

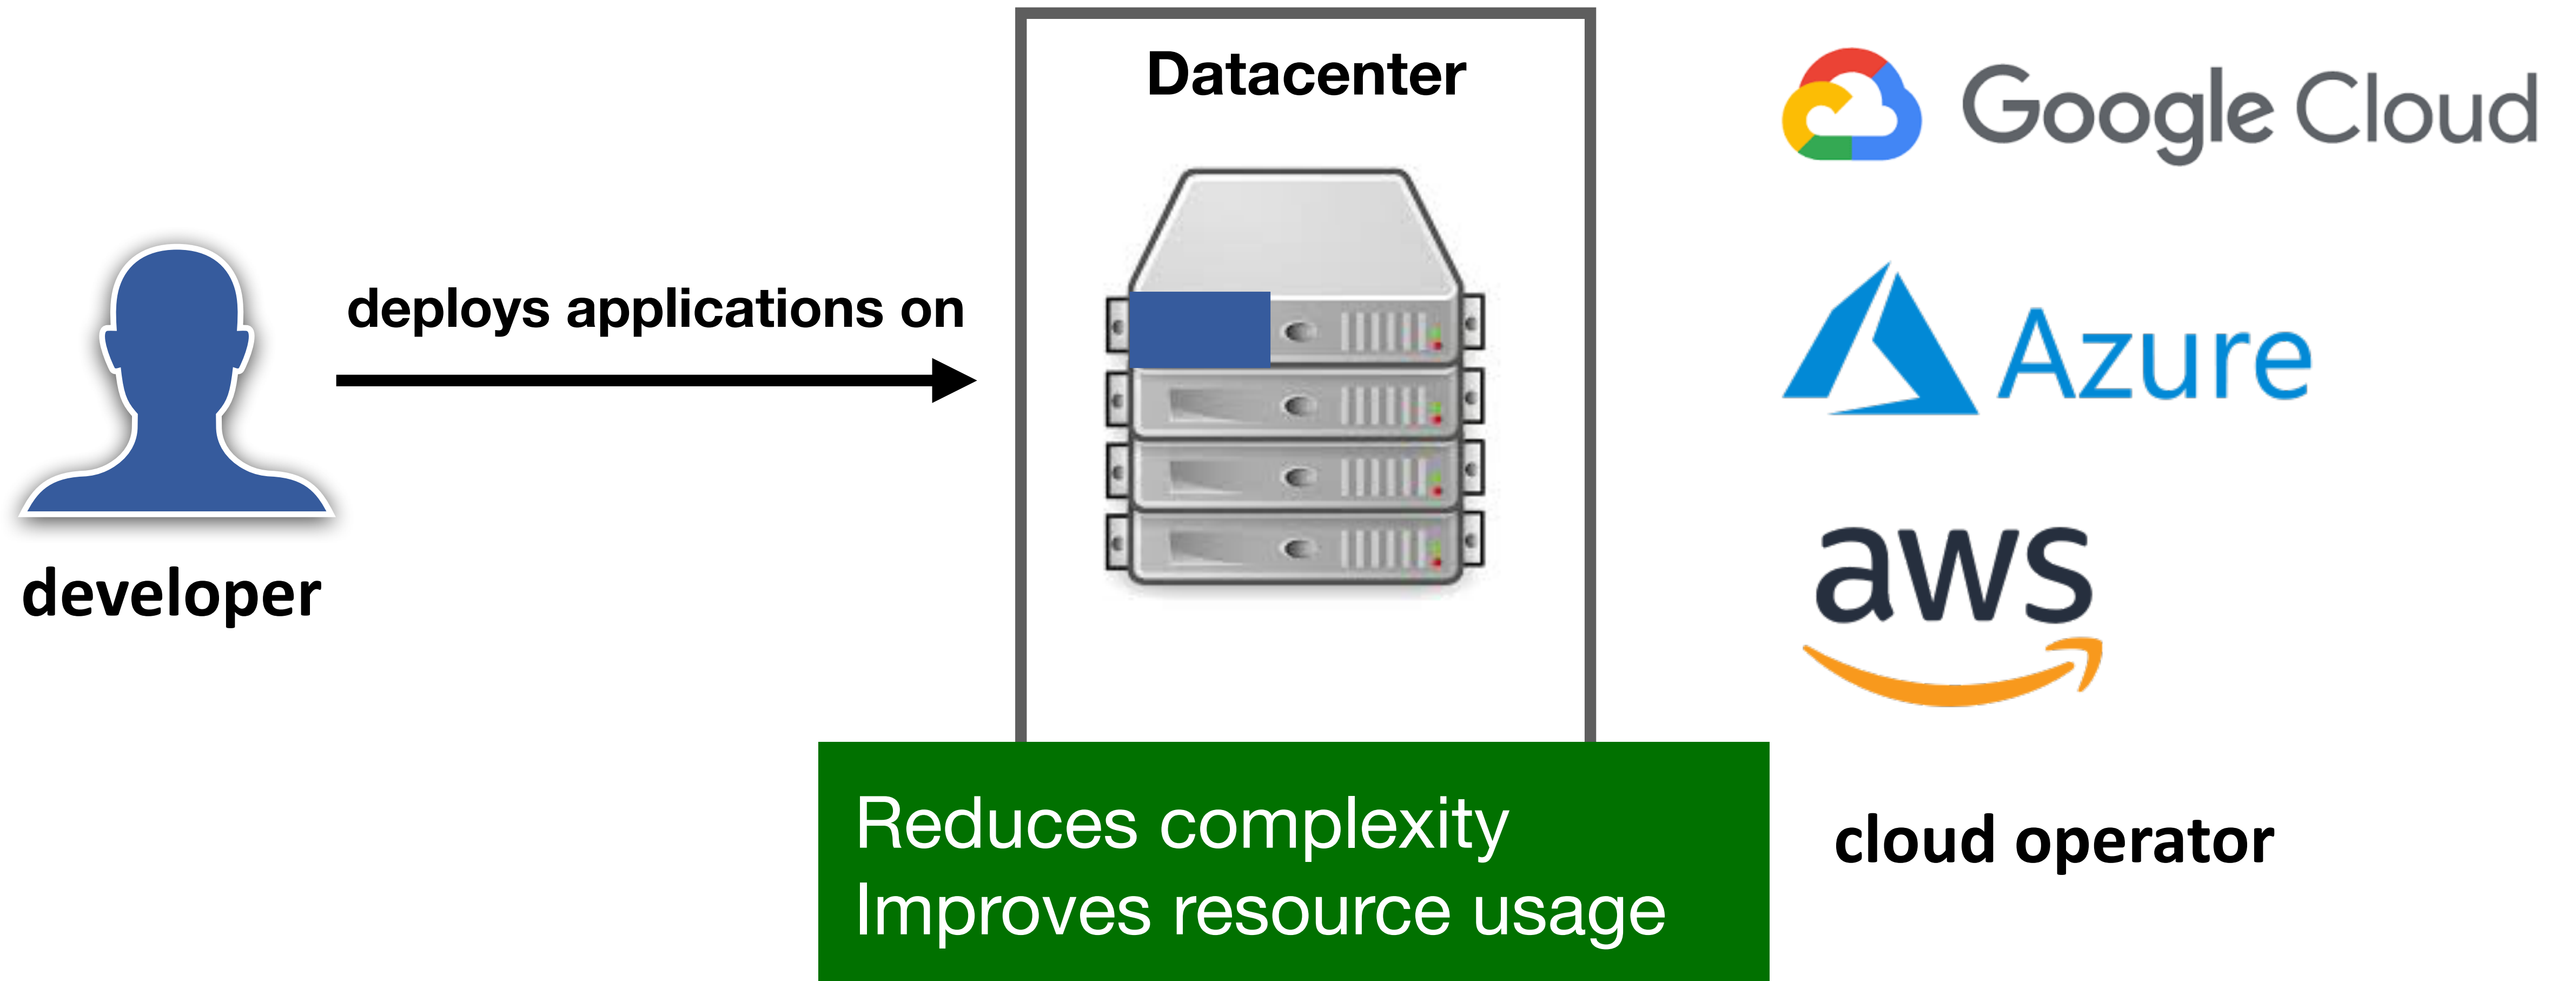


# The Design, Productization, and Evaluation of a Serverless Workflow-Management System

Erwin van Eyk




# Cloud Computing






# Complexity and resource utilization remain key challenges






A photograph of a server room. The ceiling is white with a grid of fluorescent lights. Below the ceiling, there are rows of server racks. The racks are filled with various electronic components, including circuit boards and cables. Some cables are green, some are black, and some are red. The racks are arranged in a long row, and the perspective is looking down the aisle between them. A semi-transparent white box with black text is overlaid on the middle of the image.

**Serverless Computing** is a form of cloud computing which allows users to run event-driven and granularly-billed applications without having to deal with operational logic.



# Division of operational concerns

cloud operator manages...



developer manages...

Application

Application  
Middleware

Cluster Resource Mgt

Virtualization

Hardware

DIY

Application

Application  
Middleware

Cluster Resource Mgt

Virtualization

Hardware

Cloud / IaaS

Application

Application  
Middleware

Cluster Resource Mgt

Virtualization

Hardware

serverless




# The “Serverless” market

currently: **\$ 5** billion

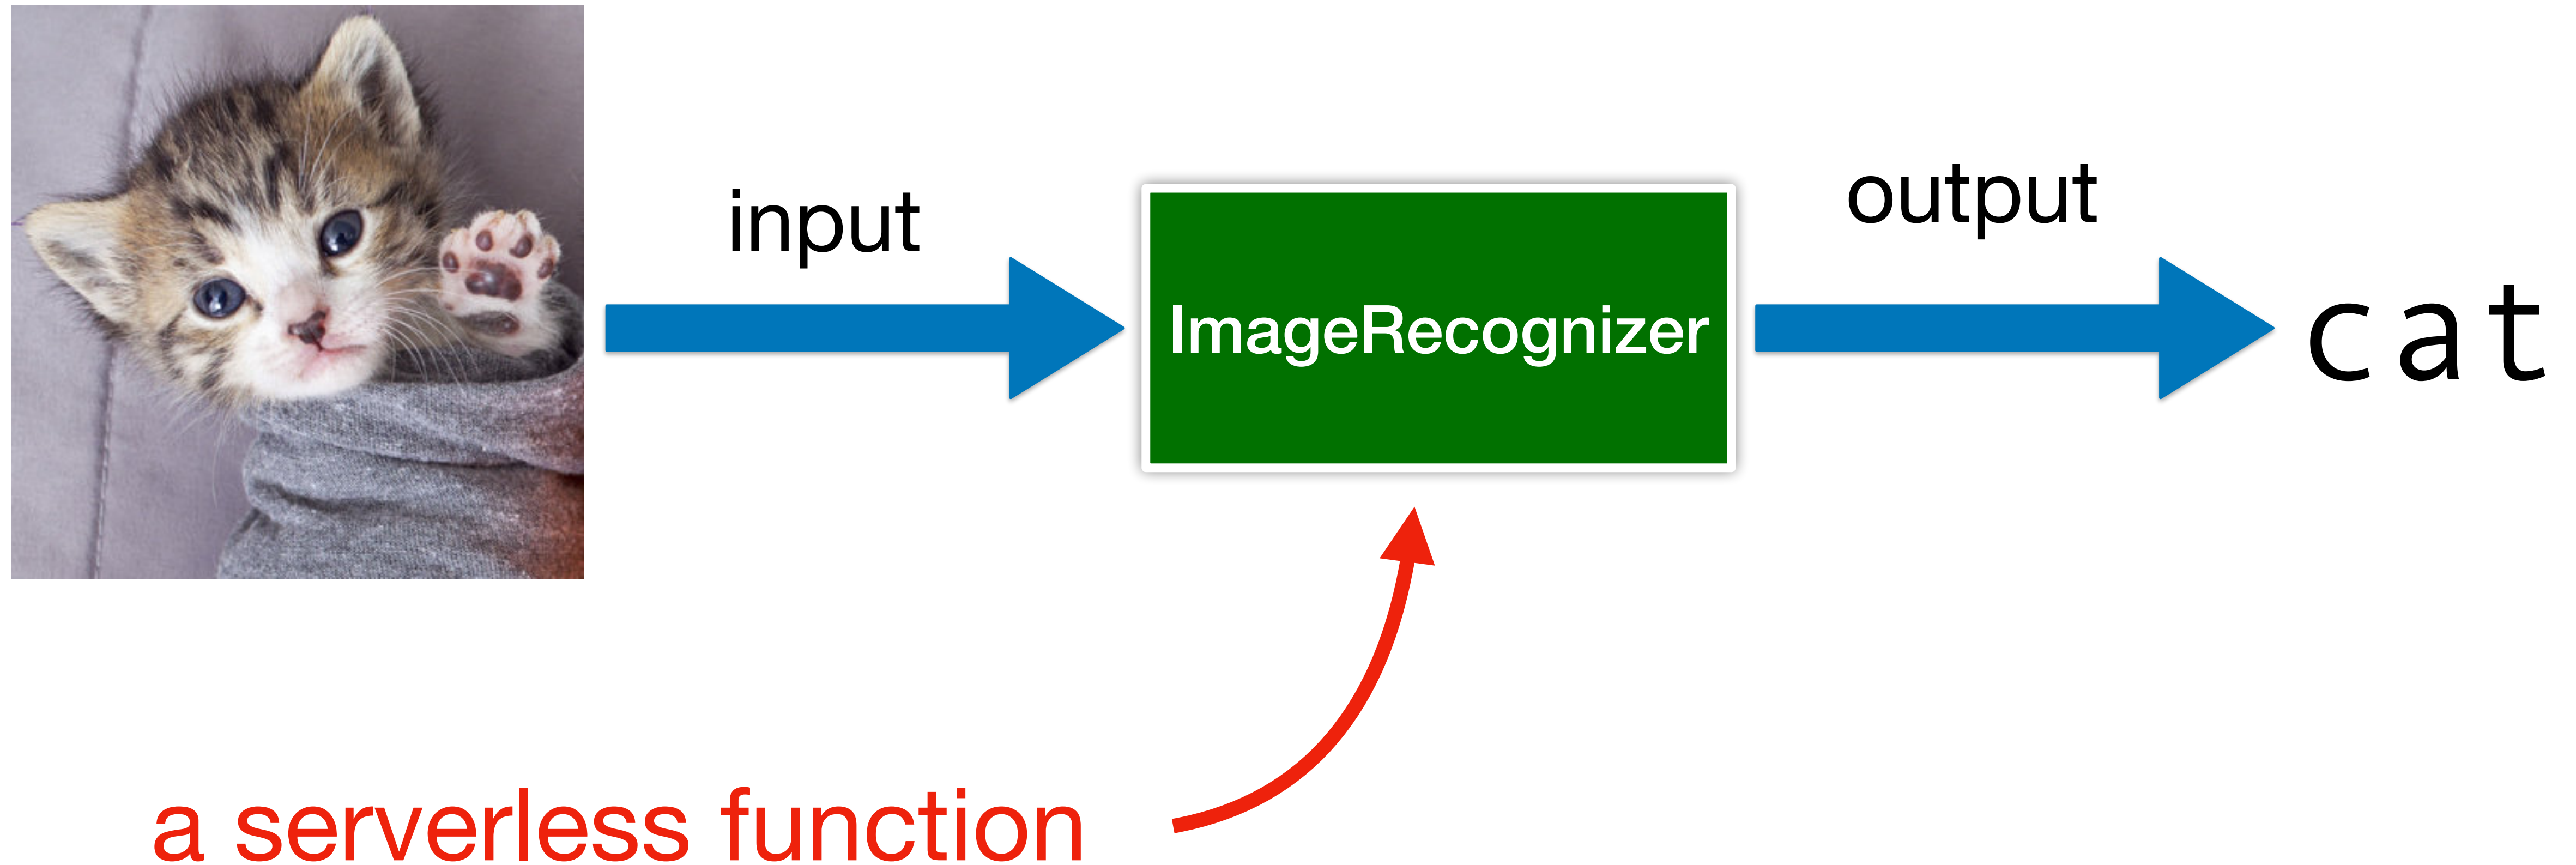
by 2023: **\$ 15** billion

(predicted)



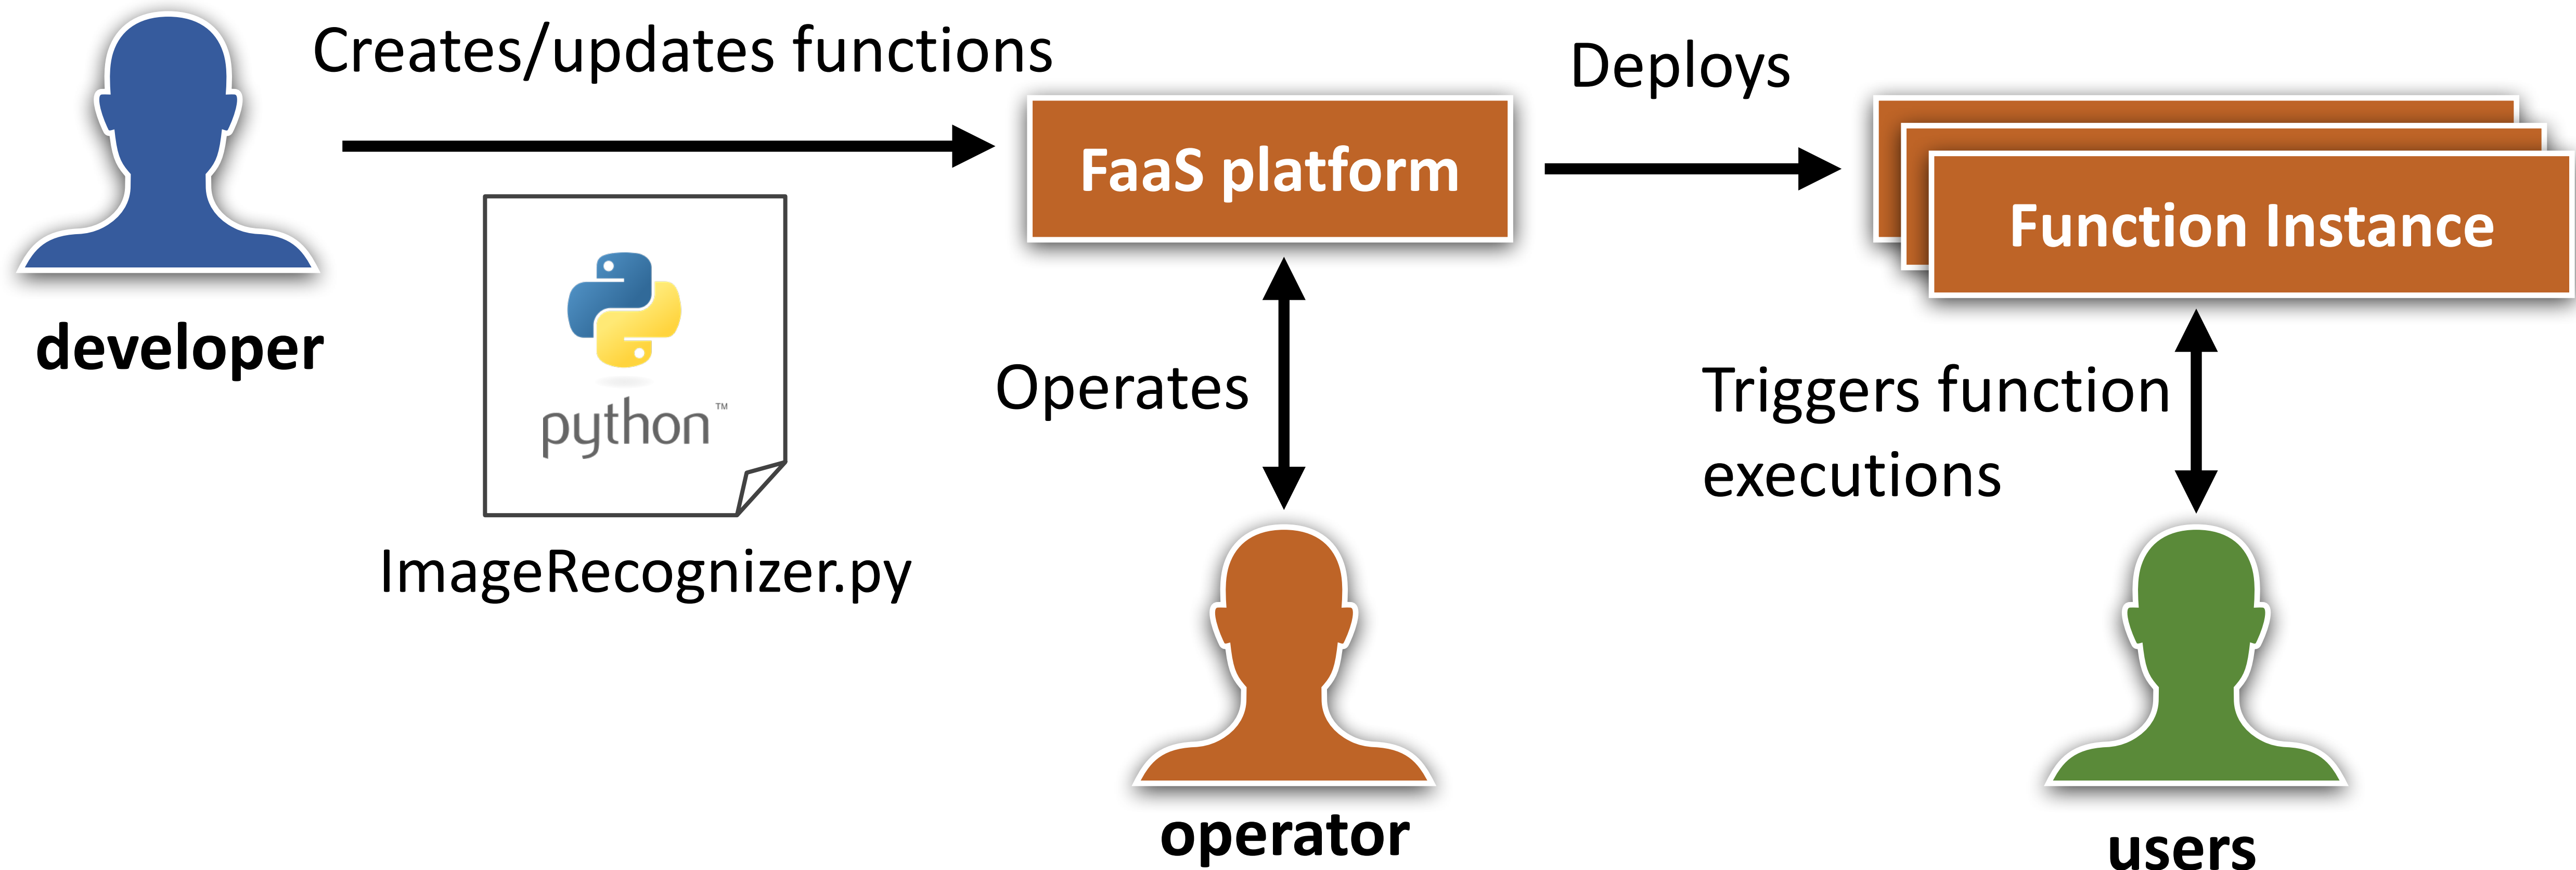


# Function-as-a-Service (FaaS)





# Function-as-a-Service (FaaS) in a nutshell



# New technology, new problems

- Undefined terminology
- Lacking fundamental models and principles
- Absent real-world data and workload traces
- Missing well-established systems, tooling and processes



