge-To-Cloud FaaS Applications

M. Malekabbasi, T. Pfandzelter, D. Bermbach | Scalable Software Systems

EINSTEIN CENTER
Digital Future

3S

The Challenge: Risky Releases in Edge-to-Cloud FaaS

- Users demand high QoS[1]; frequent releases increase risk.
  - Amazon: Deploying every 11.6 seconds[2].
  - Siemens: Over 4 production deployments per month by GitLab and AWS

- **Edge(-to-Cloud) a different beast:** Standard cloud live testing (A/B, canary) doesn't fit.
  - Lacks **geo-awareness** for distributed edge nodes.
  - FaaS platforms offer limited native live testing features.

- **Result:**
  Developers hard-code complex, error-prone solutions, or skip thorough testing.

Elfic, a comedy juggler

1. Quality of Service
2. DevOps Research and Assessment (DORA) State of DevOps Report

Umbilical Choir (UC): Automated Live Testing for the Edge

- A novel framework for **serverless edge-to-cloud live testing**.

- **Idea: Tree-like Hierarchy**
  - **Release Managers (RMs):**
    Plan & delegate strategies across regions.
  - **Agents:**
    Execute tests on FaaS platforms
    (via a Proxy Function for traffic control).

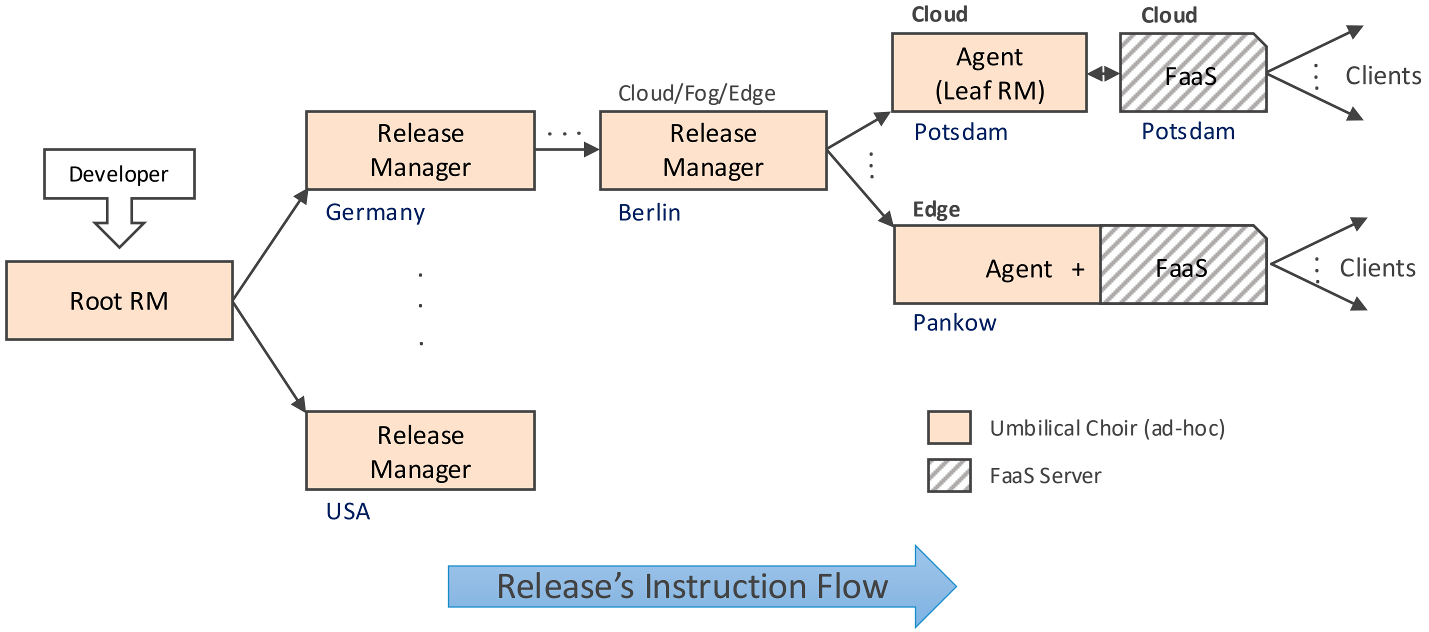

- **Key Design:** Scalable, geo-aware by structure,
  FaaS platform-agnostic (through a Proxy function),
  works with edge network limitation (no push/ssh access).