# A multi-level approach

**This thesis!**



**High-level, community problems**

Terminology

Challenges & perspectives

FaaS reference architecture




*and others*

**Conceptual problems (in related domains)**

Emergence of serverless

Scheduling in FaaS

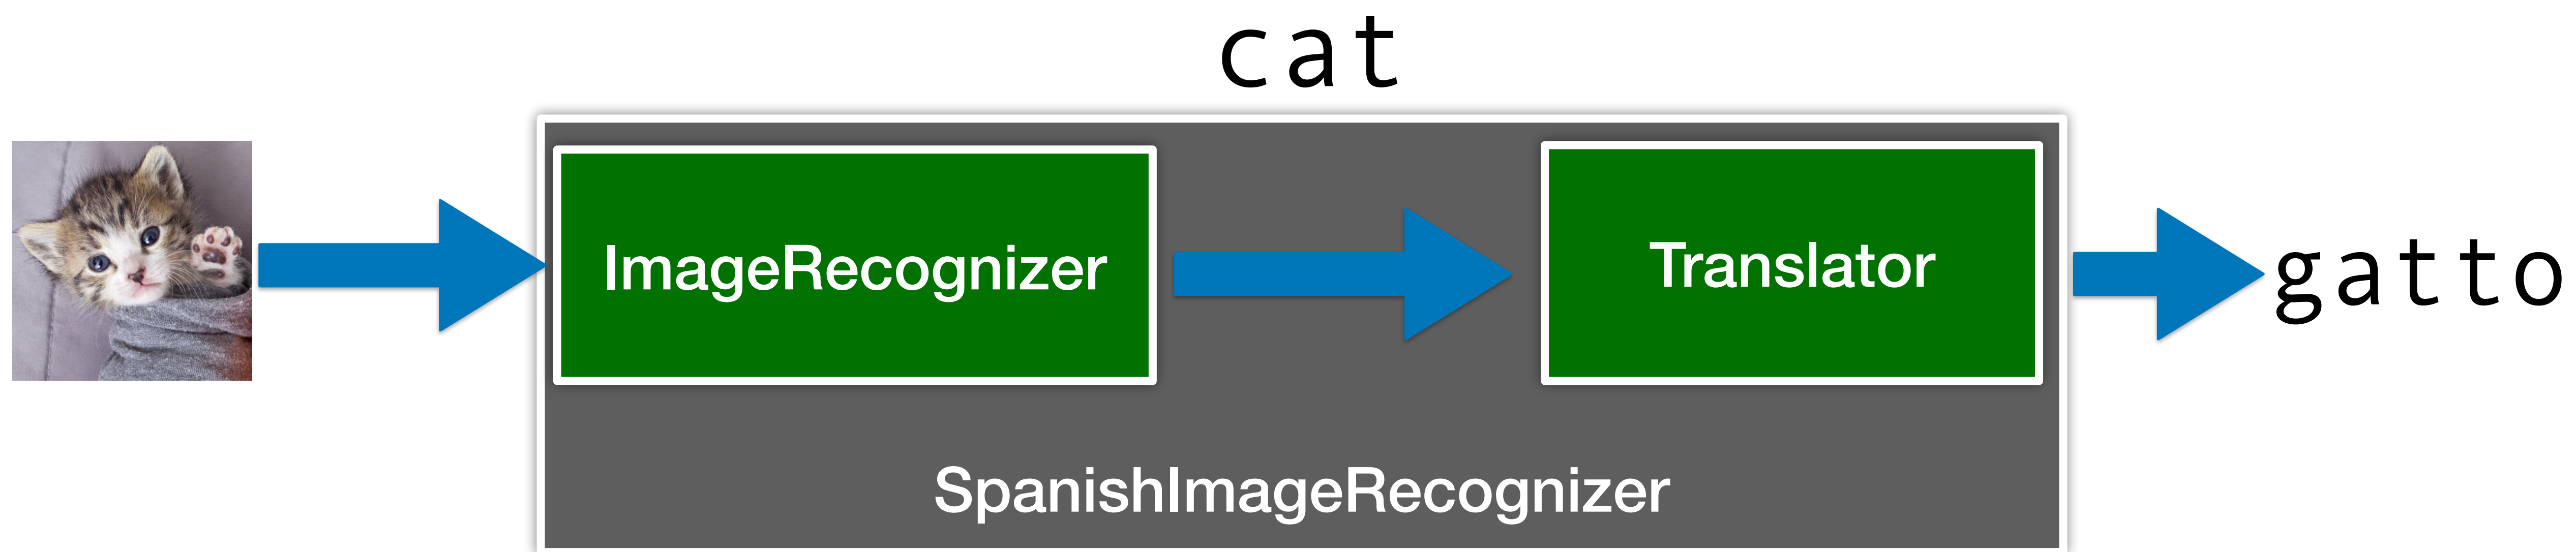


**A specific conceptual and technical problem**

Serverless function composition

# Serverless function composition

*“Orchestrating existing functions into new more complex, composed functions.”*




RQ: How to support serverless function composition through a distributed systems approach?



# Requirements analysis

RQ1: What are the requirements to support serverless function compositions?





# Requirements

1. Support long-running processes.
2. Minimize performance overhead.
3. Ensure reliable executions.
4. Scale to frequent workflow invocations.

# Constraints

1. Treat FaaS functions as black-boxes.
2. Rely only on common FaaS functionality.
3. Follow the serverless function development and execution model.

*(Summarised; see thesis for the complete requirements and constraints)*



# Surveying the State-of-the-Art

RQ2: How do approaches for the composition of cloud functions compare?

1. FaaS platforms
2. Serverless function composition approaches
3. Workflow management systems
4. Workflow languages


# (1) FaaS platform survey

Serverless Cloud Native Landscape  
v20180525

See the serverless interactive landscape at [s.cncf.io](https://s.cncf.io)

Greyed logos are not open source

Tools



Security




Framework



Hosted

Installable

Platform




source: <https://github.com/cncf/wg-serverless>



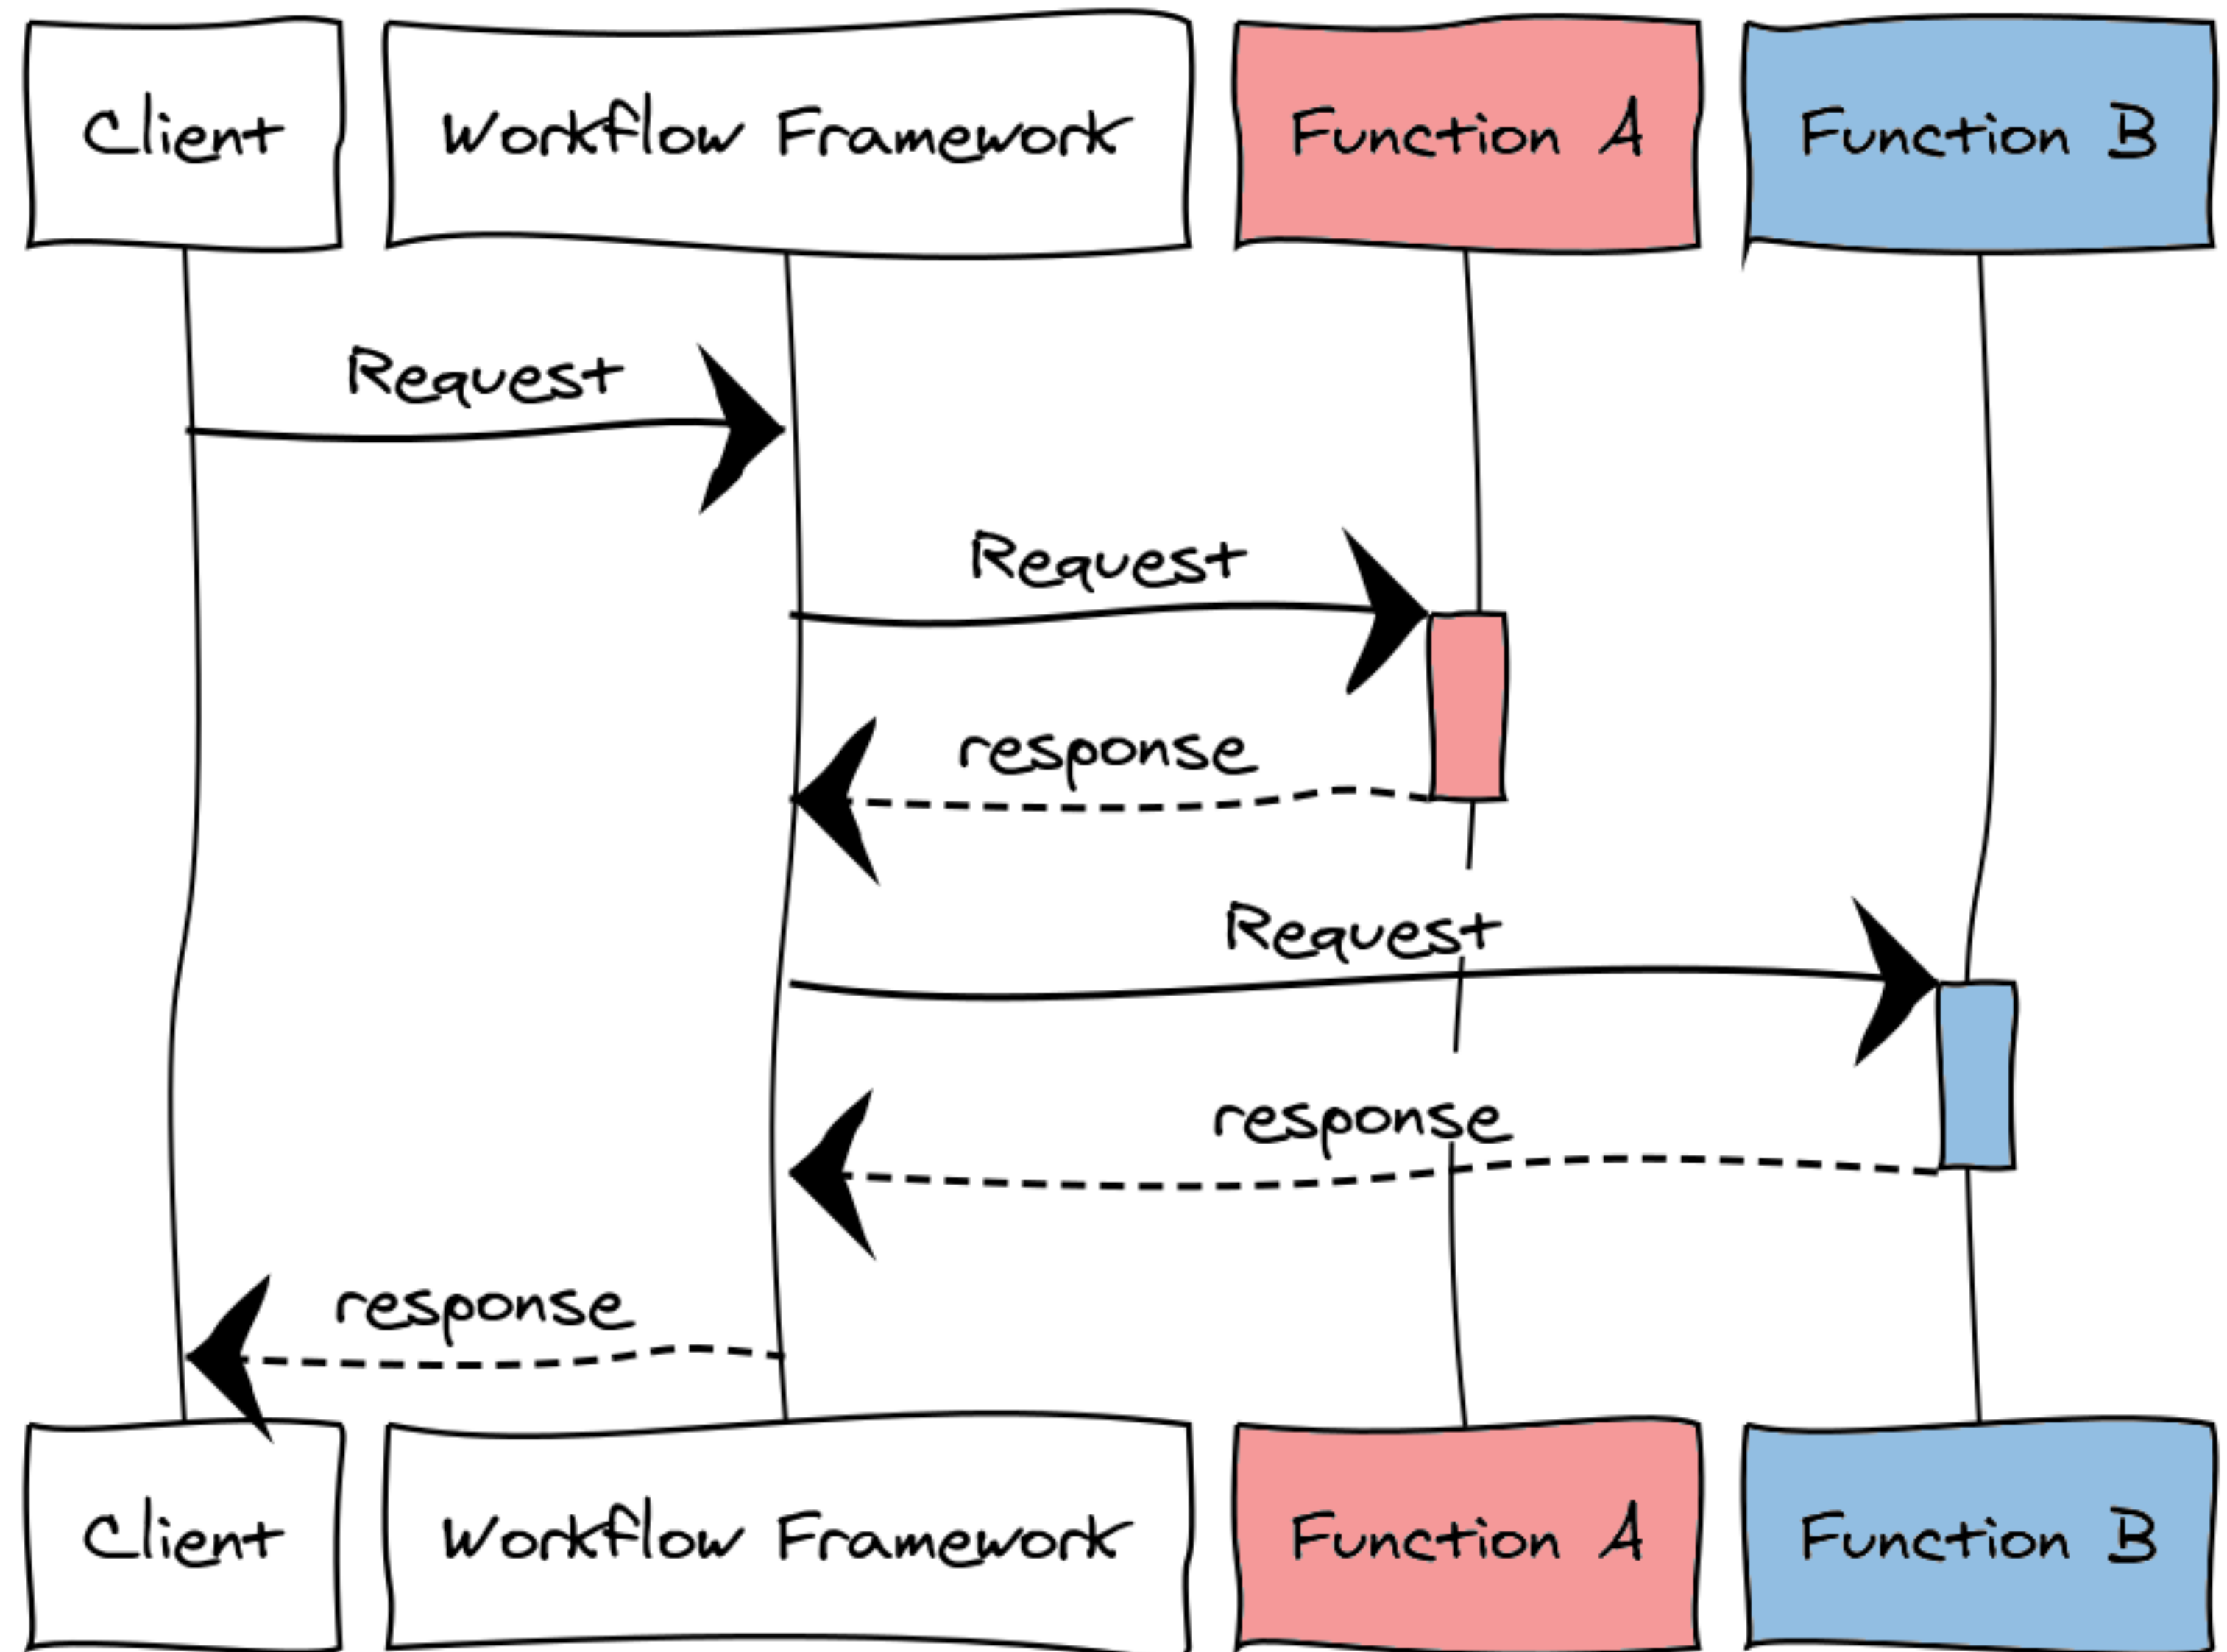
## **(2) Approaches to function composition survey**

1. Direct composition
2. Compiled composition
3. Coordinator-based composition
4. Event-driven composition
5. Workflow-based composition

# Workflow-based composition




**Workflow definition**






# (3+4) Survey of workflow systems and languages



Azure Logic Apps



Node-RED



Apache  
Airflow



**MISTRAL**  
*an OpenStack Community Project*



AWS Step Functions




Pegasus



- Most serverless-like WMS are closed-source.
- Workflow languages are often tightly coupled to the WMS.
- Workflow languages contain system-level assumptions.
- Reuse common language syntax and system architecture.

# Design of a serverless workflow management system

RQ3: How to design a system for composing cloud functions?



**SWL**

Serverless Workflow Language

(see thesis)

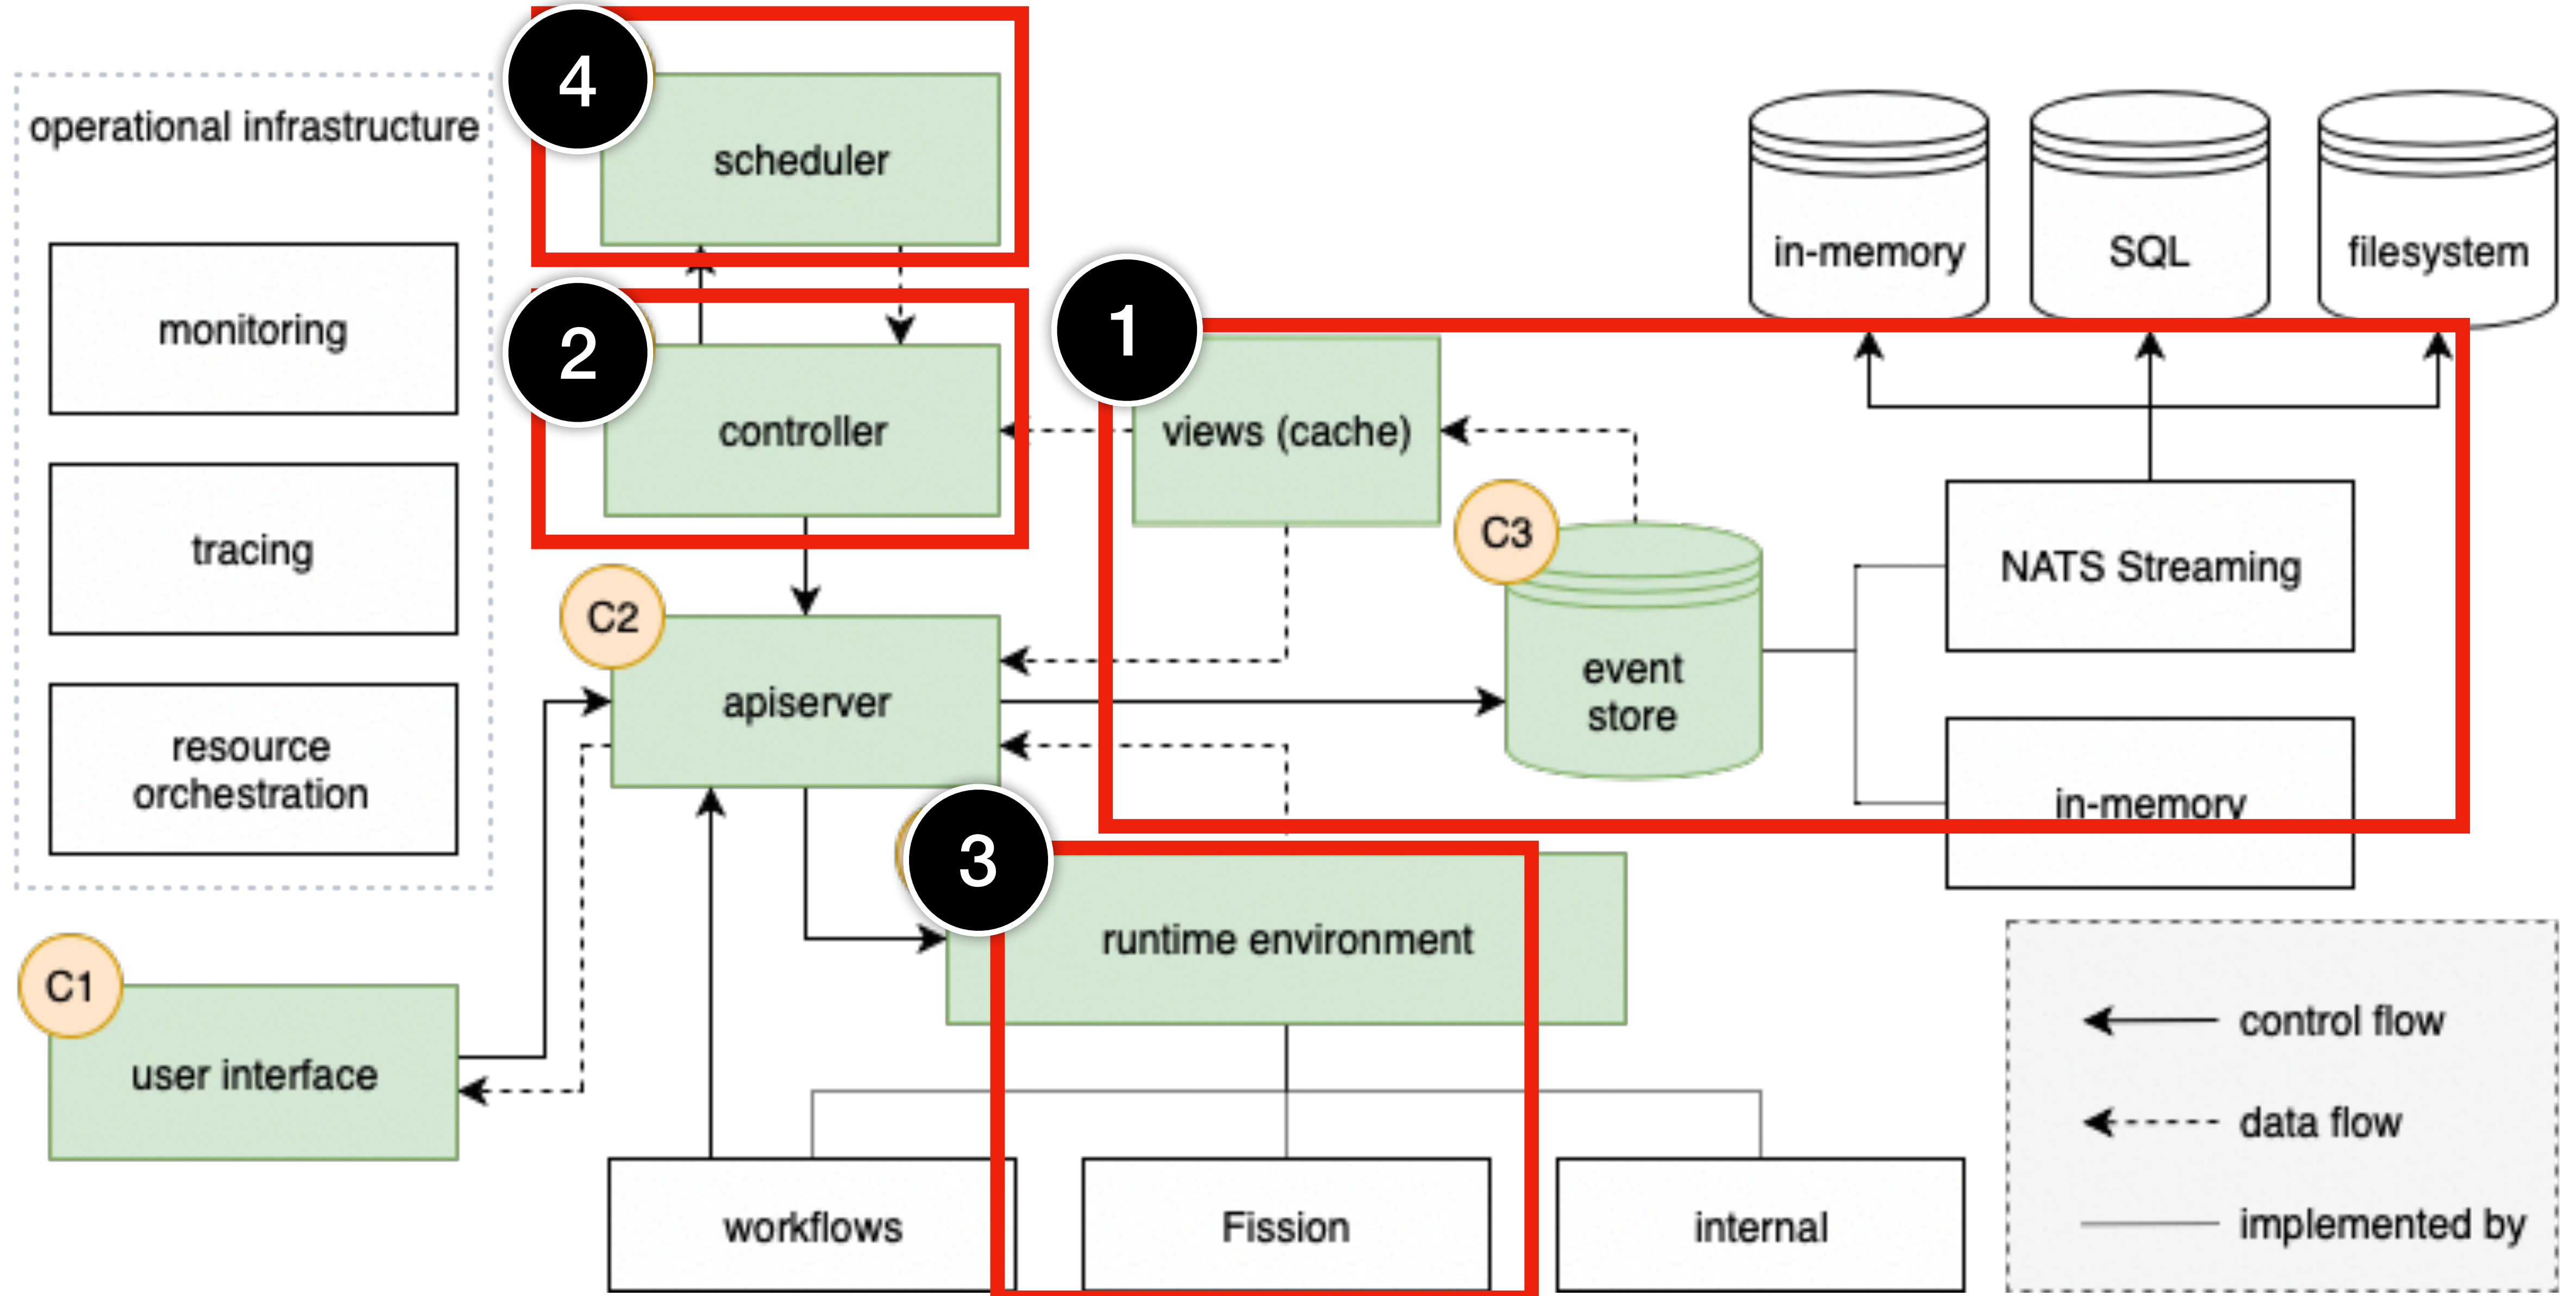


**fission**  
workflows

Serverless workflow management system




# Fission Workflows



# Event sourcing and persistence

The workflows are stored as a sequence of events



## Impact

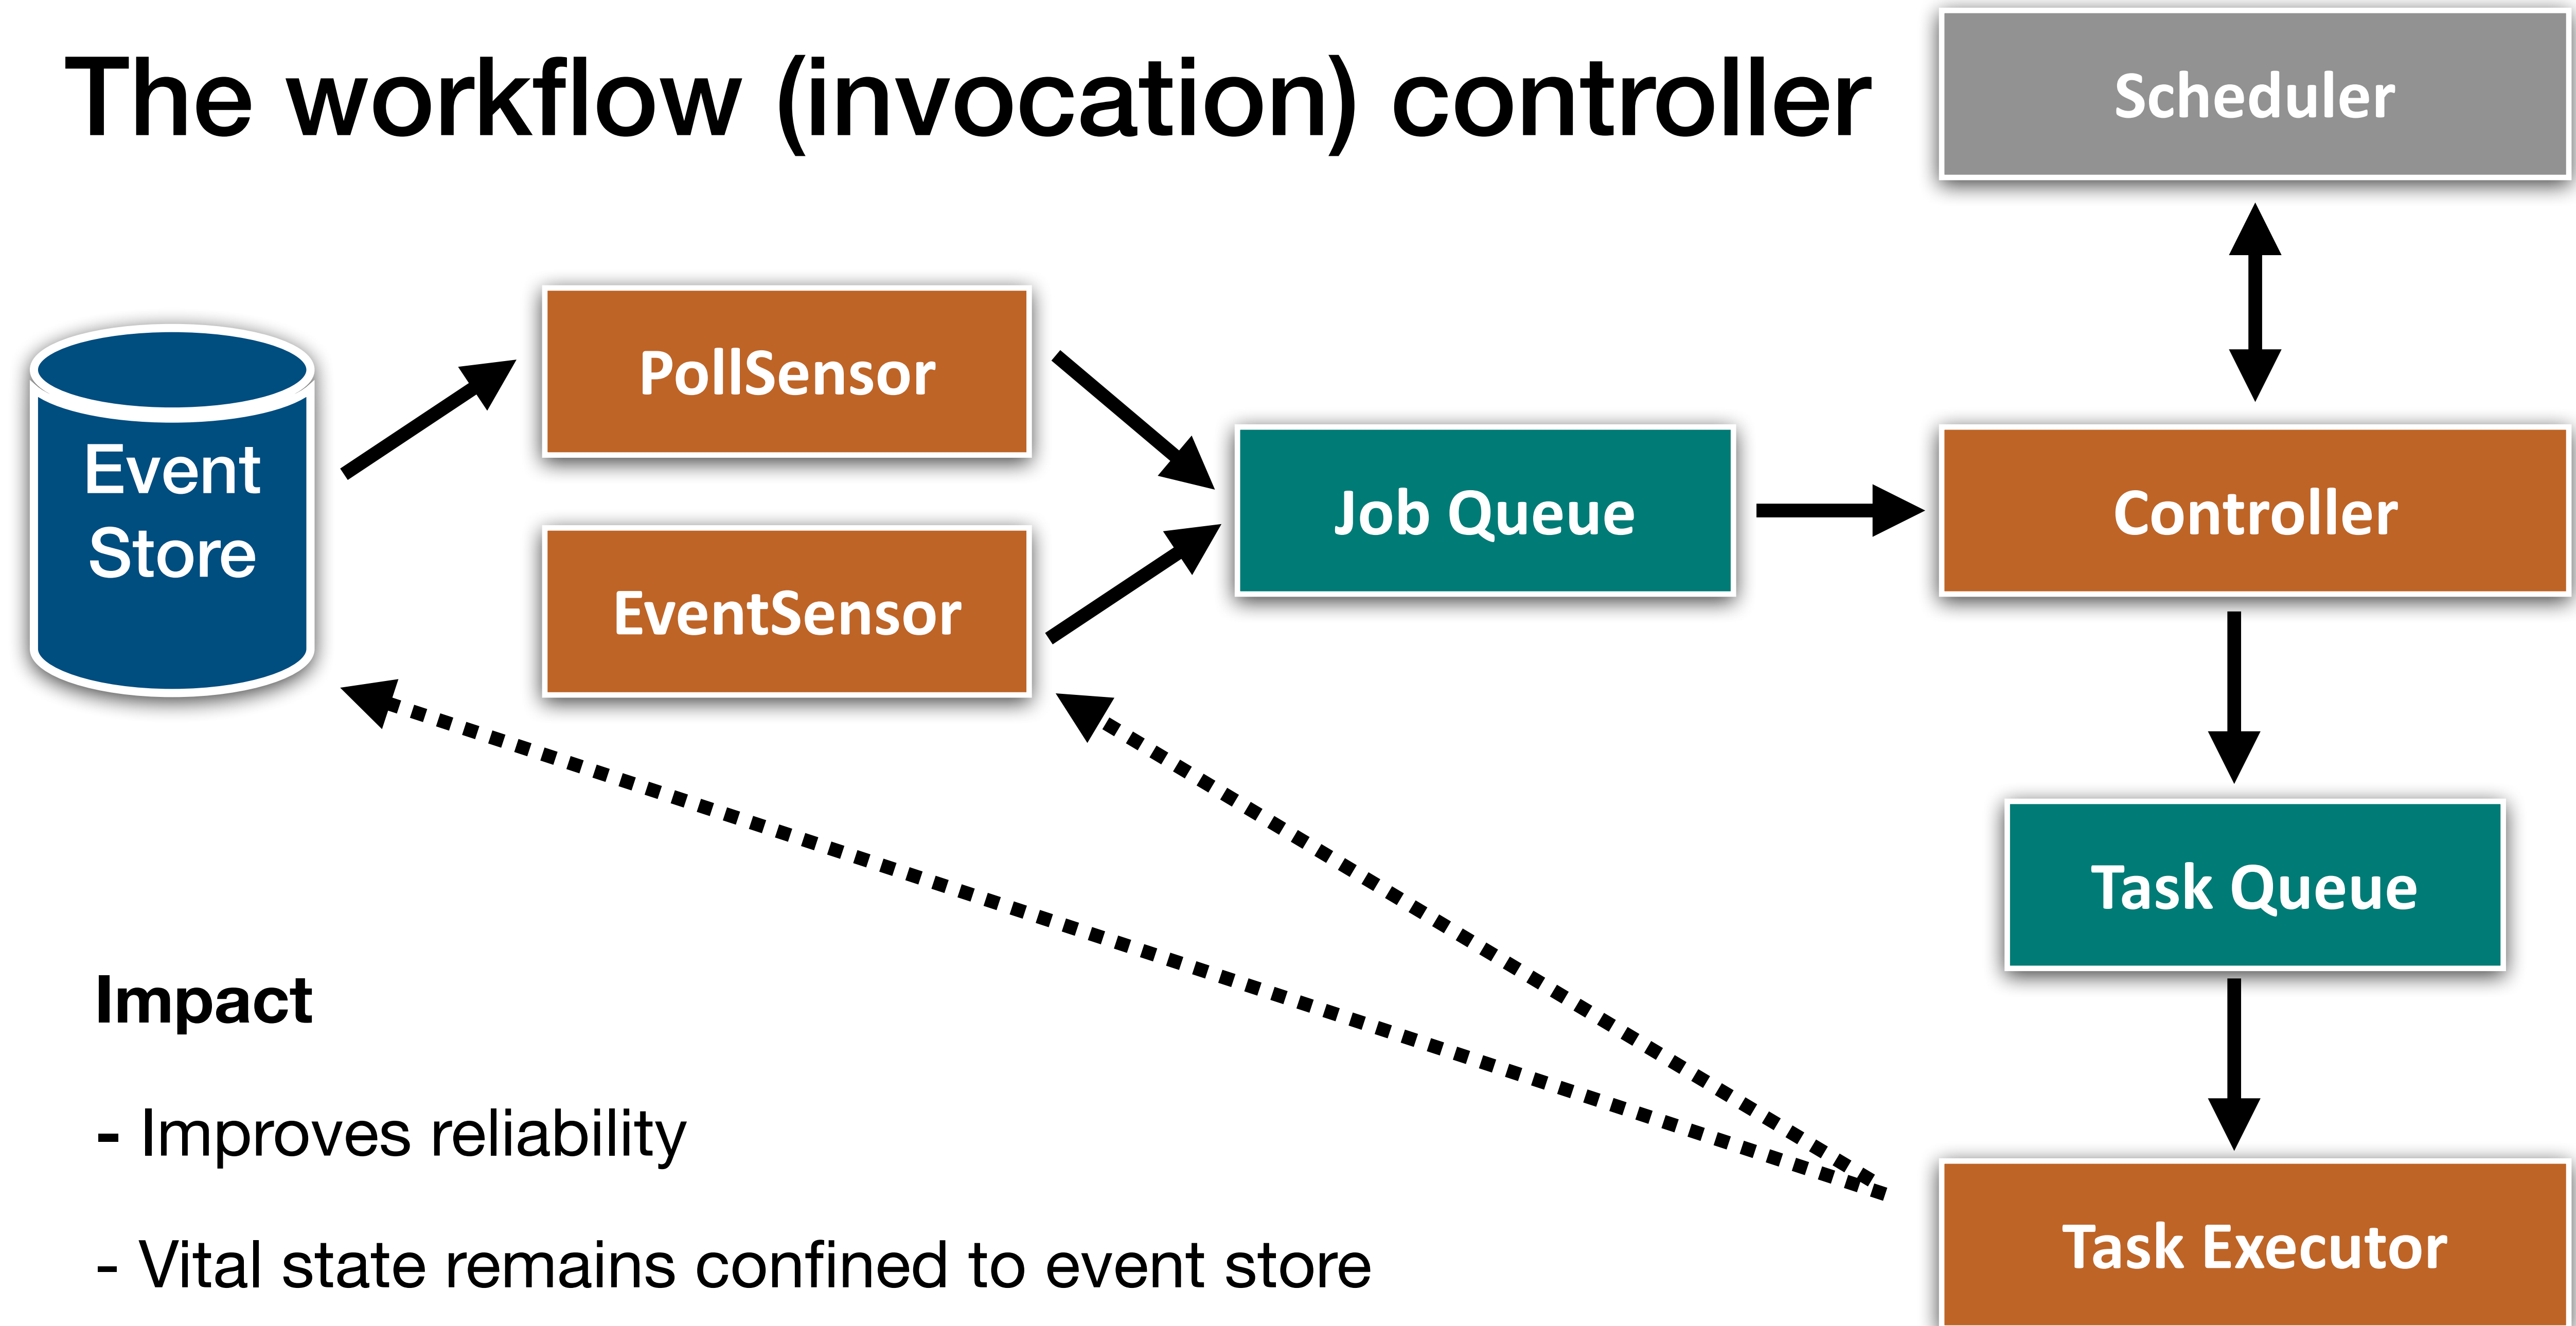
- Retain the time dimension of the data
- Simple, append-only persistence model

## Event store implementations

- In-memory
- NATS streaming



# The workflow (invocation) controller



## Impact

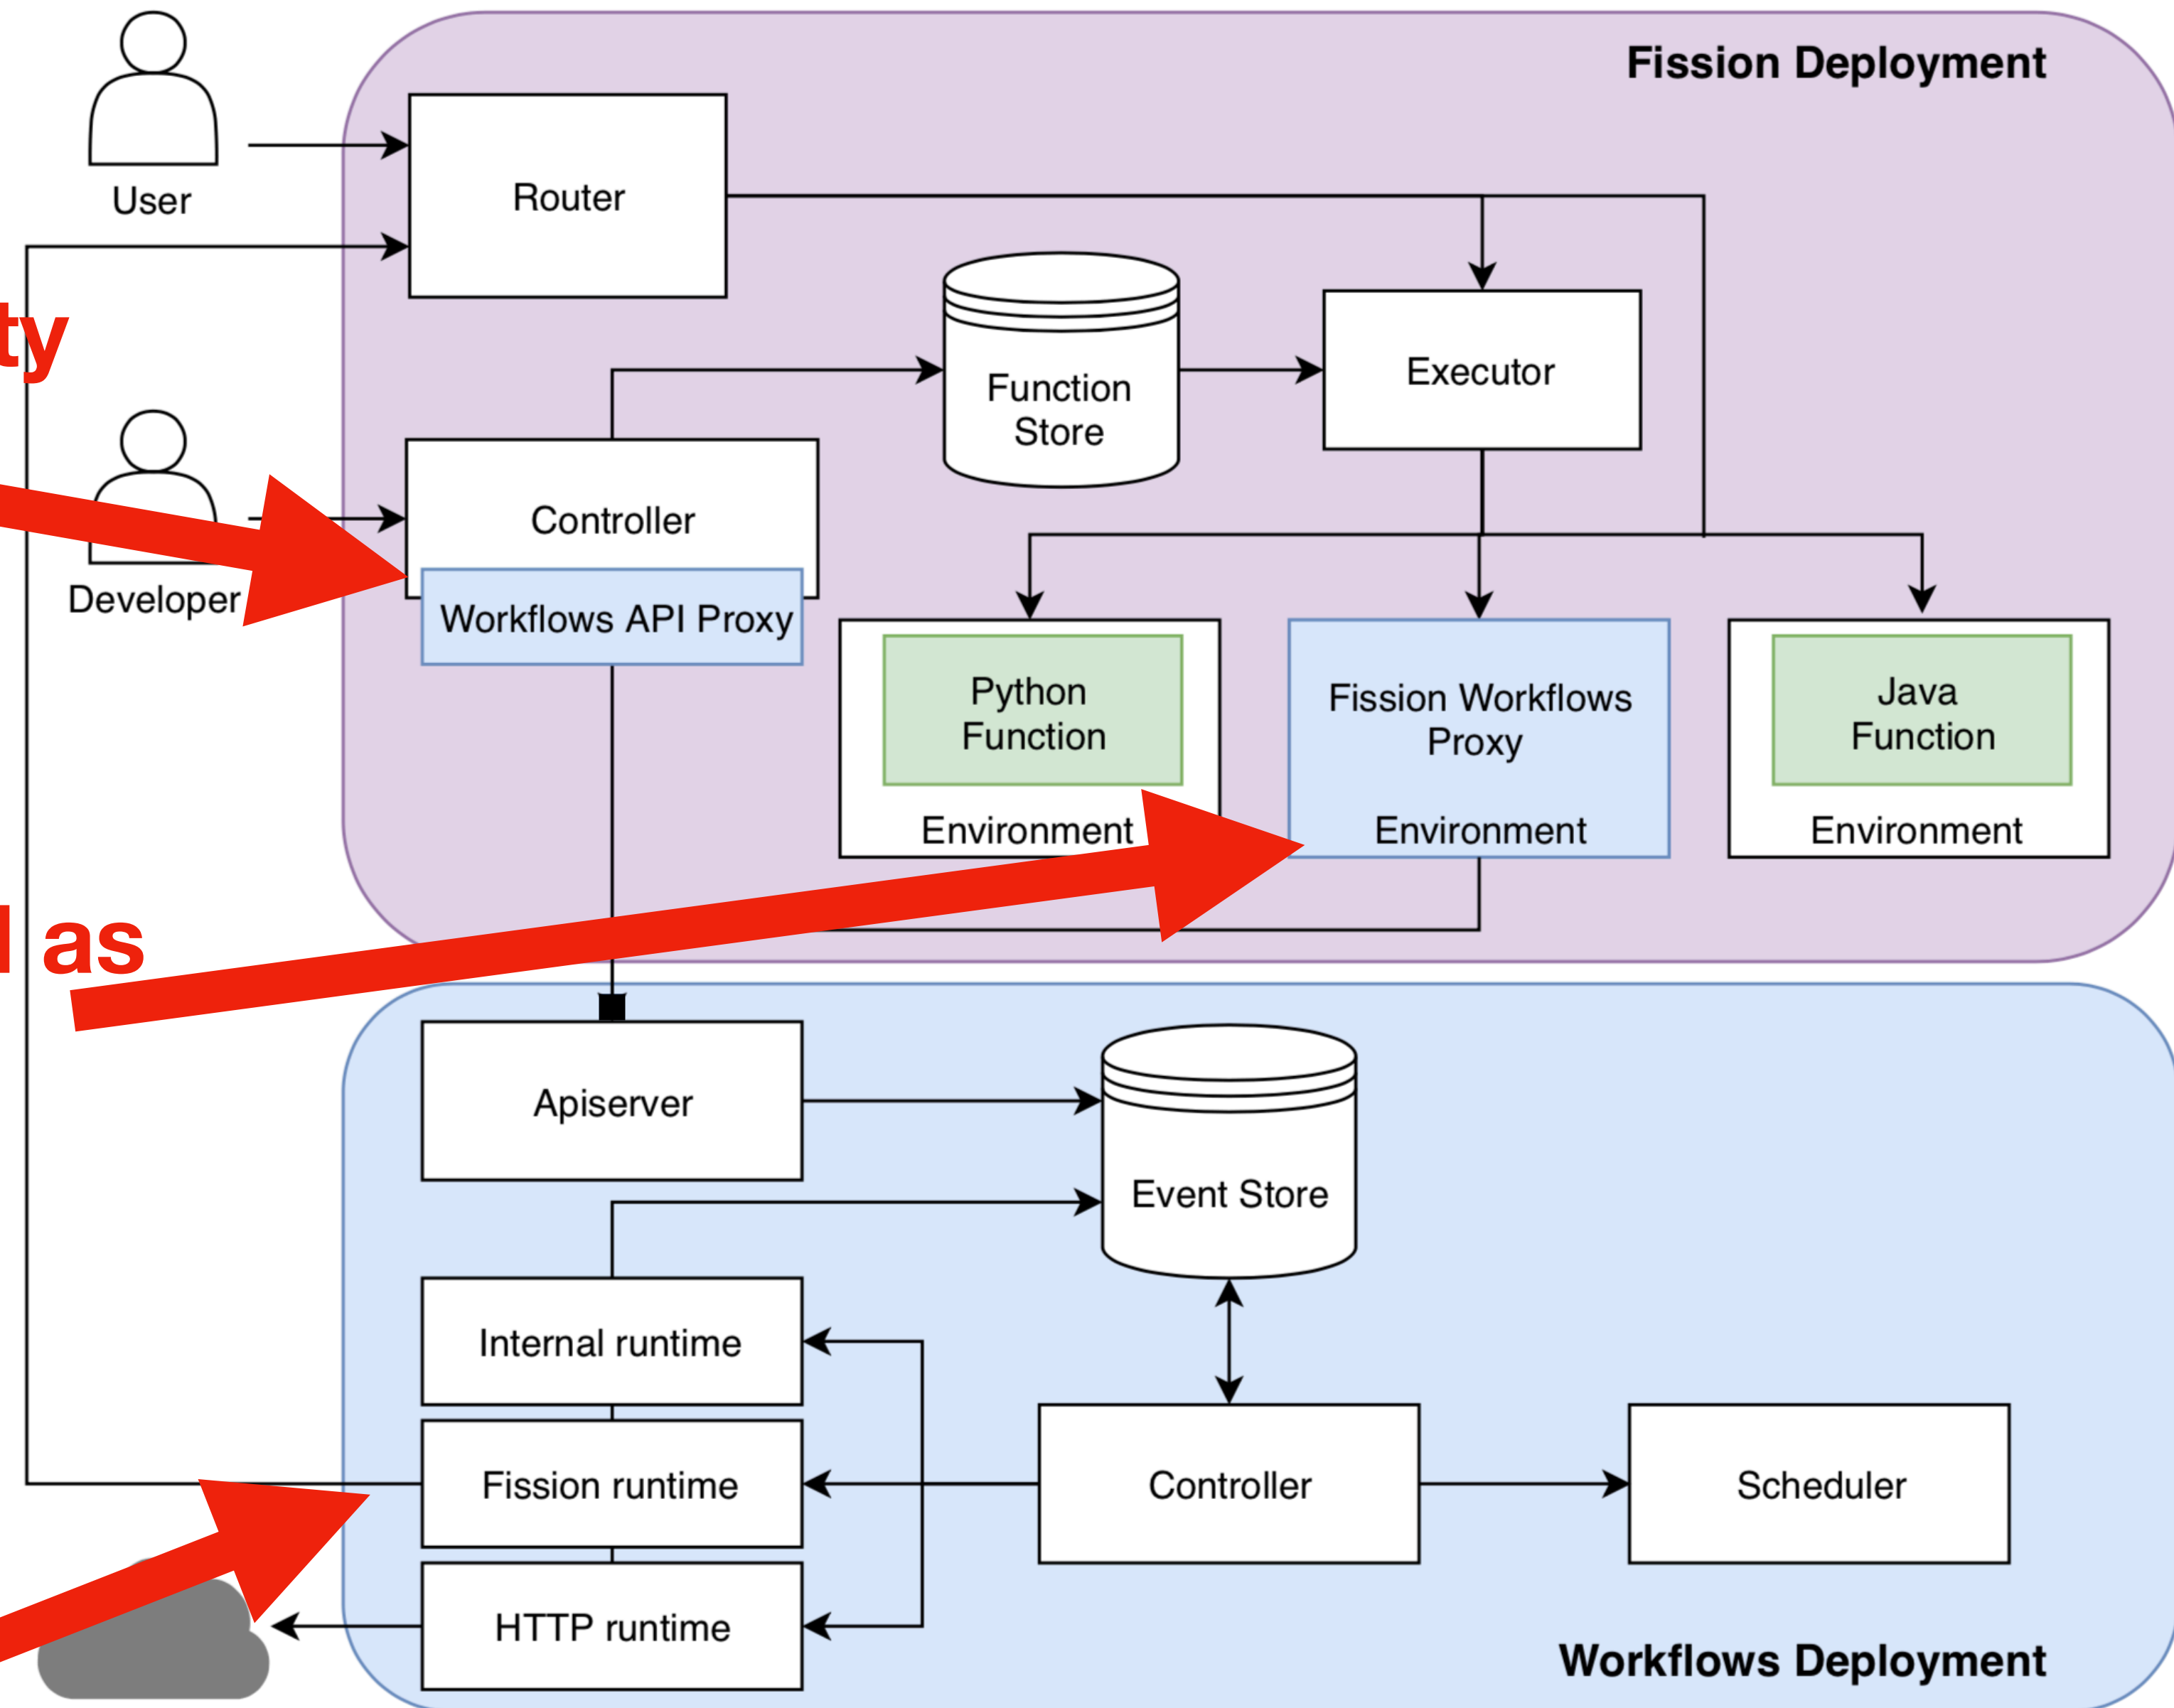
- Improves reliability
- Vital state remains confined to event store