UC: Declarative & Flexible Live Testing

- **Declarative Strategies:** Define multi-stage releases.
  - Supports: A/B, Canary, Dark Launch, Gradual Rollout.
  - Metric-based conditions for progression/rollback.
- **Defining Rollouts in Geo-Dist. Environment:**
  - **Global Incremental:** (Good for critical patches).
  - **Local Sequential:** (Minimizes risk).
  - **Regional Incremental:** Global Incremental applied to a specific region first.
  - **Regional Sequential:** Local Sequential applied within a specific region.
- These can be combined for complex scenarios.

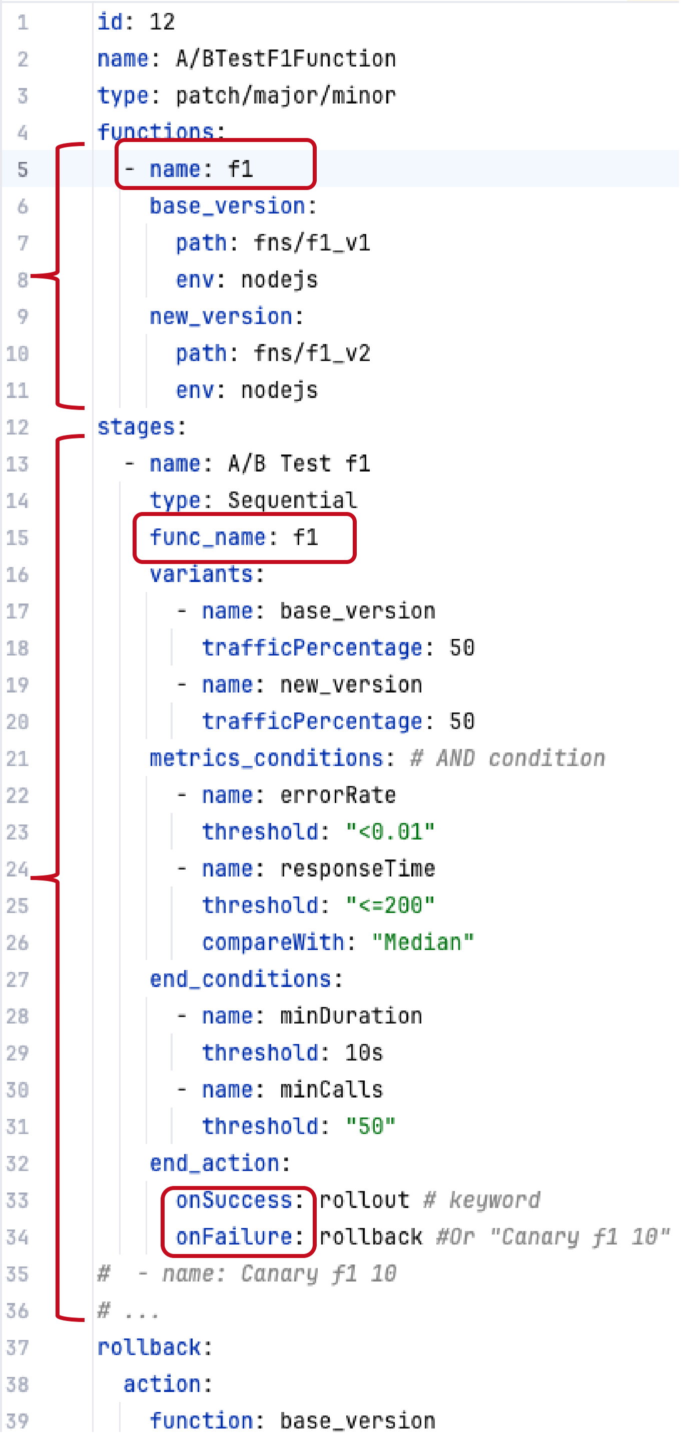

A simple "Release Strategy"

```yaml
id: 12
name: A/BTestF1Function
type: patch/major/minor
functions:
  - name: f1
    base_version:
      path: fns/f1_v1
      env: nodejs
    new_version:
      path: fns/f1_v2
      env: nodejs
stages:
  - name: A/B Test f1
    type: Sequential
    func_name: f1
    variants:
      - name: base_version
        trafficPercentage: 50
      - name: new_version
        trafficPercentage: 50
    metrics_conditions: # AND condition
      - name: errorRate
        threshold: "<0.01"
      - name: responseTime
        threshold: "<=200"
        compareWith: "Median"
    end_conditions:
      - name: minDuration
        threshold: 10s
      - name: minCalls
        threshold: "50"
    end_action:
      onSuccess: rollout # keyword
      onFailure: rollback #Or "Canary f1 10"
#   - name: Canary f1 10
# ...
rollback:
  action:
    function: base_version
```

Umbilical Choir Works in Practice

- **Open-Source Prototype (Go):** Supports major FaaS providers (e.g., GCP Cloud Run) and edge FaaS (e.g., tinyFaaS).
- **Proxy Overhead:** Measurable, but enables crucial features.
- **Complex Scenario Success:** Successfully orchestrated a multi-stage, geo-distributed release (canary -> A/B -> gradual rollout) across edge (Pi+tinyFaaS) and cloud (GCP Cloud Run).

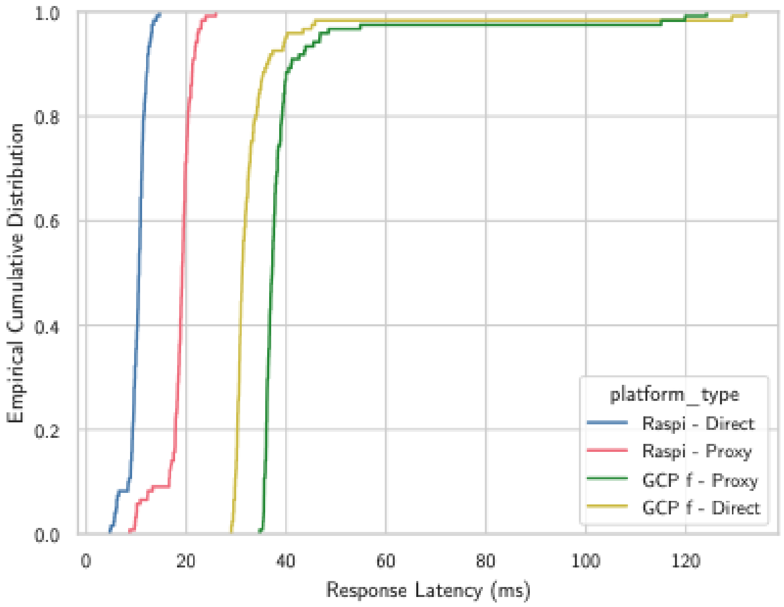

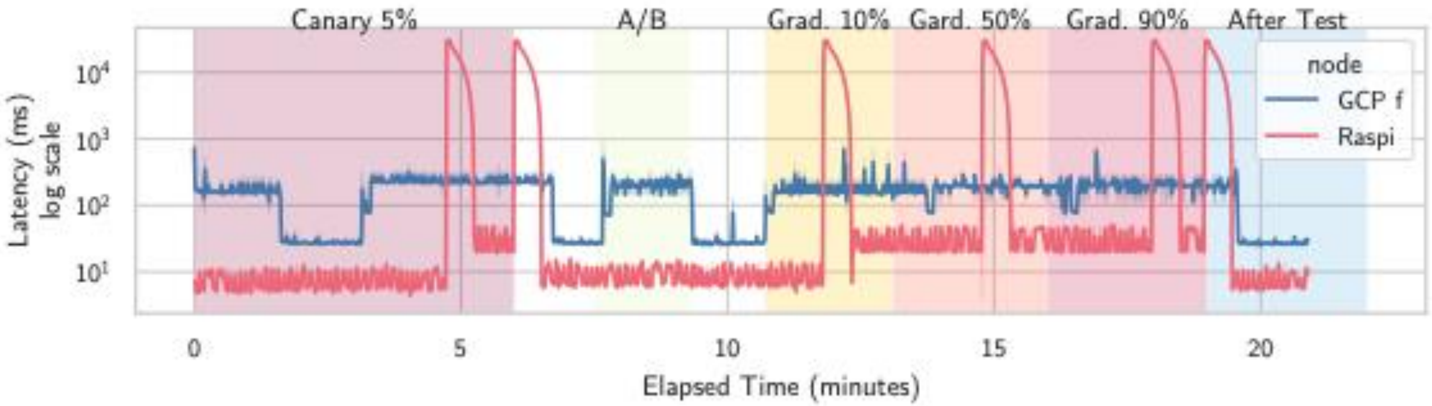

Conclusion: Enabling Safer Edge-to-Cloud Releases

- ➢ Edge-to-cloud FaaS needs robust, geo-aware live testing.