# Integration with the FaaS platform

**FaaS can support workflow-specific functionality (optional)**

**Workflows can be executed as any serverless function**

**Fission Workflows uses the same user-level FaaS API**






# Workflow scheduling

**Motivation:** use workflows to mitigate cold starts of serverless functions.


horizon policy




**Prewarming:** predictively deploying serverless functions based on expected demand.

# Workflow scheduling


## prewarm-all policy




Prewarm all non-active tasks in the workflow




**Not started**



**Started**



**Finished**





**Prewarmed**




# Workflow scheduling


## prewarm-horizon policy




**Not started**



**Started**



**Finished**



**Prewarmed**

Prewarm all tasks that are “up next” in the workflow

# Experimental evaluation of the prototype

RQ4: How to evaluate systems for function composition experimentally?

**Goal:** Evaluate the workflow system prototype based on the following metrics:

- Reliability
- Performance
- Cost

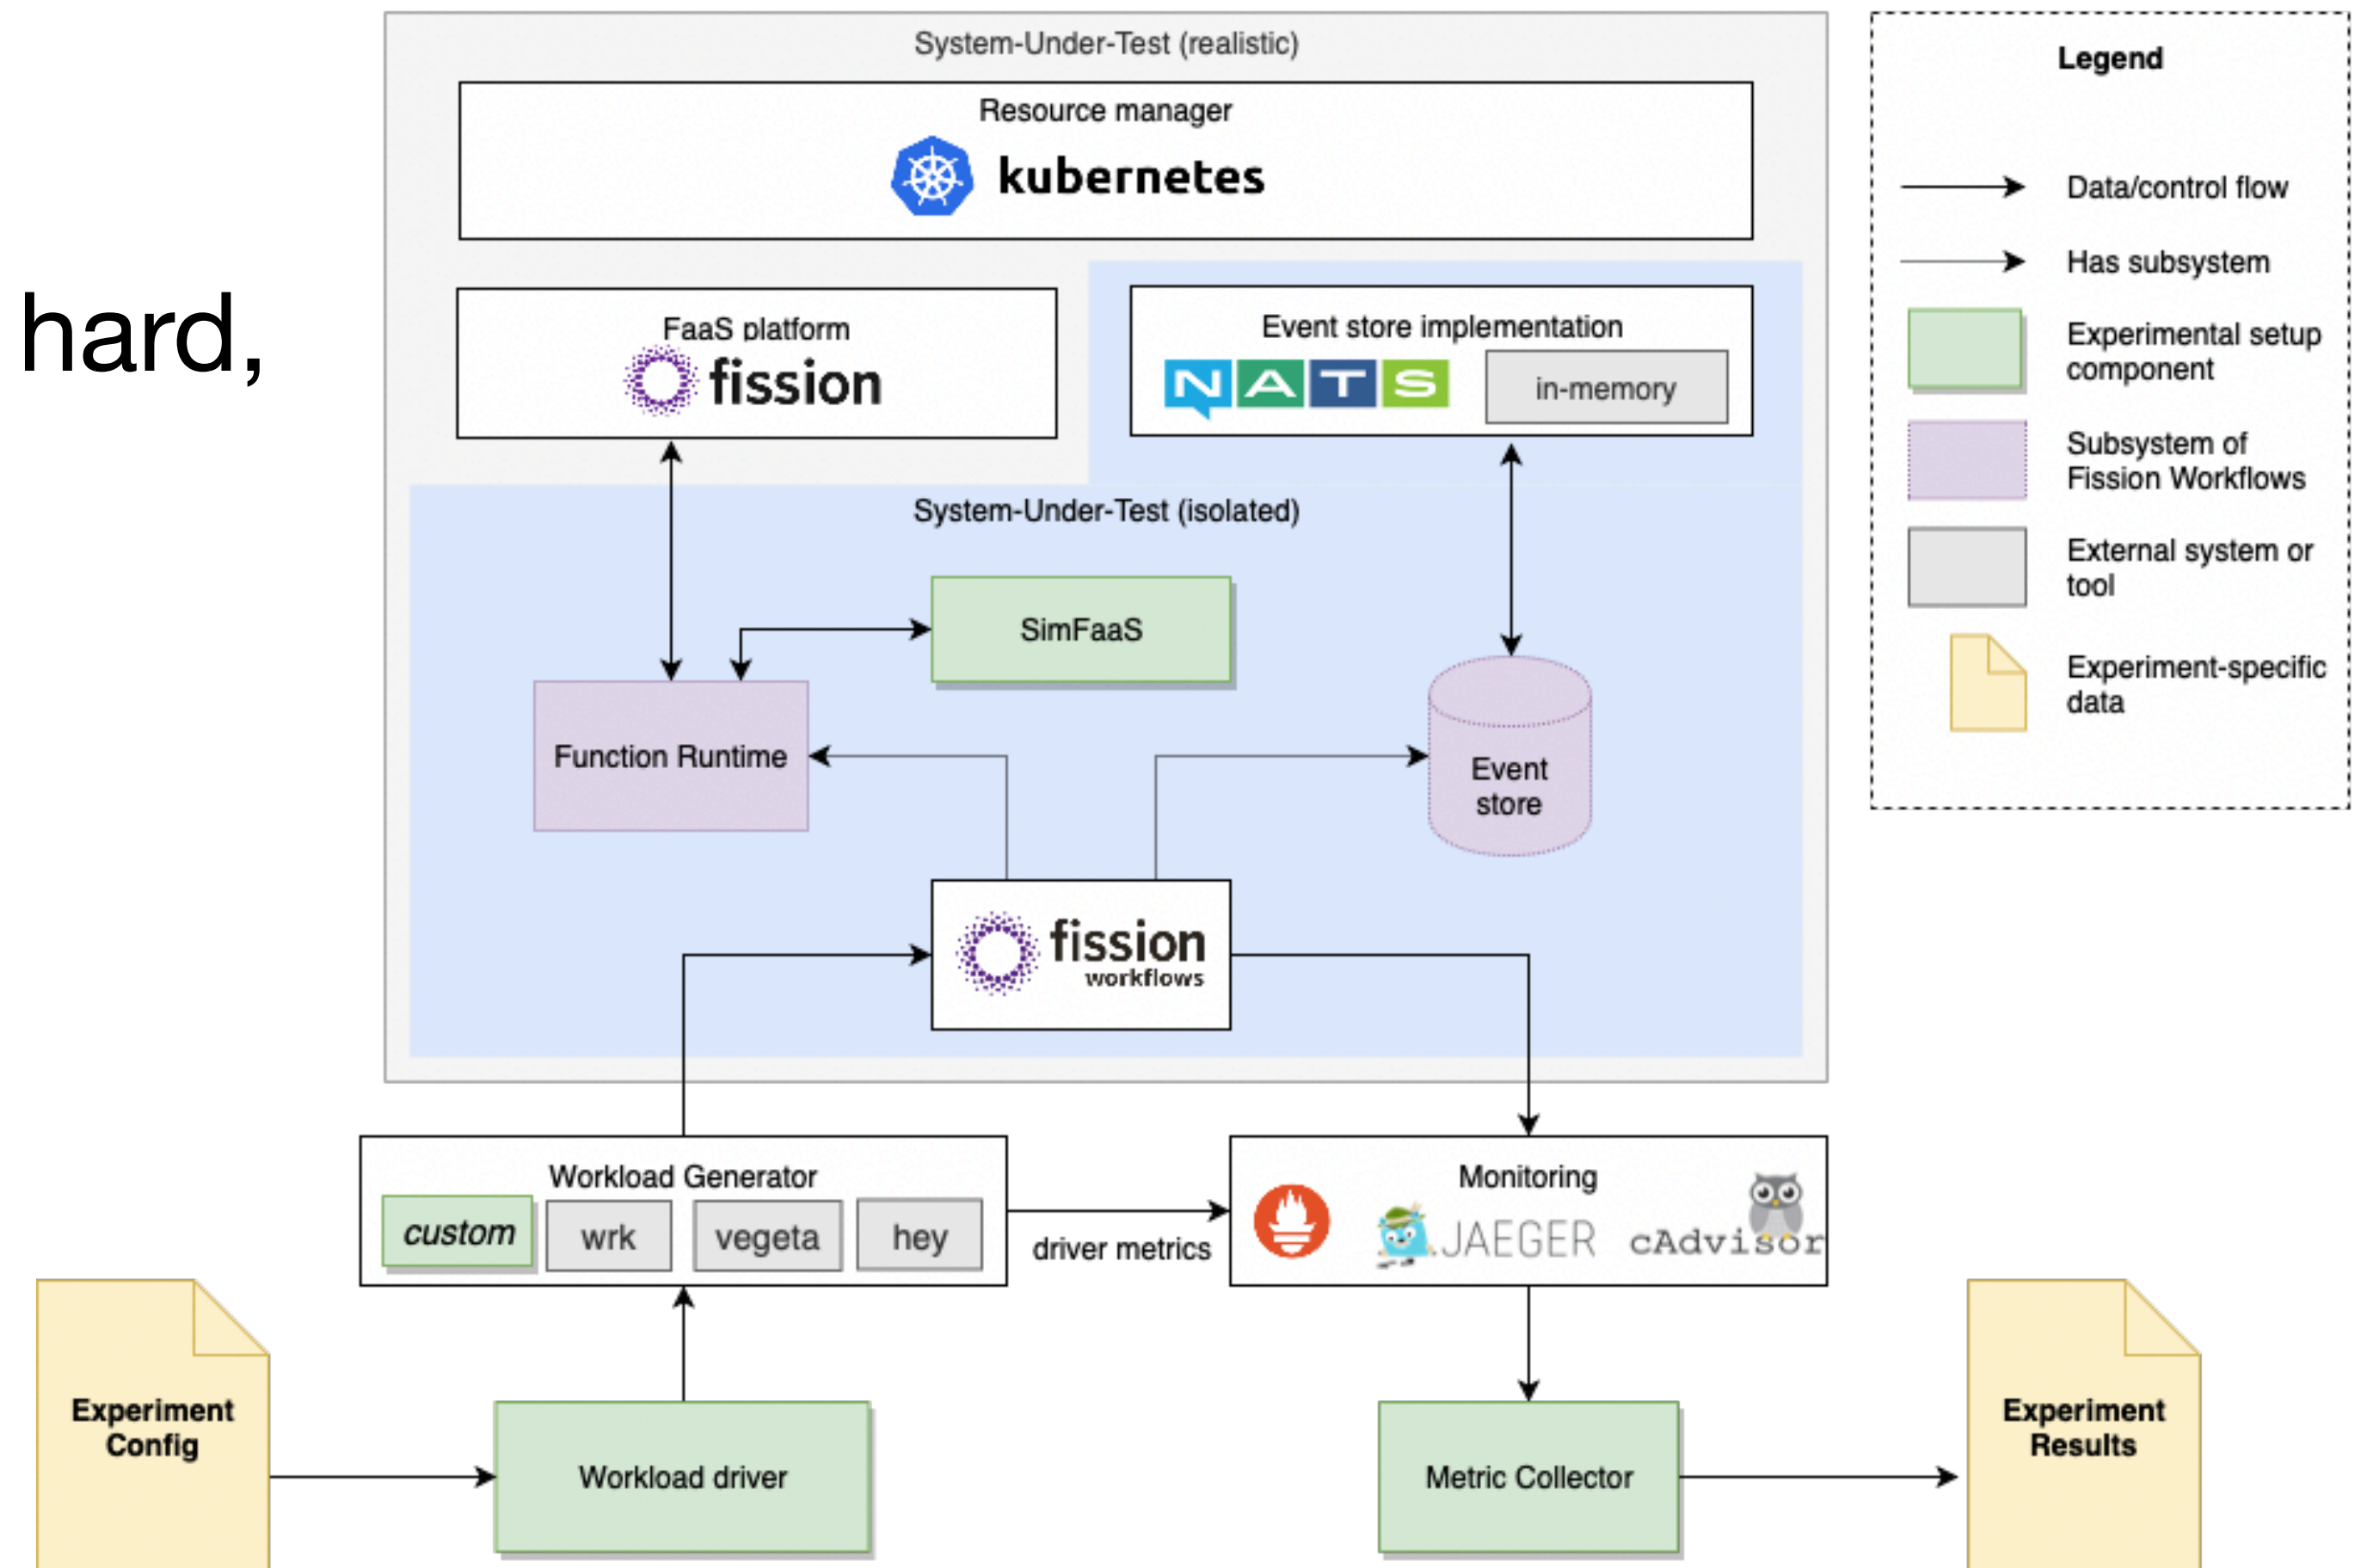


# Experimental setup

Developing distributed systems is hard,  
testing them is even harder.

Two types of experiments


- **Isolated** experiments
- **Realistic** experiments





# SimFaaS: a FaaS emulator

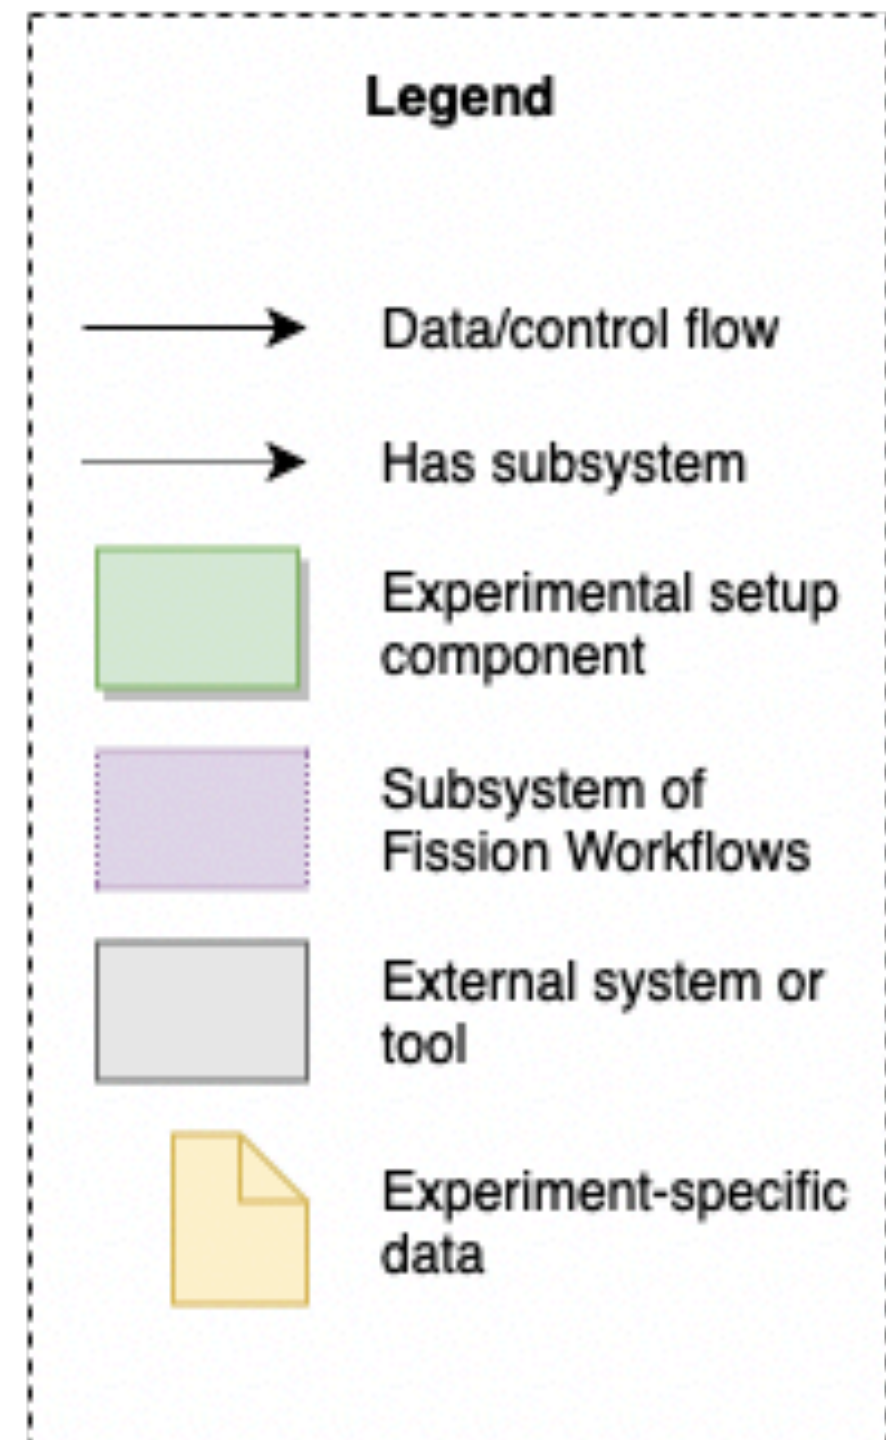
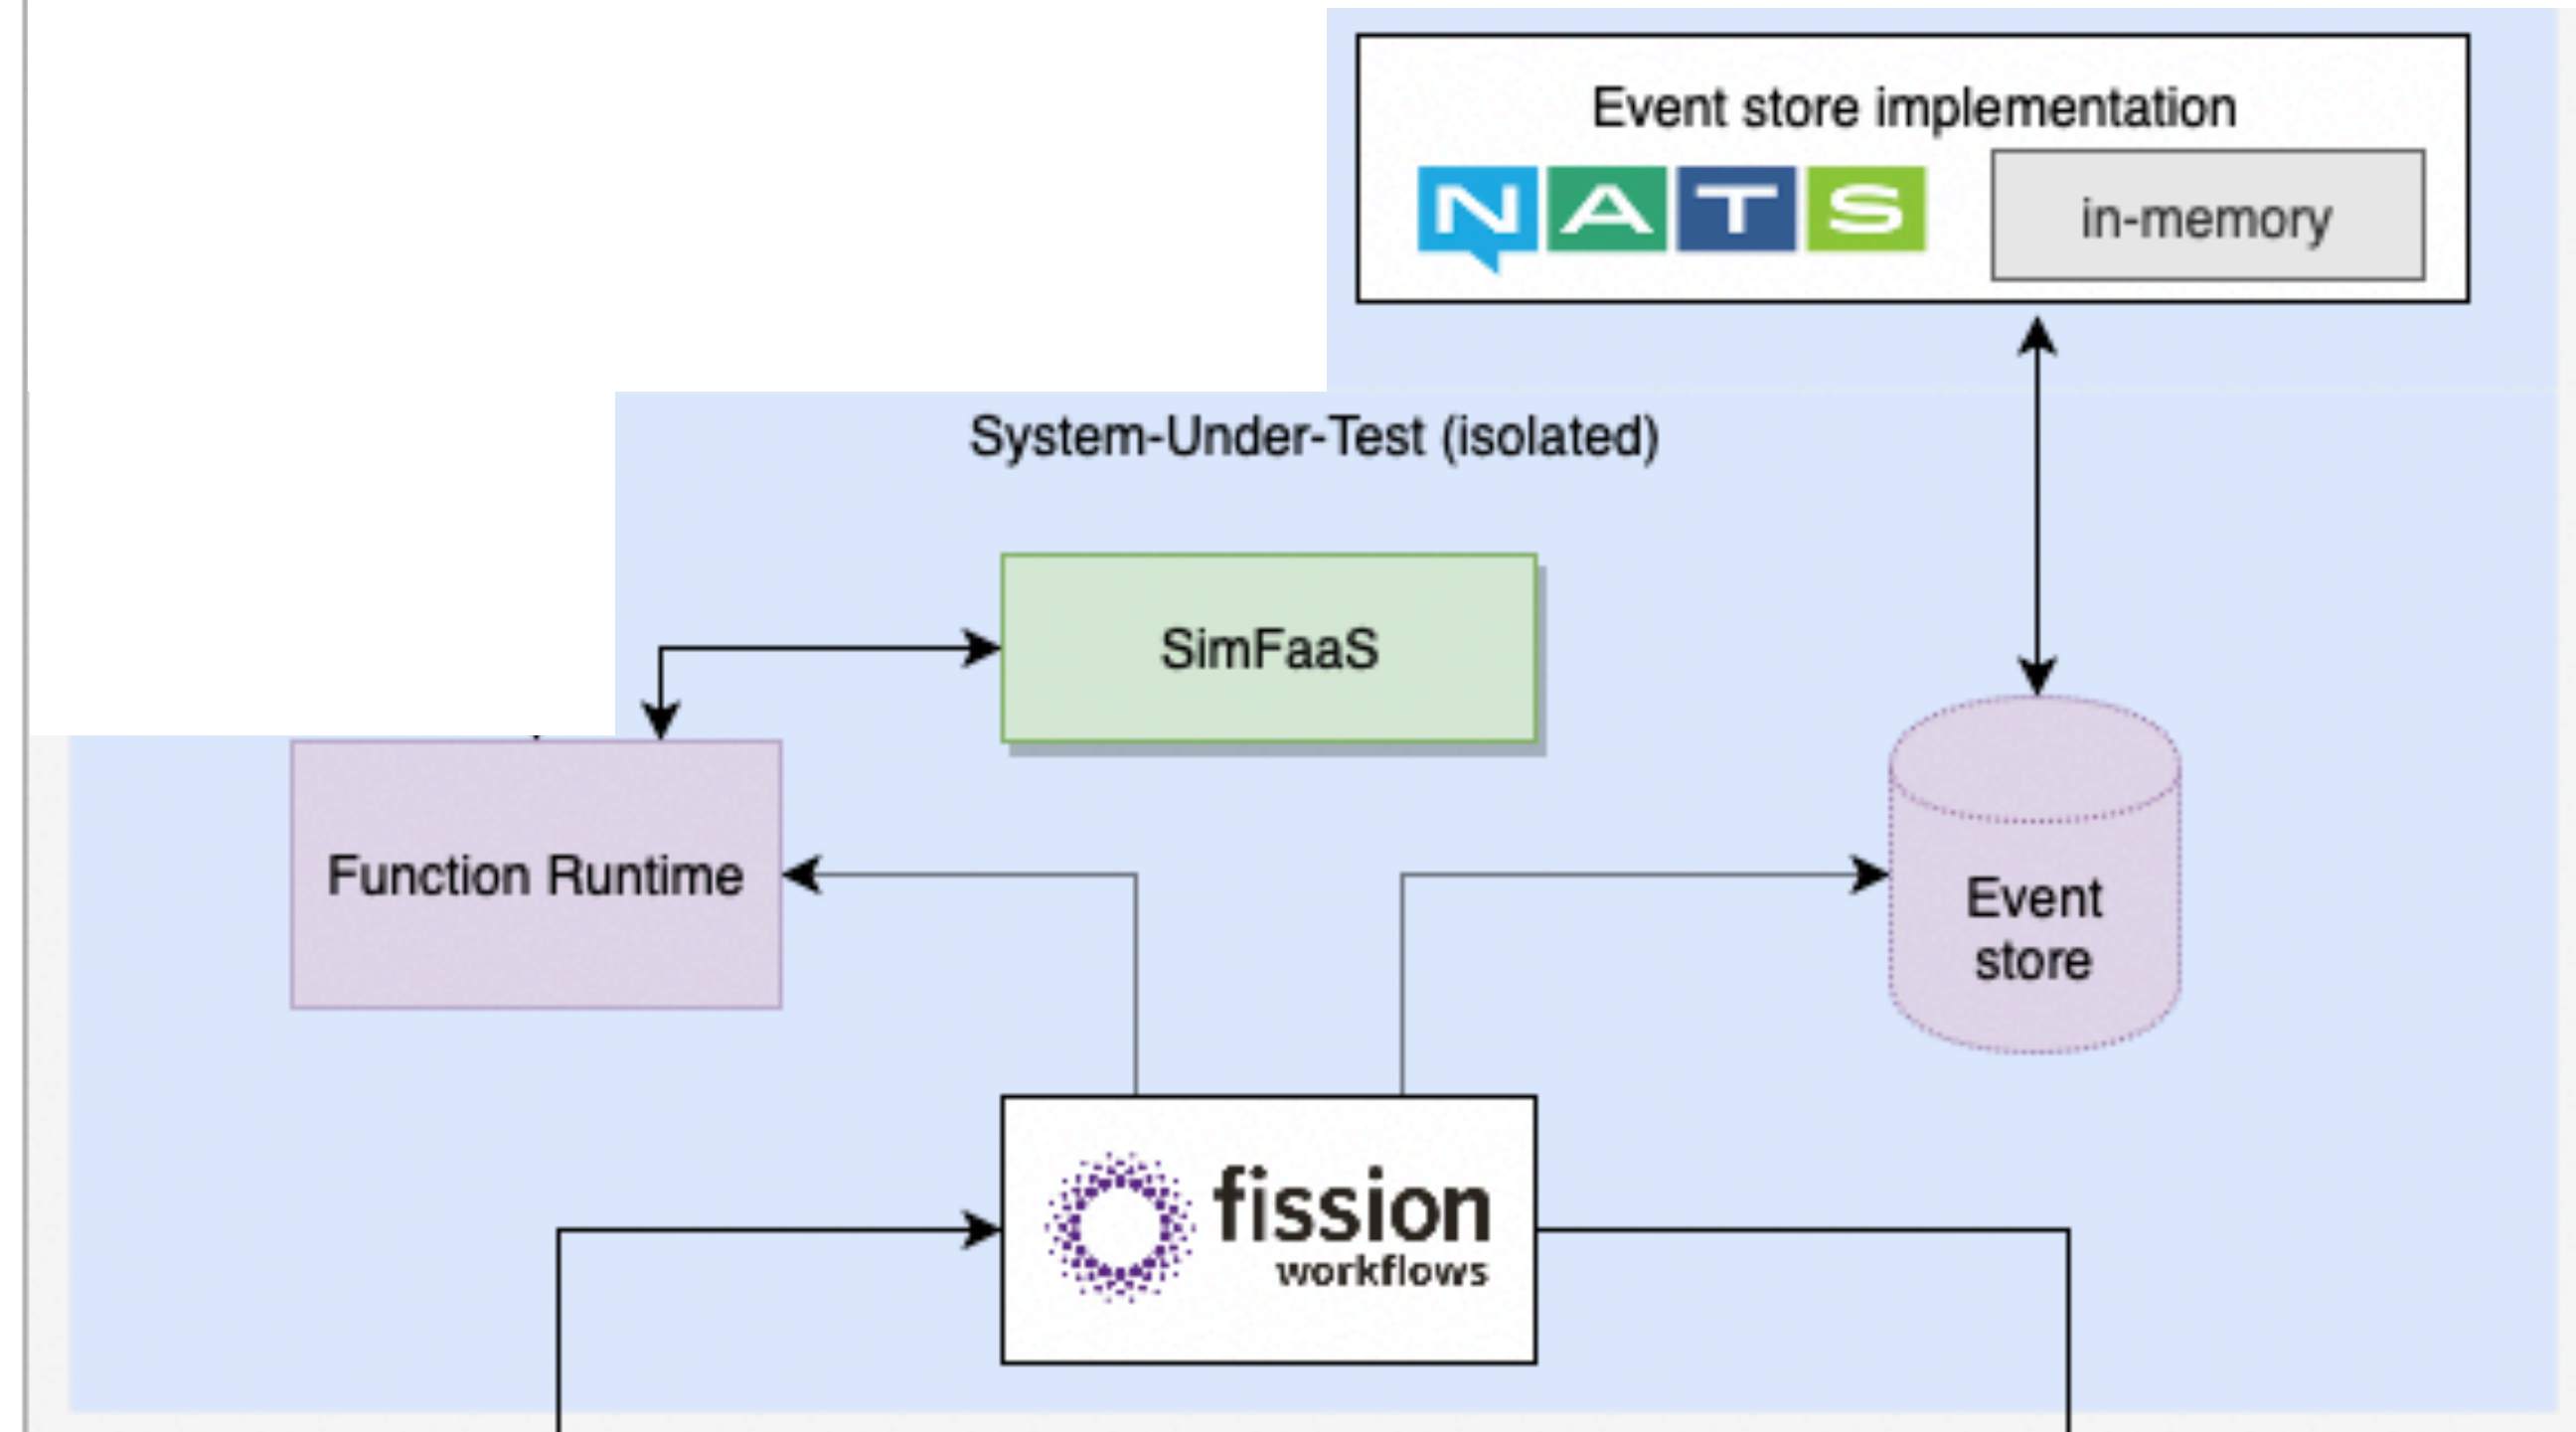
**Goal:** versatile FaaS emulation, which scales (with negligible performance overhead) well beyond the workloads in the other experiments.



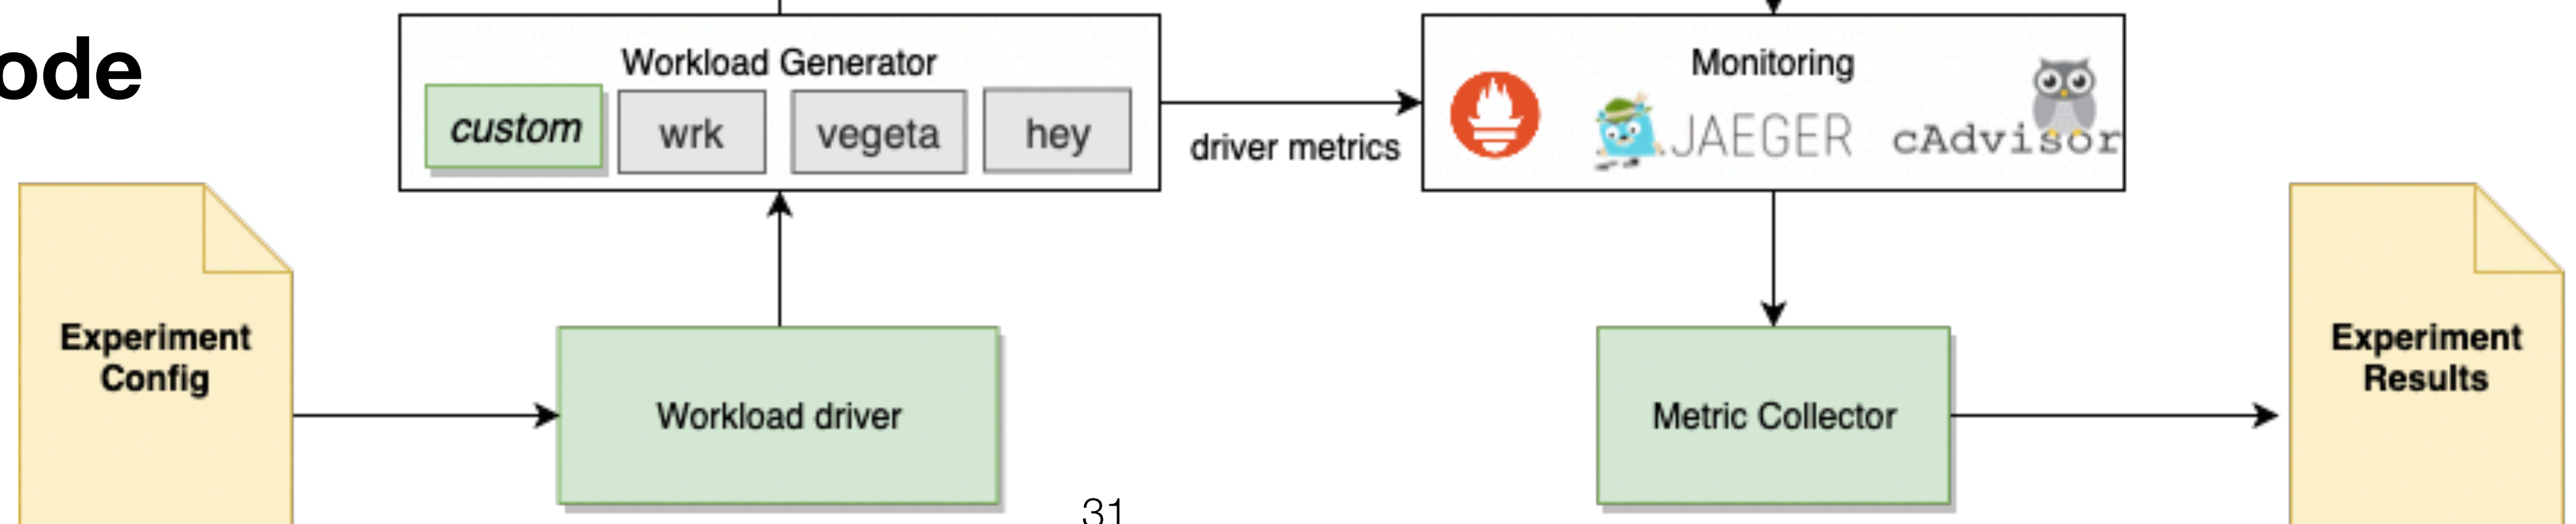


# Isolated experiment setup

## Worker Node



## Driver Node



# Overview of the isolated experiments

## Fault tolerance experiments

- Unavailable event store
- Unavailable FaaS runtime
- **Fail-stop crashes of Fission Workflows with different configurations**

## Scalability experiments

- Event store implementation performance overhead
- Workflow submission
- Workflow throughput
- Workflow parallelism
- Workflow length

## Scheduling experiments

# Fault-tolerance experiments


**Goal:** Evaluate the fault-tolerance of the prototype under different failure scenarios.

Scenarios evaluated:

- Unavailable event store
- Unavailable FaaS runtime
- **Fail-stop crashes of Fission Workflows**



# Fault-tolerance: crash of the controller



# Scheduling experiments


**Goal:** Evaluate the performance and resource consumption of the (prewarm-focused) scheduling policies:

- horizon
- prewarm-horizon
- prewarm-all

**Approach:** Use SimFaaS to emulate and control cold starts and resource usage.




# Performance of scheduling policies





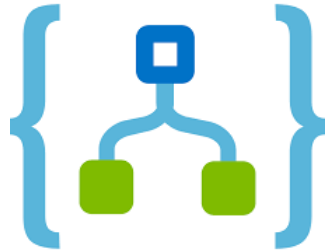
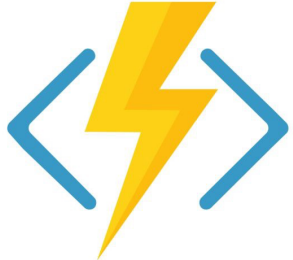






# Resource consumption of scheduling policies



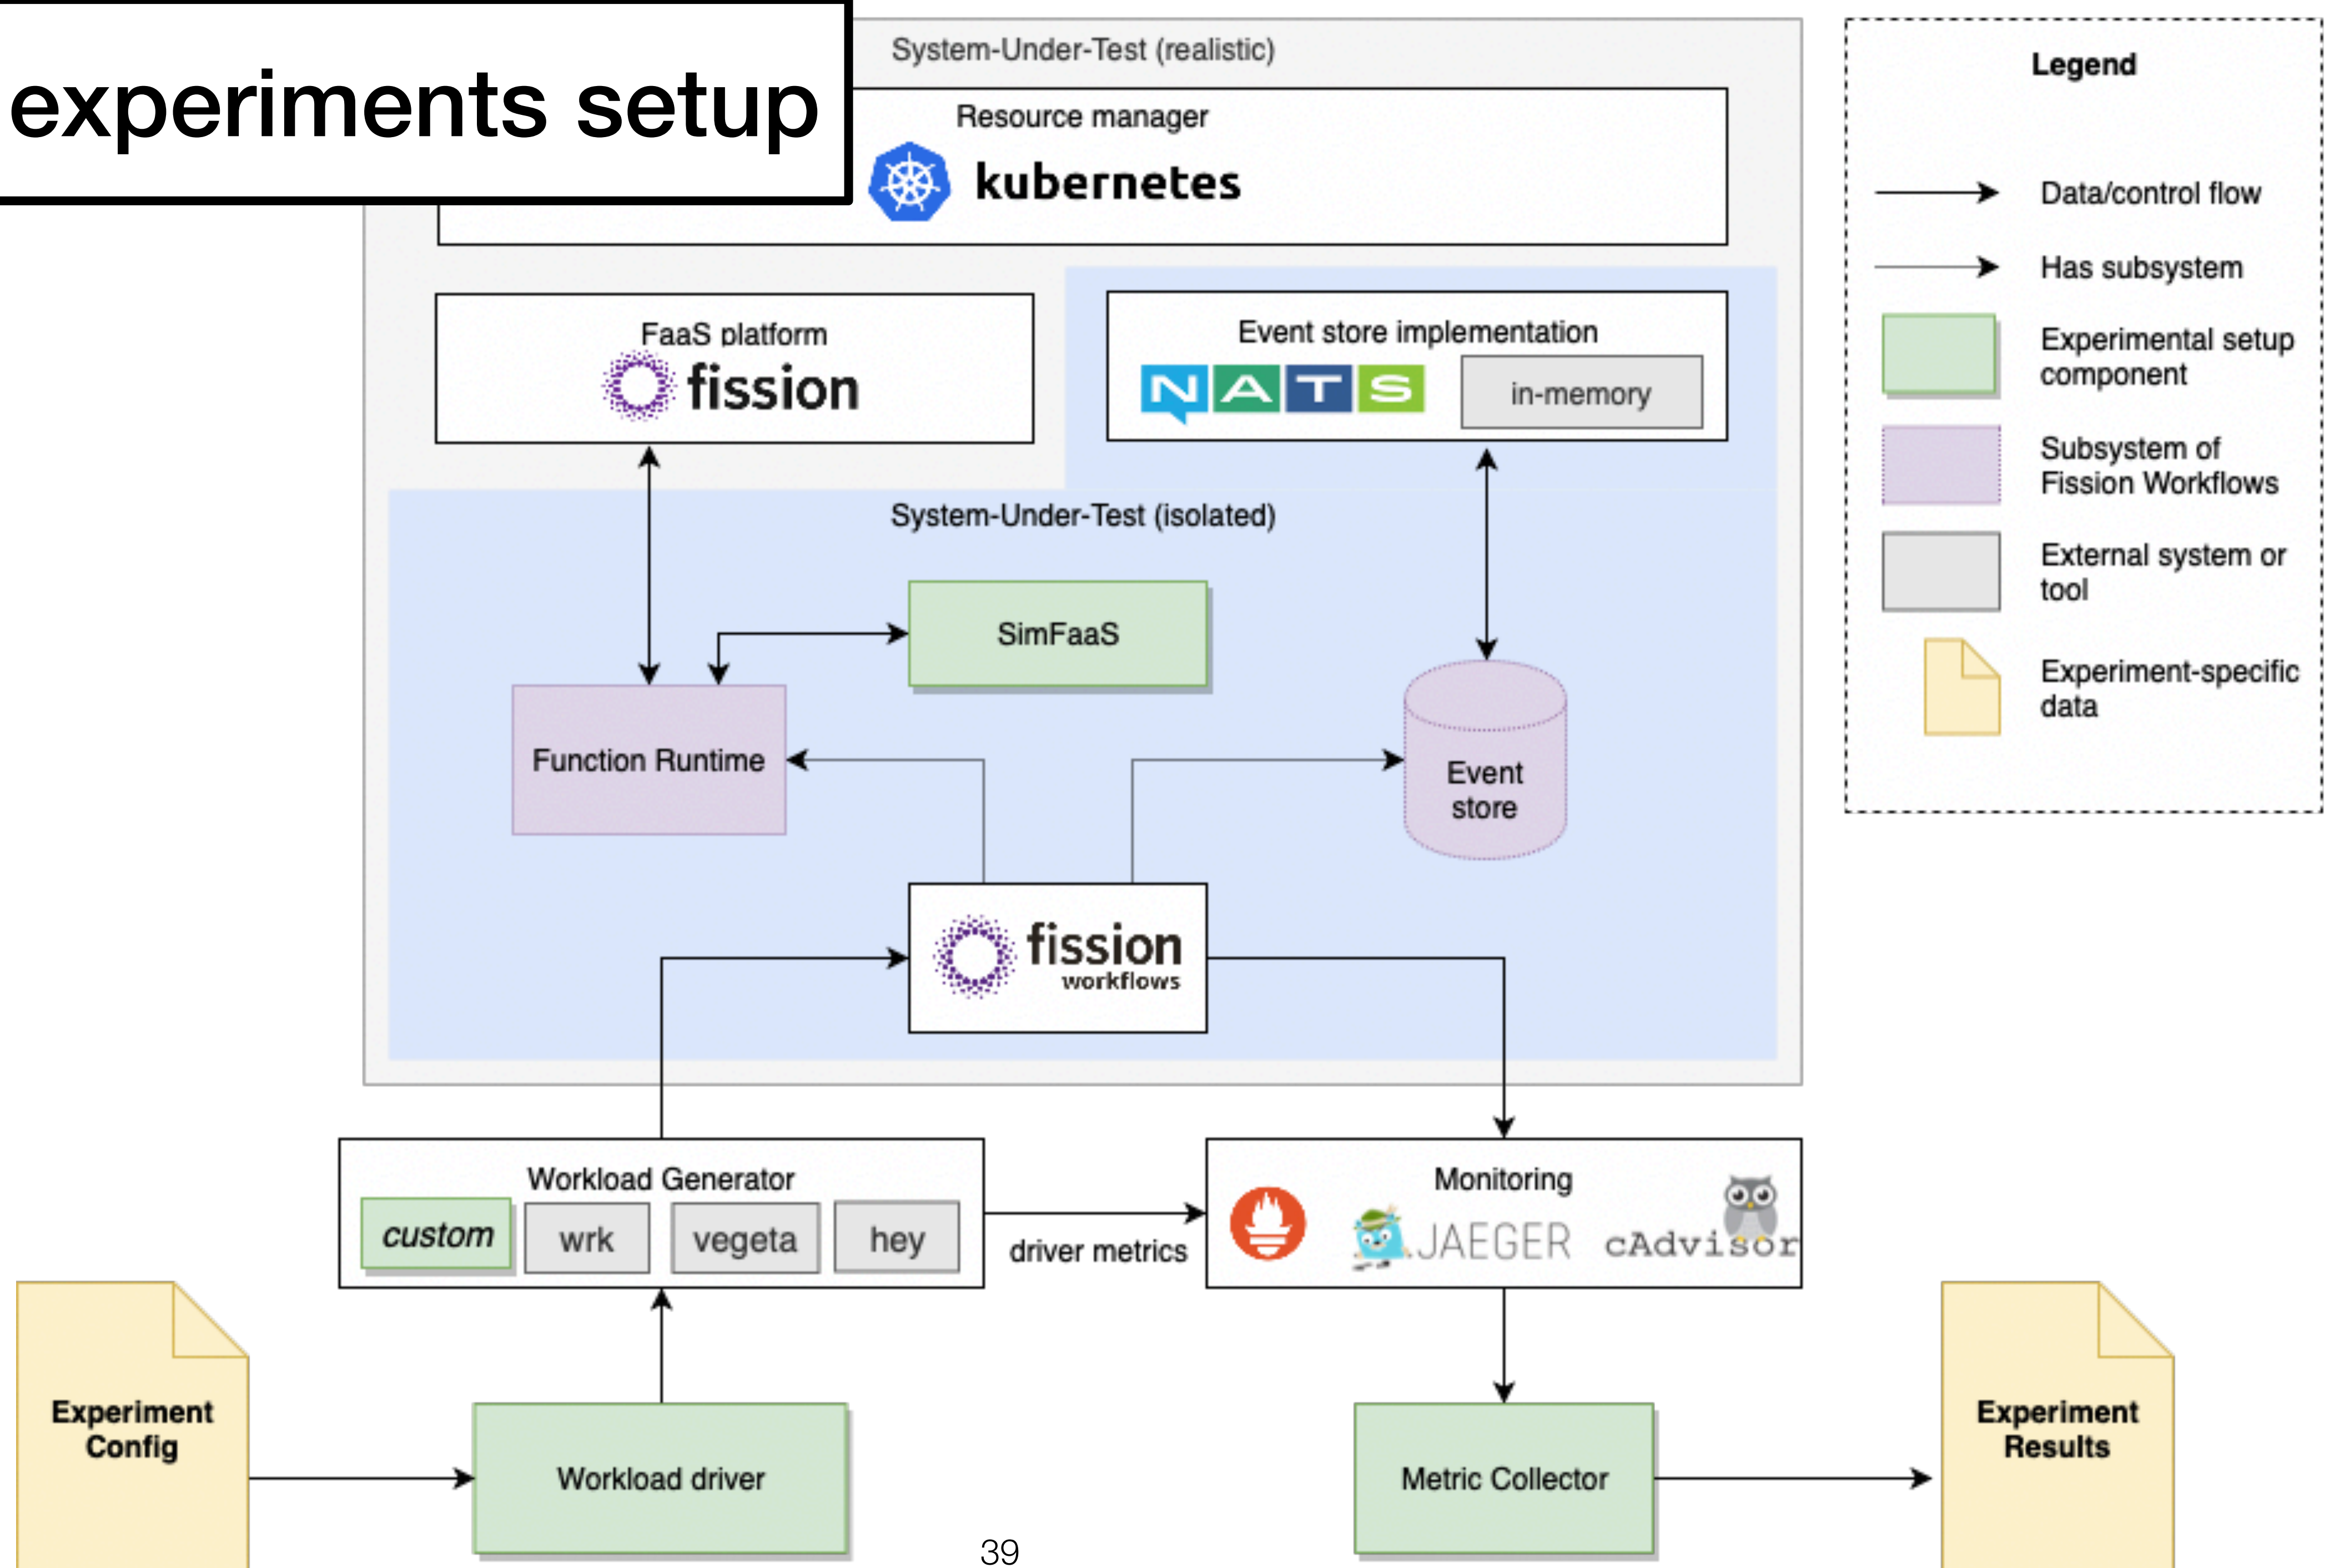
# Realistic experiments

**Goal:** evaluate the performance and cost of Fission Workflows under realistic circumstances, comparing it to the state-of-the-art.

| Provider            | Workflow system  | FaaS platform   |
|---------------------|--|---|
| Google Cloud        |  Google Cloud Composer        |  Google Cloud Functions |
| Microsoft Azure     |  Azure Logic Apps            |  Azure Functions       |
| AWS                 |  AWS Step Functions          |  AWS Lambda            |
| Fission (Workflows) |  <b>fission</b><br>workflows |  <b>fission</b>        |






# Realistic experiments setup






# Chronos workload




# Realistic experiment: FaaS overhead






# Realistic experiment: performance





# Realistic experiment: cost





An aerial photograph of a coastal landscape. In the foreground, a steep, rocky cliff covered in dense green vegetation drops down to a dark blue ocean. To the right, a wide, sandy beach stretches out, bordered by a road and more green hills. In the background, a series of mountains rise under a clear blue sky. A small bridge is visible on a road in the distance. The word "Process" is overlaid in the center of the image.