- ➢ **Umbilical Choir provides a solution:**
  - ➢ Automates complex serverless release strategies (incl. A/B testing, dark launches, canary releases with geo-aware strategies)
  - ➢ Works across any FaaS platform
  - ➢ Flexible Routing Methods (supports sticky sessions for tests).

- ➢ **Impact:** Empowers developers to automatically evaluate new FaaS application releases in the evolving edge-to-cloud landscape.

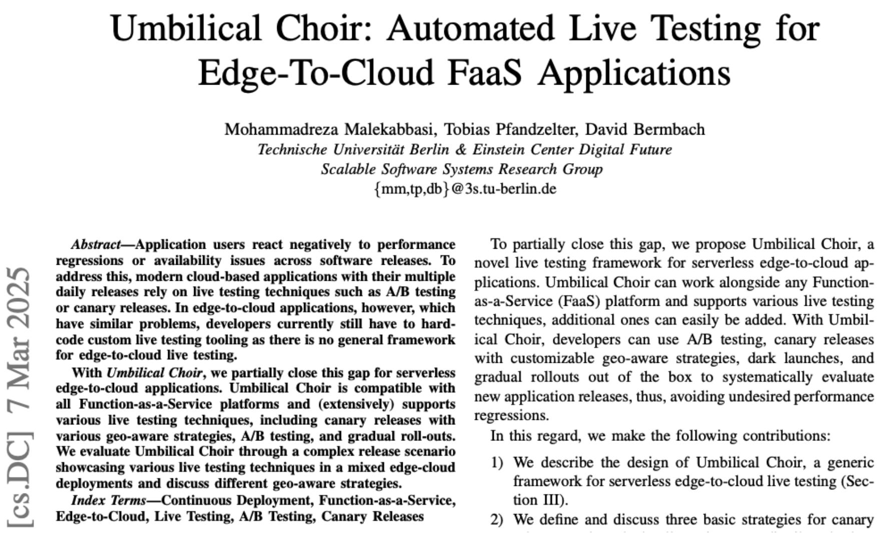

Umbilical Choir: Automated Live Testing for Edge-To-Cloud FaaS Applications

Mohammadreza Malekabbasi, Tobias Pfandzelter, David Bermbach
Technische Universität Berlin & Einstein Center Digital Future
Scalable Software Systems Research Group
{mm,tp,db}@3s.tu-berlin.de

[cs.DC] 7 Mar 2025

***Abstract*—Application users react negatively to performance regressions or availability issues across software releases. To address this, modern cloud-based applications with their multiple daily releases rely on live testing techniques such as A/B testing or canary releases. In edge-to-cloud applications, however, which have similar problems, developers currently still have to hard-code custom live testing tooling as there is no general framework for edge-to-cloud live testing.**

With *Umbilical Choir*, we partially close this gap for serverless edge-to-cloud applications. Umbilical Choir is compatible with all Function-as-a-Service platforms and (extensively) supports various live testing techniques, including canary releases with various geo-aware strategies, A/B testing, and gradual roll-outs. We evaluate Umbilical Choir through a complex release scenario showcasing various live testing techniques in a mixed edge-cloud deployments and discuss different geo-aware strategies.

***Index Terms*—Continuous Deployment, Function-as-a-Service, Edge-to-Cloud, Live Testing, A/B Testing, Canary Releases**

To partially close this gap, we propose Umbilical Choir, a novel live testing framework for serverless edge-to-cloud applications. Umbilical Choir can work alongside any Function-as-a-Service (FaaS) platform and supports various live testing techniques, additional ones can easily be added. With Umbilical Choir, developers can use A/B testing, canary releases with customizable geo-aware strategies, dark launches, and gradual rollouts out of the box to systematically evaluate new application releases, thus, avoiding undesired performance regressions.

In this regard, we make the following contributions:

1) We describe the design of Umbilical Choir, a generic framework for serverless edge-to-cloud live testing (Section III).
2) We define and discuss three basic strategies for canary

To be published at ICFEC 2025

Thank You & Questions!

Contribution and Technical Meat