# Process



**~59k lines of code**

**~150 hours of experiments**  
**~105 GB of experiment data**

**~689 cups of coffee**



**10** publications

**4** as lead author

**60** citations





**Internship @ Platform9**





Login

- Startups
- Apps
- Gadgets
- Videos
- Audio
- Extra Crunch NEW
- Newsletters

- Events
- Advertise
- More

Search

- Apple
- Transportation
- Enterprise
- Cybersecurity 101

# Platform9's Fission Workflows makes it easier to write complex serverless applications

Frederic Lardinois @fredericl / 2 years ago Comment



En | 中文 | 日本 | Fr | Br  
1,274,004 May unique visitors

- Development
- Architecture & Design
- AI, ML and Data Engineering
- Culture & Methods
- DevOps
- Videos with Transcripts

NEW

FEATURED: Streaming Machine Learning Reactive Microservices Containers Observability Continuous Delivery QCon is

InfoQ Homepage > Articles > Four Techniques Serverless Platforms Use To Balance Performance And Cost

CLOUD

## Four Techniques Serverless Platforms Use to Balance Performance and Cost

LIKE 1 BOOKMARKS

FEB 13, 2019 • 25 MIN READ

by  
Erwin van Eyk

FOLLOW

reviewed by  
Richard Seroter

FOLLOW

### Key Takeaways

- The cost and performance models are two of the key drivers of the popularity of serverless and Function-as-a-Service (FaaS).
- Cold starts have gone down a lot, from multiple seconds to 100s of milliseconds, but there is still much space for improvement.
- There are various techniques that are being used to improve the performance of serverless functions, most of which focus on reducing or avoiding cold starts.
- These optimizations are not free; it is a trade-off between performance and cost. which depends on the requirements of your application.

### RELATED CONTENT

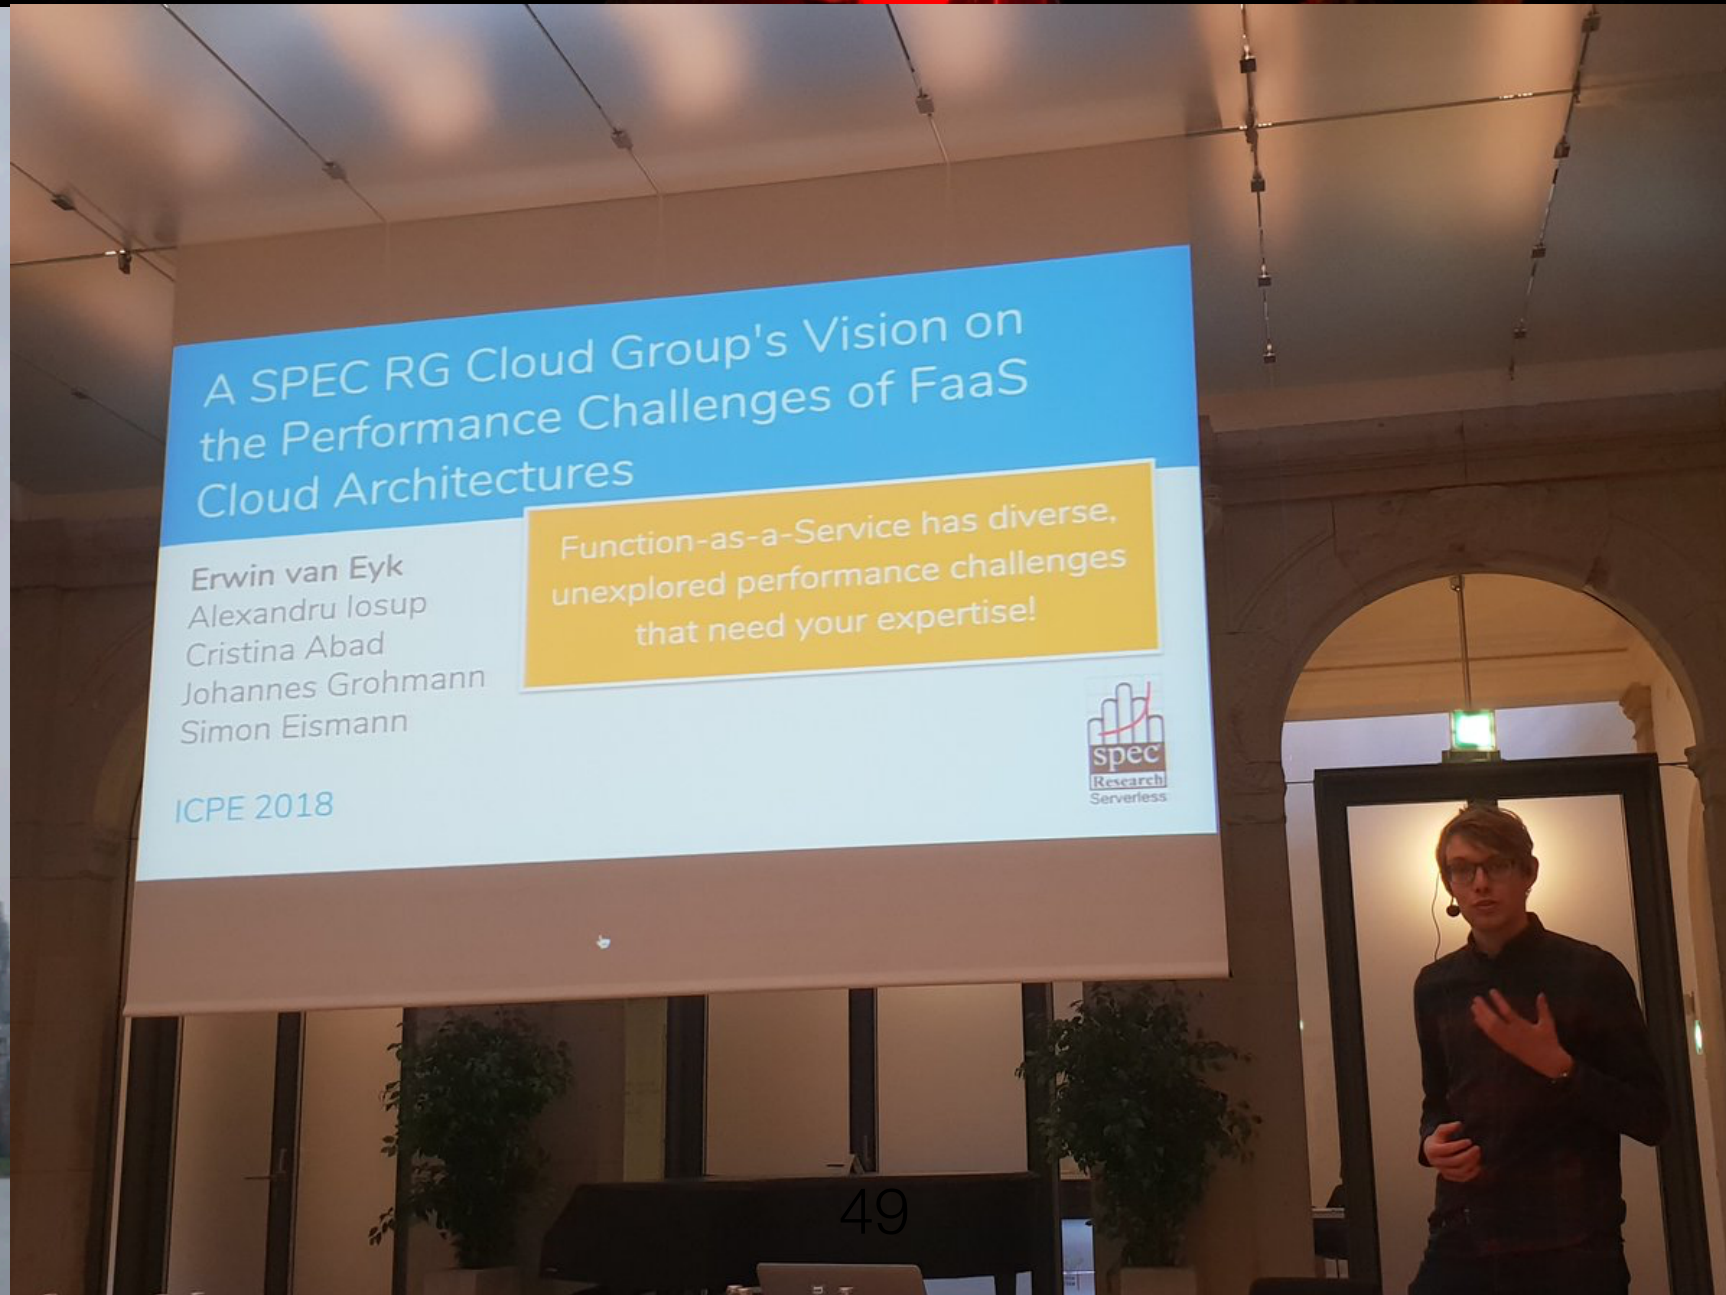
**Akamai IoT Edge Connect Serverless Edge Platform**  
JUN 15, 2019

**Cloudflare Releases Worker Value Store at the Edge**  
MAY 28, 2019

**Building Responsive Serverless, Event-Driven**  
JAN 05, 2019

Industry publications







# Conclusion

1. Workflows are key to enabling serverless function composition.
2. Fission Workflows demonstrates the possibilities of serverless workflows; and highlights the opportunities of prewarming.
3. The prototype is on-par performance-wise, and cheaper than the state-of-the-art.
4. Industry interest in the Fission Workflows product emphasises the need for serverless workflow systems.

**More research is needed in serverless computing!**



# Thanks!



**SimFaaS**

**Thesis + slides**

<https://github.com/fission>

<https://github.com/erwinvaneyk/simfaas>

<https://erwinvaneyk.nl/thesis>

**Erwin van Eyk**

Software Engineer @ Platform9 Systems  
Researcher @ AtLarge Research  
Co-chair @ SPEC CLOUD RG Serverless

@erwinvaneyk  
erwinvaneyk@gmail.com  
<https://erwinvaneyk.nl>





# Additional Slides







# Photo credits

- Empty datacenter: <https://www.datacenterknowledge.com/manage/wave-data-center-consolidation-different-first-one>
- Copyright of logos used (Google, AWS, Azure, SPEC RG CLOUD, TU Delft, Platform9, Python) belongs to the respective organizations
- Logo for SWL: <https://pixabay.com/vectors/hexagon-symbol-gui-internet-2307350/>








# Performance overhead of SimFaaS (1)






# Performance overhead of SimFaaS (2)






# Fault-tolerance: event store unavailability






# Fault-tolerance: FaaS runtime unavailability






# Fault-tolerance: crash of bundled workflow engine






# Fault-tolerance: crash of distributed workflow engine



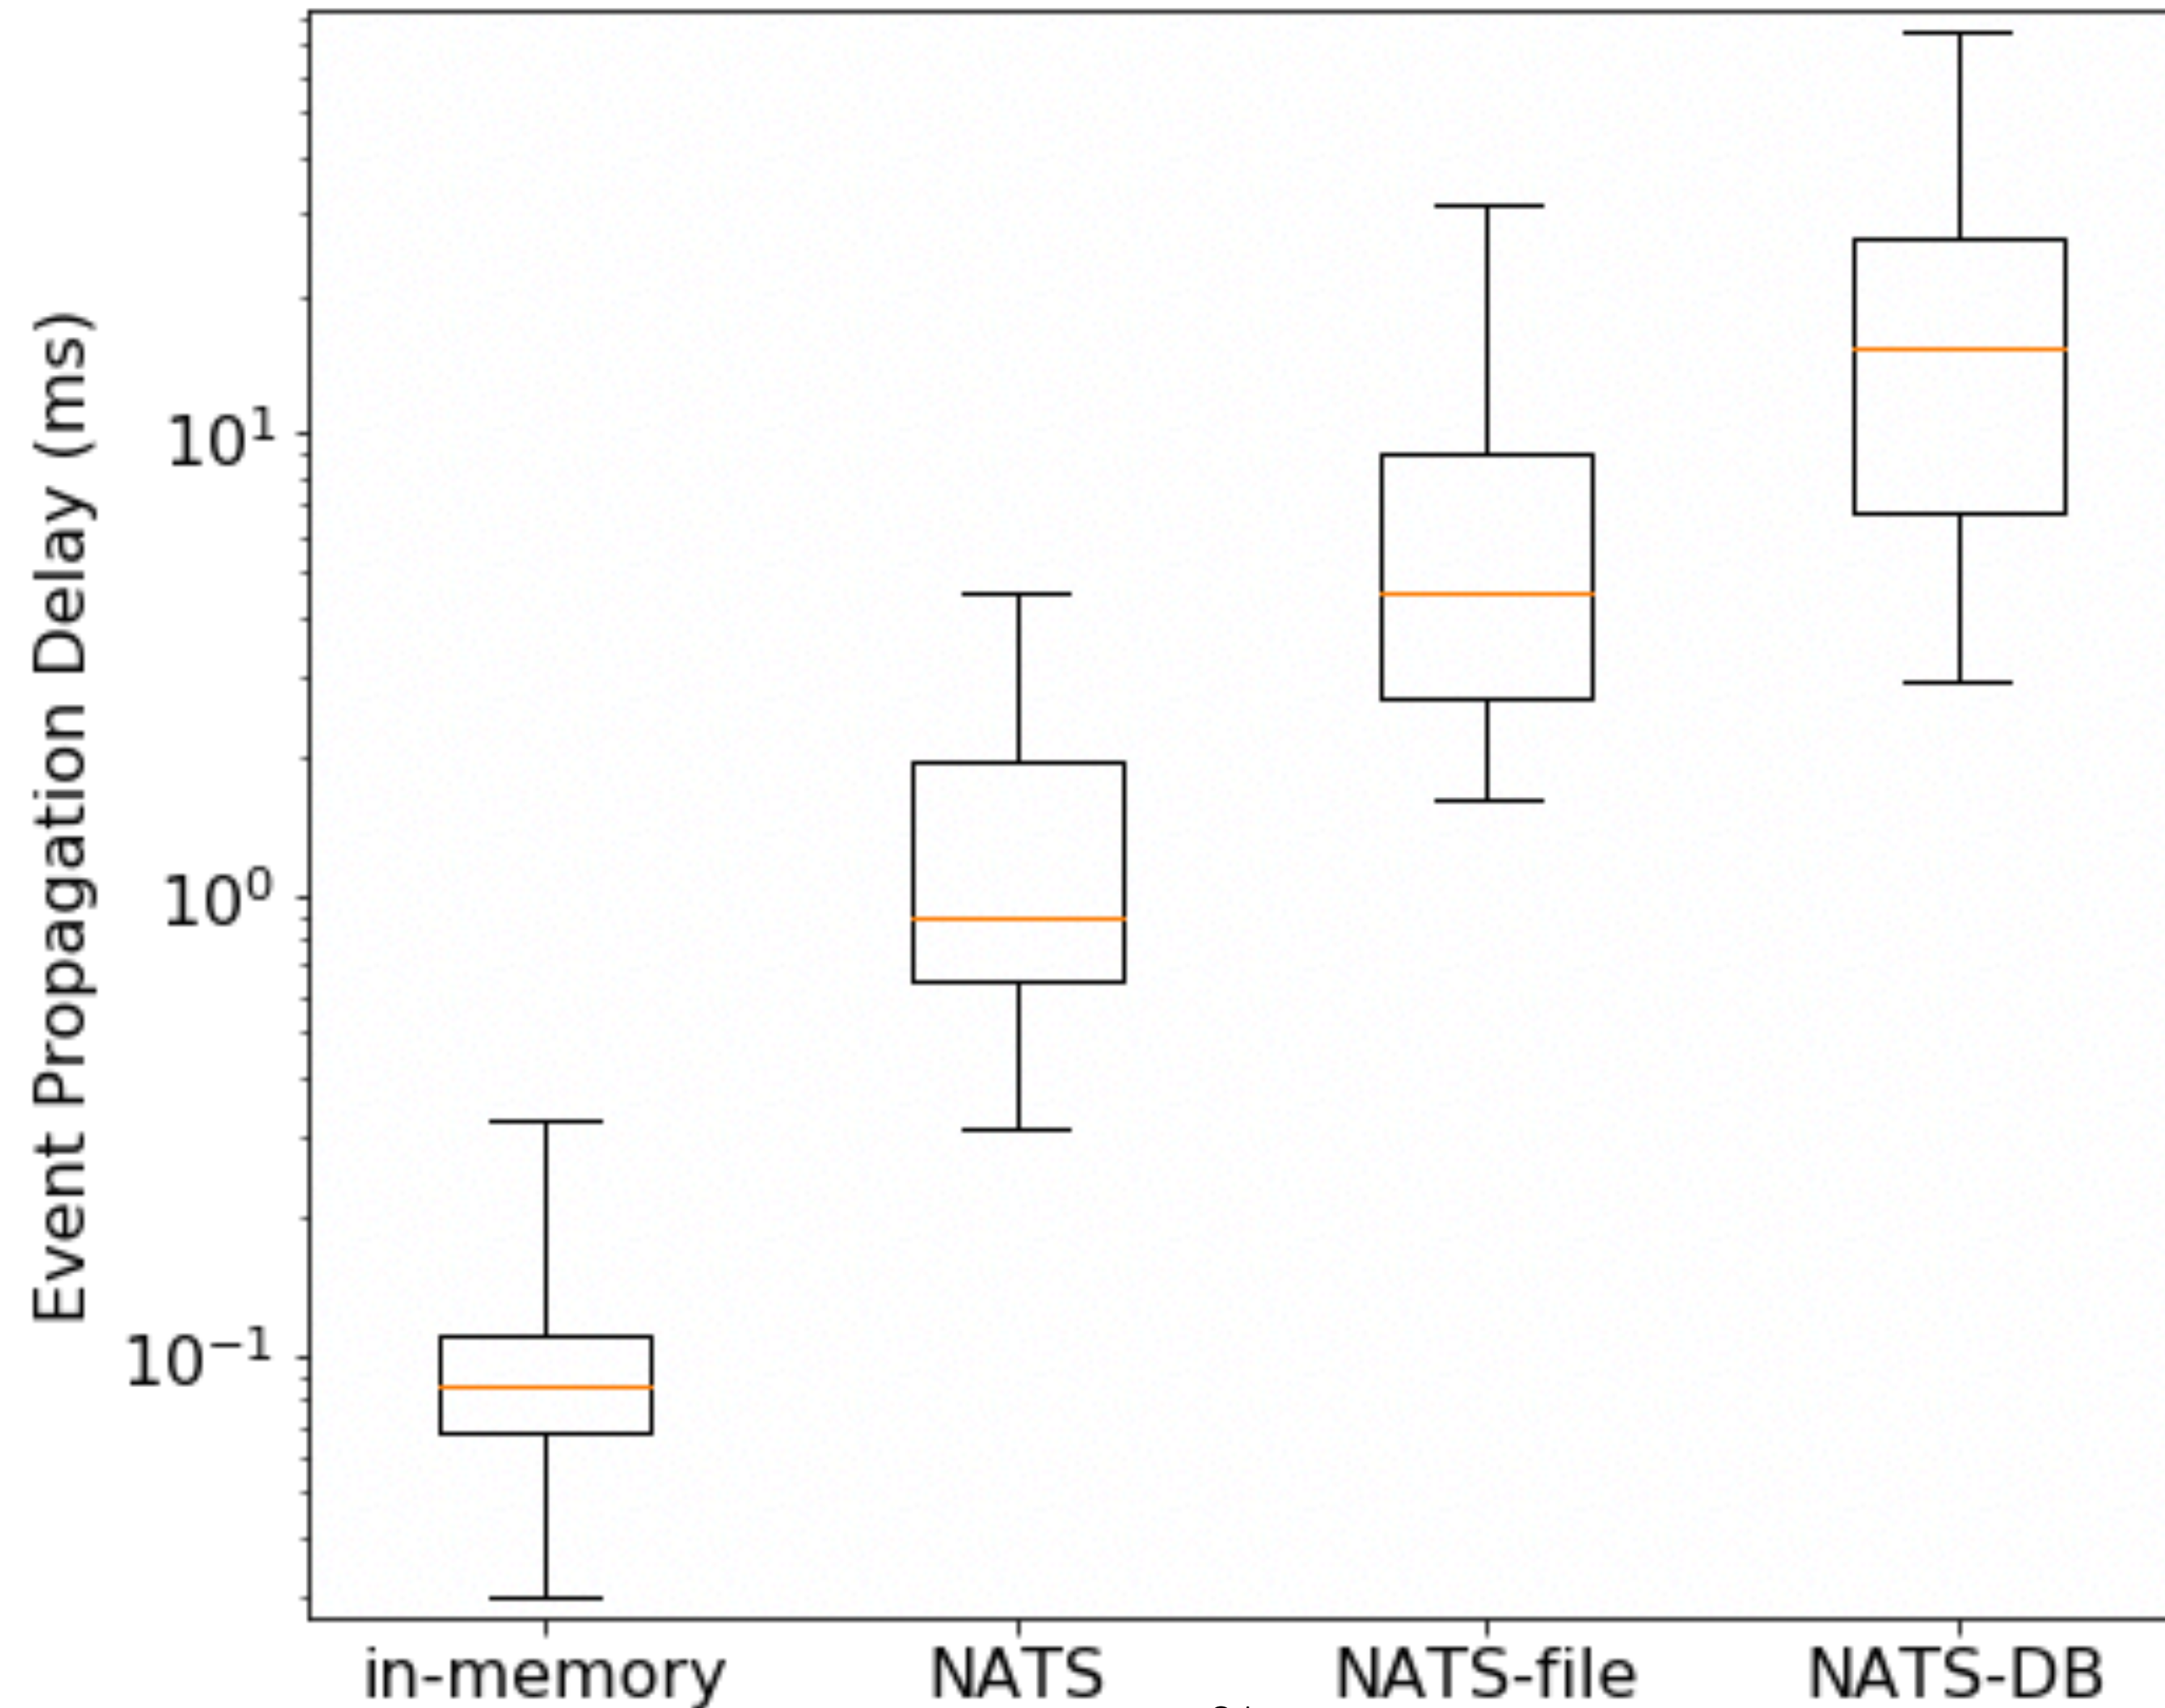


# Fault-tolerance: crash of workflow engine with an in-memory event store










# Performance overhead of event store implementations



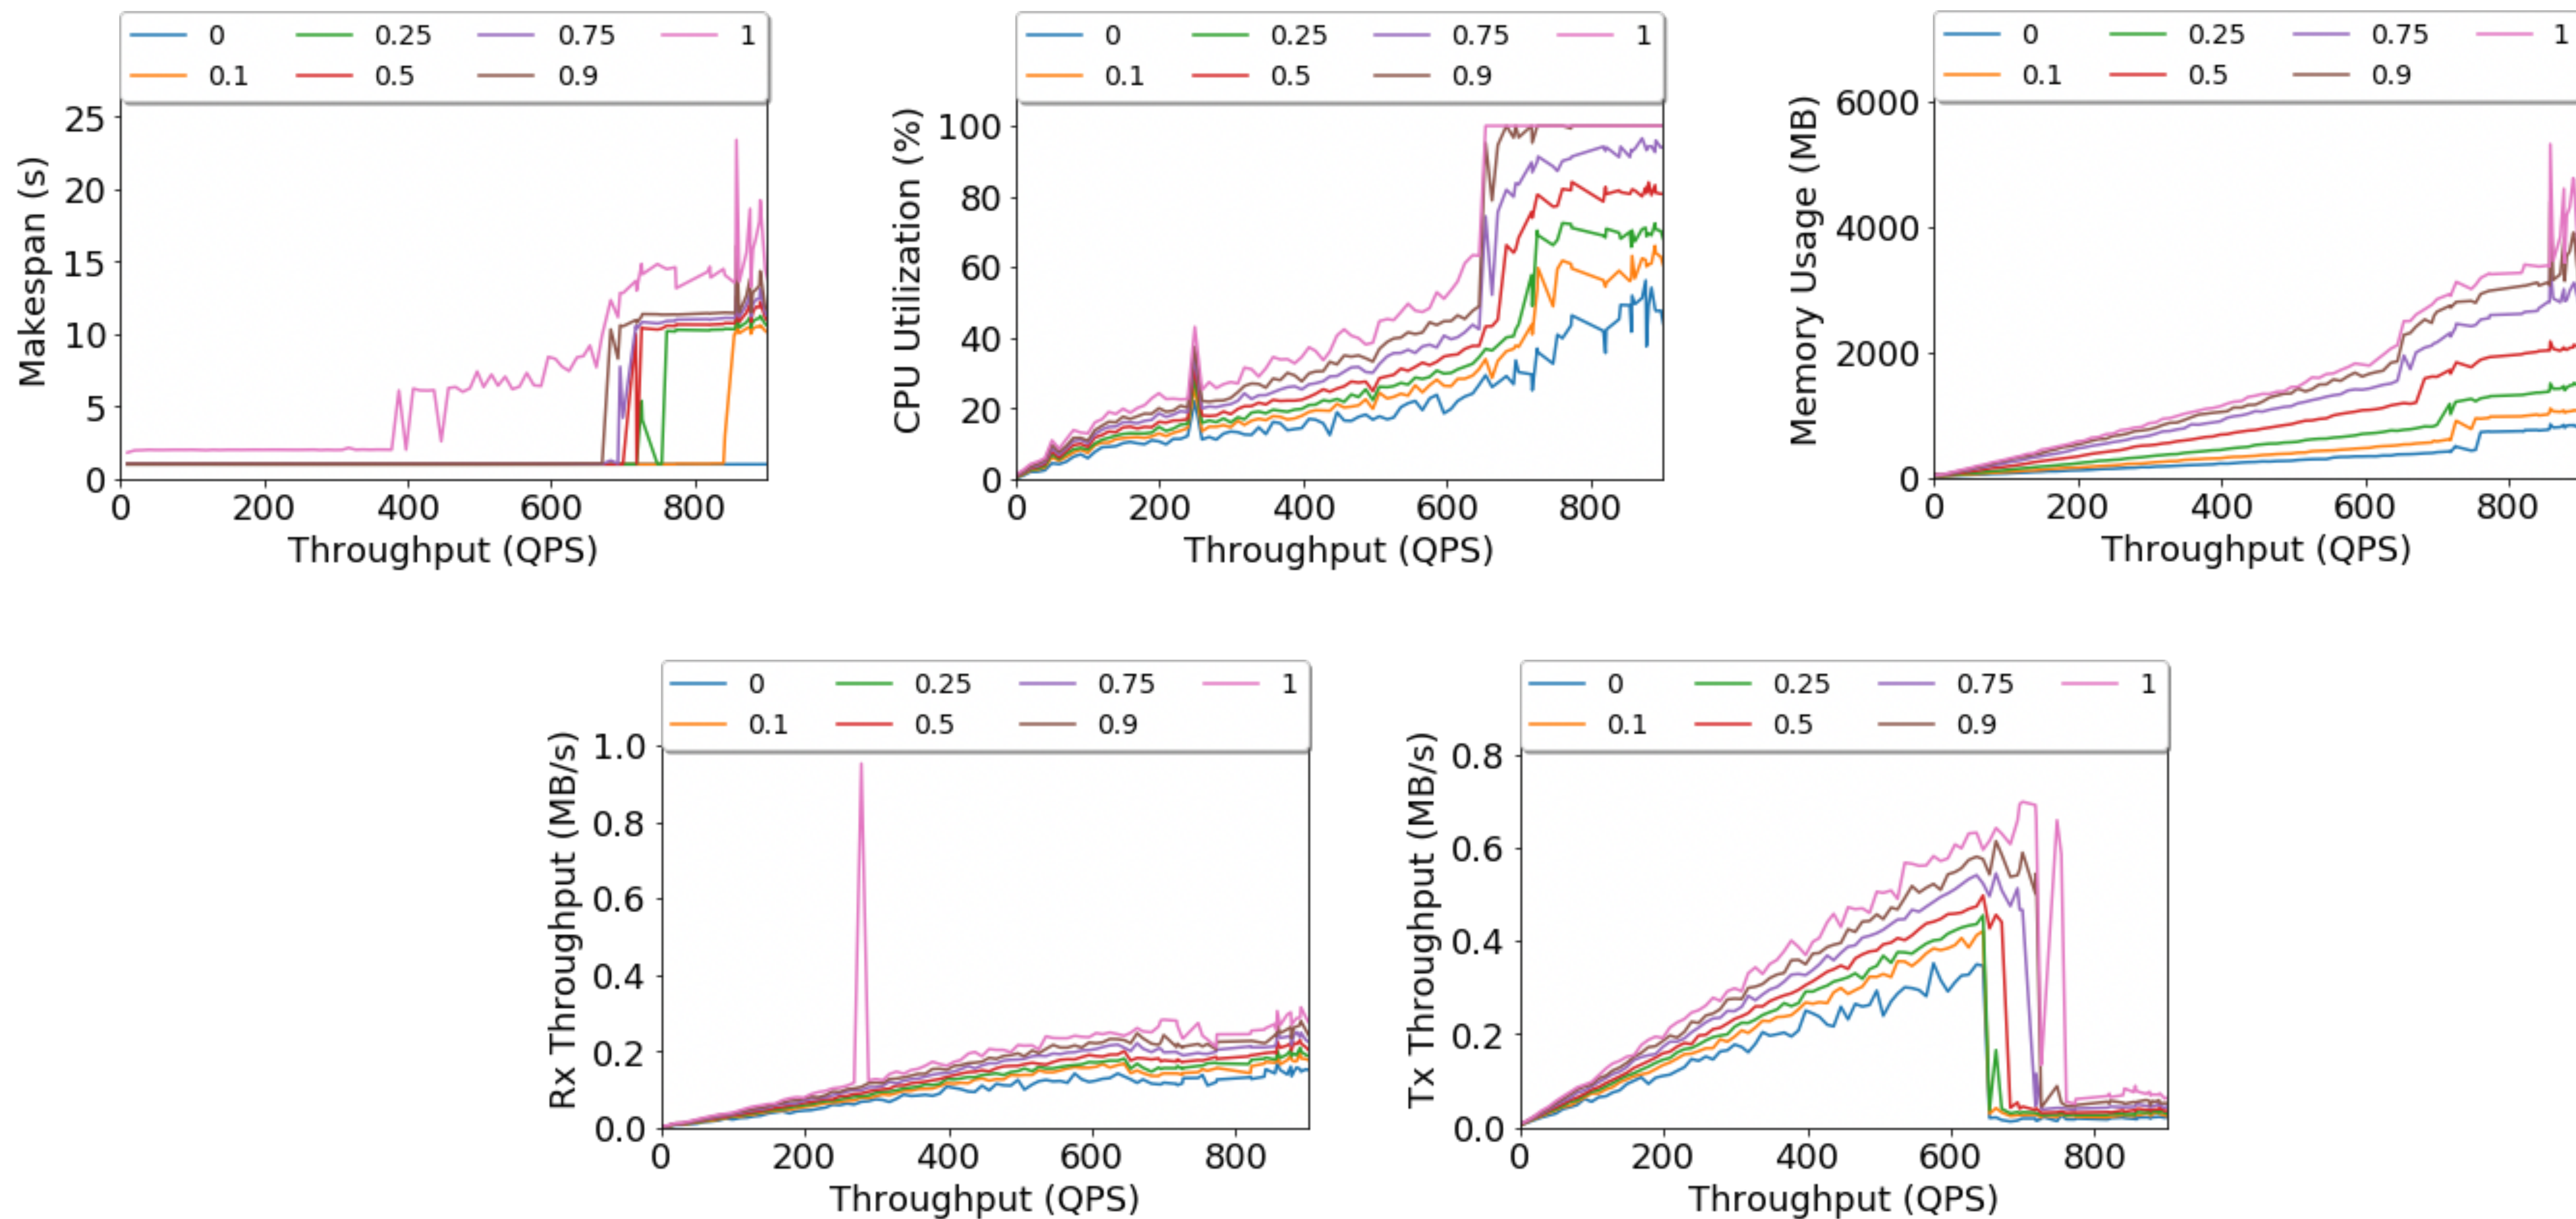


# Scalability: workflow invocation submission



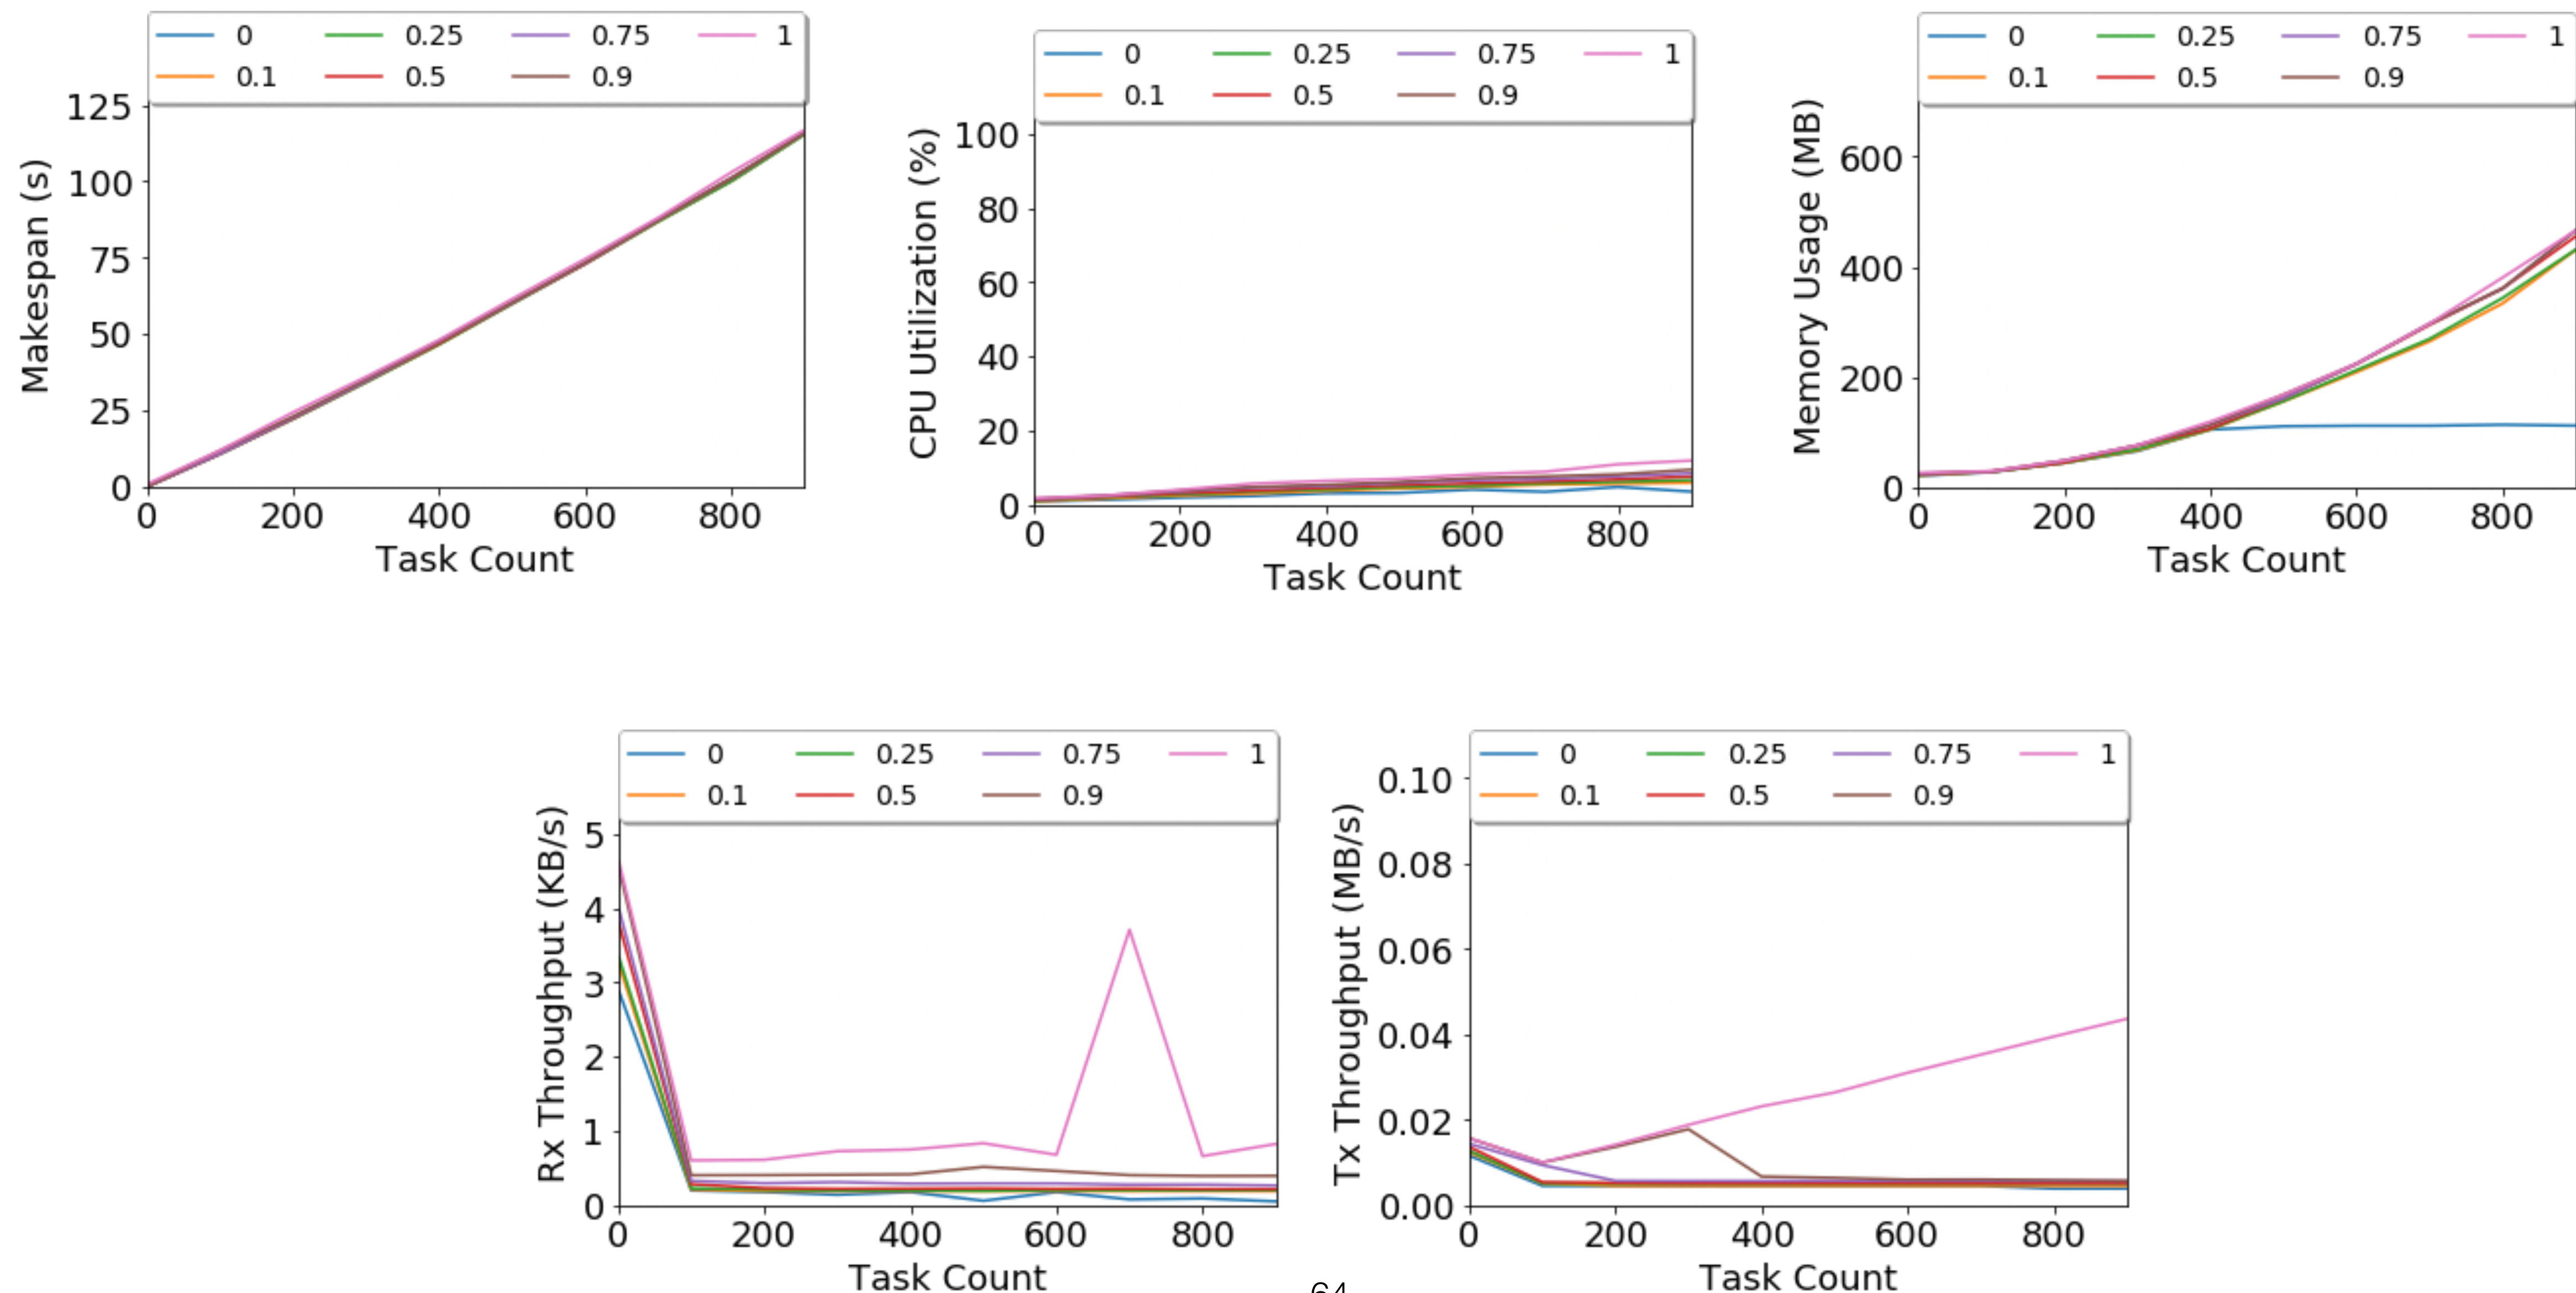


# Scalability: 1-task workflows










# Scalability: serial/long-running workflows






# Scalability: parallel workflows






# Resource consumption of scheduling policies






# Performance of scheduling policies





# Slowdown of the serverless workflow engines





# Cost breakdown of the evaluated serverless platforms

| Cost item                       | Fission Workflows | AWS        | Azure       | Google Cloud |
|---------------------------------|-------------------|------------|-------------|--------------|
| Compute nodes                   | 4                 | -          | -           | 4            |
| Node cost per hour              | 0.0475            | -          | -           | 0.0475       |
| Other hourly costs              | -                 | -          | -           | 0.4092       |
| Function cost per second        | -                 | 0.00000208 | 0.000014    | 0.00000203   |
| Function cost per execution     | -                 | 0.0000002  | 0.000000169 | 0.0000004    |
| Orchestration cost per workflow | -                 | 0.0001     | 0.00006     | -            |
| Experiment duration (seconds)   | 651.0             | 653.2      | 650.5       | 650          |
| <b>Total Costs</b>              | 0.0343            | 0.5278     | 1.455       | 0.293        |



# Fission Workflow definition example

```
1  apiVersion: 1
2  output: WhaleWithFortune
3  tasks:
4    GenerateFortune:
5      run: fortune
6
7    WhaleWithFortune:
8      run: whalesay
9      inputs:
10       body: "{ output('GenerateFortune') }"
11      requires:
12       - GenerateFortune
```



# Survey of FaaS platforms

| FaaS Platform          | F1                               | F2                               | F3                               | F4                               | F5                               | F6                               |
|------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|
| AWS Lambda             | <input type="radio"/>            | <input checked="" type="radio"/> | <input type="radio"/>            | <input type="radio"/>            | <input checked="" type="radio"/> | <input checked="" type="radio"/> |
| Azure Functions        | <input type="radio"/>            | <input checked="" type="radio"/> | <input type="radio"/>            | <input type="radio"/>            | <input checked="" type="radio"/> | <input checked="" type="radio"/> |
| Google Cloud Functions | <input type="radio"/>            | <input checked="" type="radio"/> | <input type="radio"/>            | <input type="radio"/>            | <input checked="" type="radio"/> | <input checked="" type="radio"/> |
| OpenWhisk              | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input type="radio"/>            | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input checked="" type="radio"/> |
| OpenFaaS               | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input type="radio"/>            | <input type="radio"/>            | <input checked="" type="radio"/> | <input checked="" type="radio"/> |
| OpenLambda             | <input checked="" type="radio"/> | <input type="radio"/>            | <input checked="" type="radio"/> | <input type="radio"/>            | <input type="radio"/>            | <input checked="" type="radio"/> |
| Kubeless               | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input type="radio"/>            | <input checked="" type="radio"/> | <input checked="" type="radio"/> |
| Fission                | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input checked="" type="radio"/> | <input checked="" type="radio"/> |
| Funktion               | <input checked="" type="radio"/> | <input type="radio"/>            | <input type="radio"/>            | <input type="radio"/>            | <input type="radio"/>            | <input checked="" type="radio"/> |



# Survey of composition approaches

| Approach     | Constraints |    |    |    | Requirements |    |    |    |    |    |    |
|--------------|-------------|----|----|----|--------------|----|----|----|----|----|----|
|              | C1          | C2 | C3 | C4 | R1           | R2 | R3 | R4 | R5 | R6 | R7 |
| Direct       | ●           | ●  | ●  | ●  | ●            | ○  | ○  | ○  | ●  | ●  | ●  |
| Compiled     | ●           | ○  | ●  | ●  | ◐            | ○  | ●  | ●  | ○  | ○  | ●  |
| Coordinator  | ●           | ●  | ●  | ●  | ●            | ○  | ○  | ○  | ●  | ●  | ●  |
| Event-Driven | ●           | ●  | ○  | ◐  | ○            | ●  | ●  | ●  | ○  | ◐  | ◐  |
| Workflows    | ●           | ●  | ●  | ◐  | ●            | ●  | ◐  | ●  | ●  | ●  | ●  |



# Survey of WMSs

| Approach           | Constraints |    |    |    | Requirements |    |    |    |    |    |    | OSS | Docs |
|--------------------|-------------|----|----|----|--------------|----|----|----|----|----|----|-----|------|
|                    | C1          | C2 | C3 | C4 | R1           | R2 | R3 | R4 | R5 | R6 | R7 |     |      |
| NodeRED            | ●           | ◐  | ◐  | ◐  | ●            | ◐  | ○  | ○  | ○  | ○  | ●  | ●   | ●    |
| Openstack Mistral  | ◐           | ○  | ◐  | ○  | ◐            | ●  | ◐  | ●  | ○  | ●  | ●  | ●   | ●    |
| AWS Step Functions | ◐           | ●  | ●  | ●  | ◐            | ●  | ●  | ●  | ◐  | ●  | ●  | ○   | ○    |
| Azure Logic Apps   | ●           | ◐  | ◐  | ●  | ●            | ●  | ●  | ●  | ◐  | ●  | ●  | ○   | ○    |
| Pywren             | ●           | ○  | ●  | ○  | ○            | ●  | ●  | ○  | ○  | ●  | ●  | ●   | ◐    |
| Apache Airflow     | ○           | ◐  | ○  | ○  | ◐            | ●  | ◐  | ●  | ○  | ◐  | ◐  | ●   | ●    |
| Azkaban            | ◐           | ○  | ◐  | ◐  | ○            | ●  | ◐  | ●  | ○  | ◐  | ●  | ●   | ○    |
| Luigi              | ○           | ○  | ◐  | ◐  | ◐            | ●  | ○  | ●  | ●  | ◐  | ●  | ●   | ●    |
| Pegasus            | ◐           | ◐  | ○  | ◐  | ○            | ●  | ○  | ●  | ●  | ●  | ●  | ●   | ◐    |

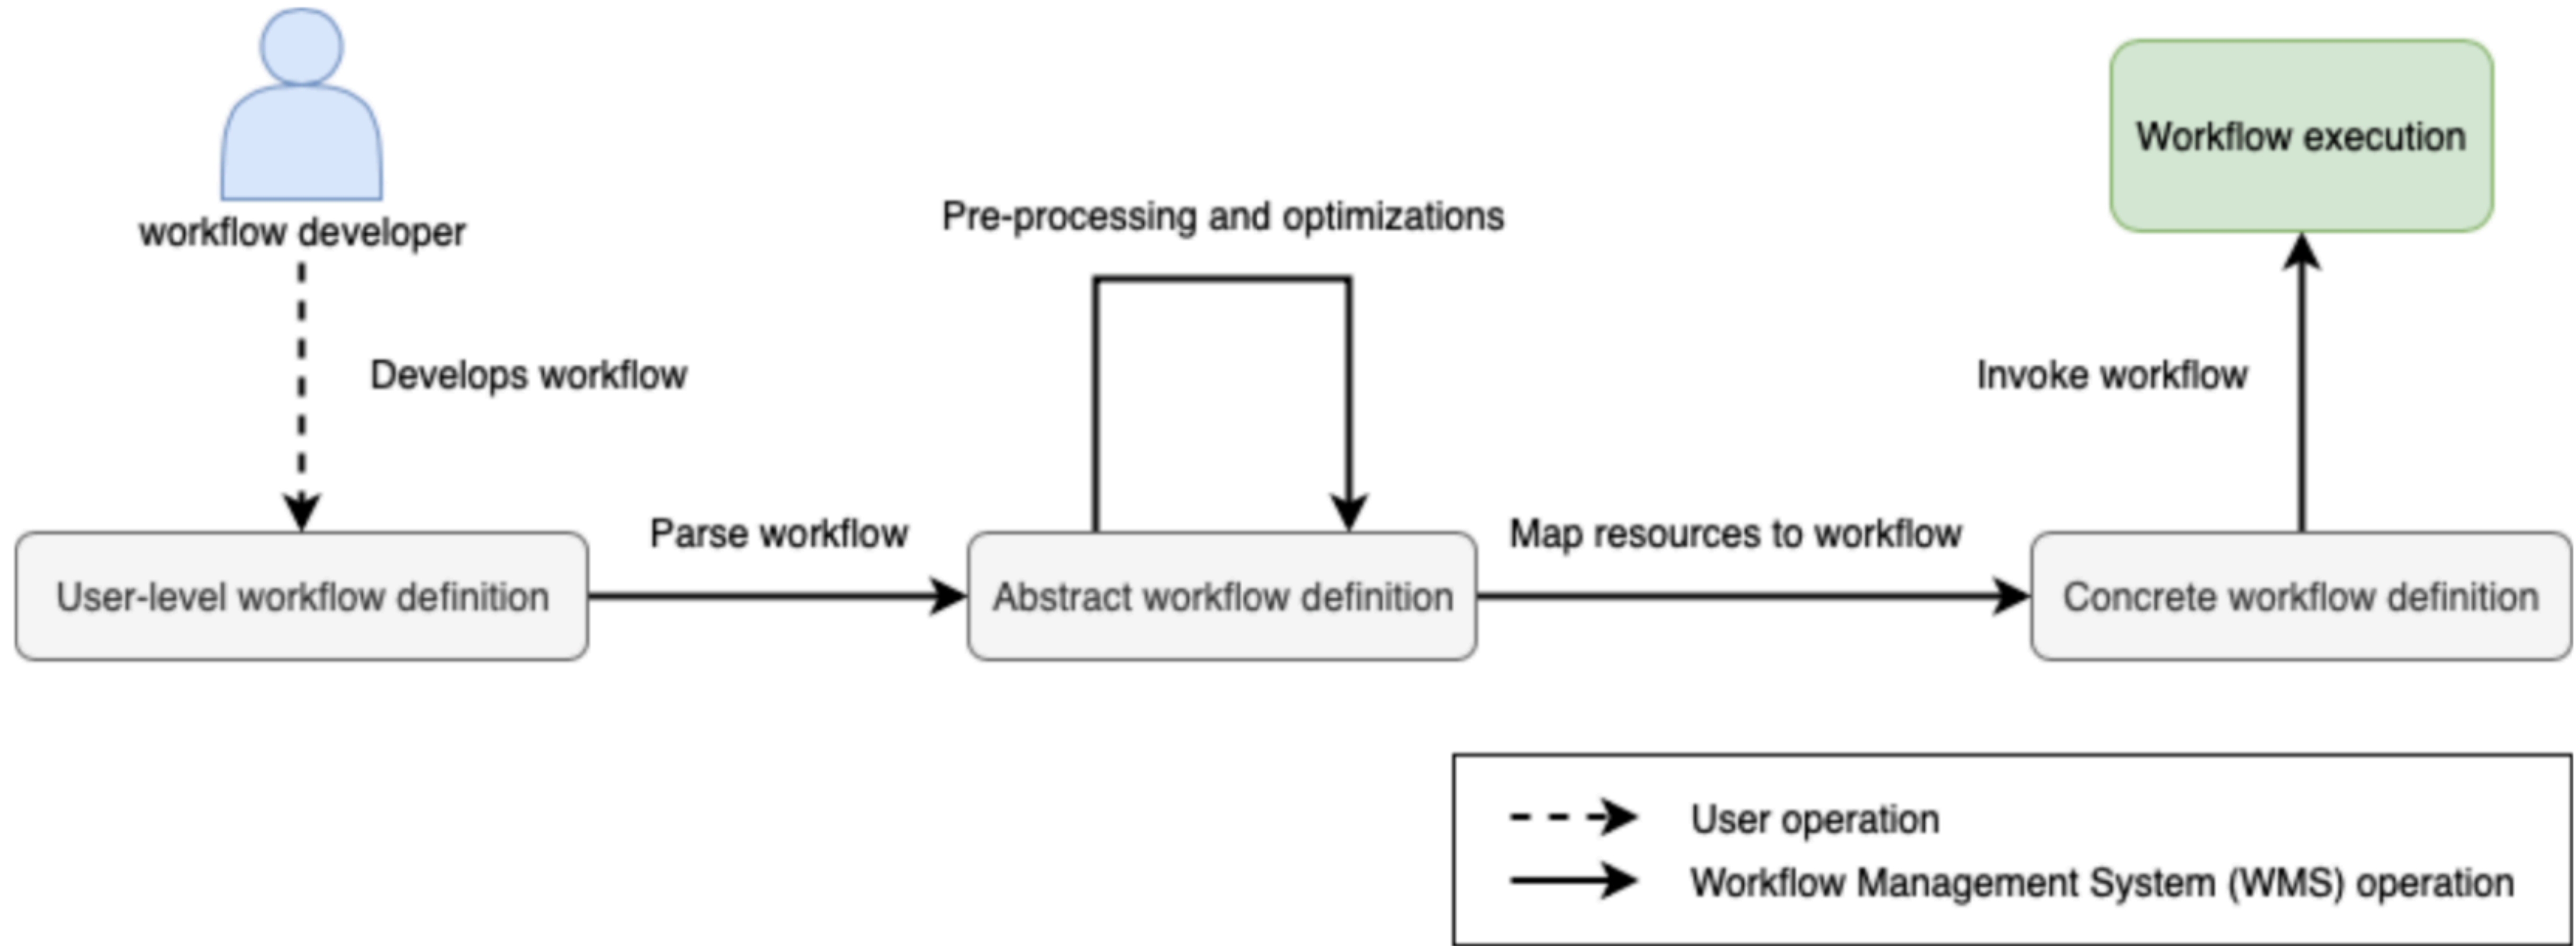


# Survey of workflow languages

| Workflow Language                  | Format     | FaaS functions | Complex Control Flows | Readability | Model fit |
|------------------------------------|------------|----------------|-----------------------|-------------|-----------|
| BPMN 2.0                           | XML        | ○              | ●                     | ○           | ○         |
| WS-BPEL                            | XML        | ○              | ●                     | ○           | ○         |
| YAWL                               | XML        | ○              | ●                     | ○           | ○         |
| WDL                                | Custom     | ○              | ○                     | ●           | ◐         |
| CWL                                | YAML, JSON | ●              | ○                     | ●           | ◐         |
| Amazon States Language             | JSON       | ●              | ◐                     | ◐           | ●         |
| Openstack Mistral Language         | YAML       | ◐              | ◐                     | ●           | ◐         |
| Apache Airflow's Workflow format   | Python     | ◐              | ●                     | ●           | ◐         |
| Azure Workflow Definition Language | JSON       | ●              | ●                     | ◐           | ◐         |
| Pegasus DAX                        | XML        | ◐              | ○                     | ○           | ◐         |
| Azkaban Language                   | Custom     | ◐              | ○                     | ◐           | ◐         |
| NodeRED Flow File                  | JSON       | ○              | ○                     | ○           | ○         |
| Luigi's Workflow Format            | Python     | ○              | ●                     | ●           | ◐         |
| Amazon Simple Workflow Framework   | Java       | ◐              | ●                     | ○           | ○         |

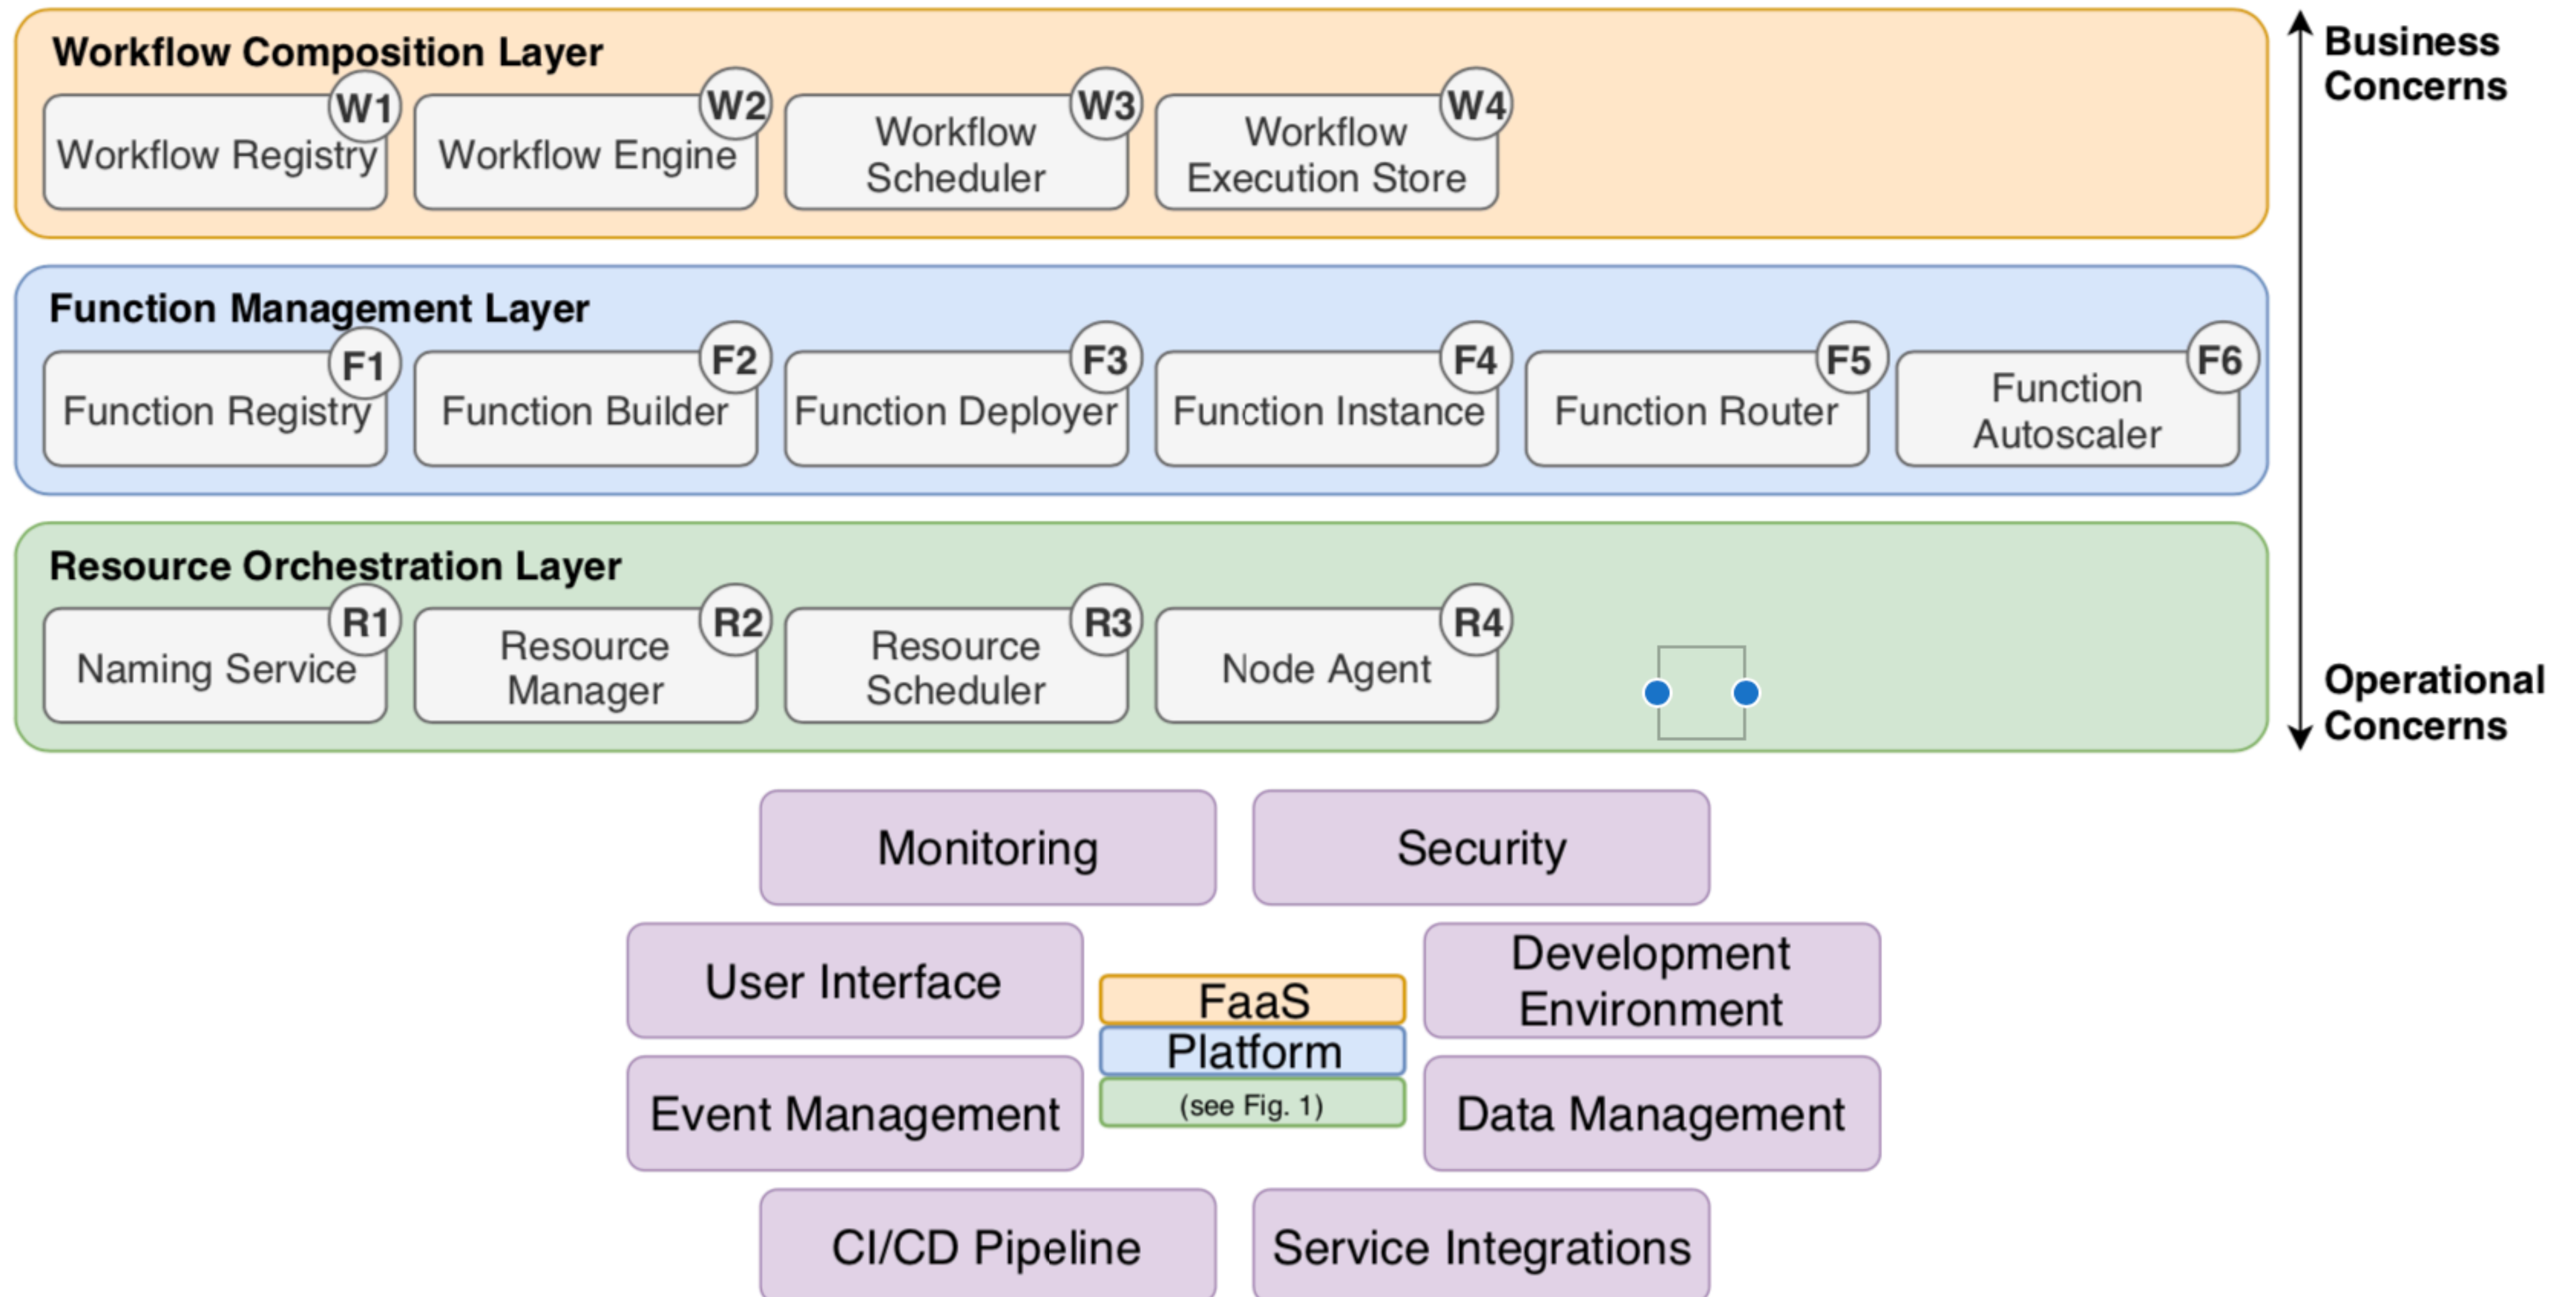


# Workflow definition lifecycle



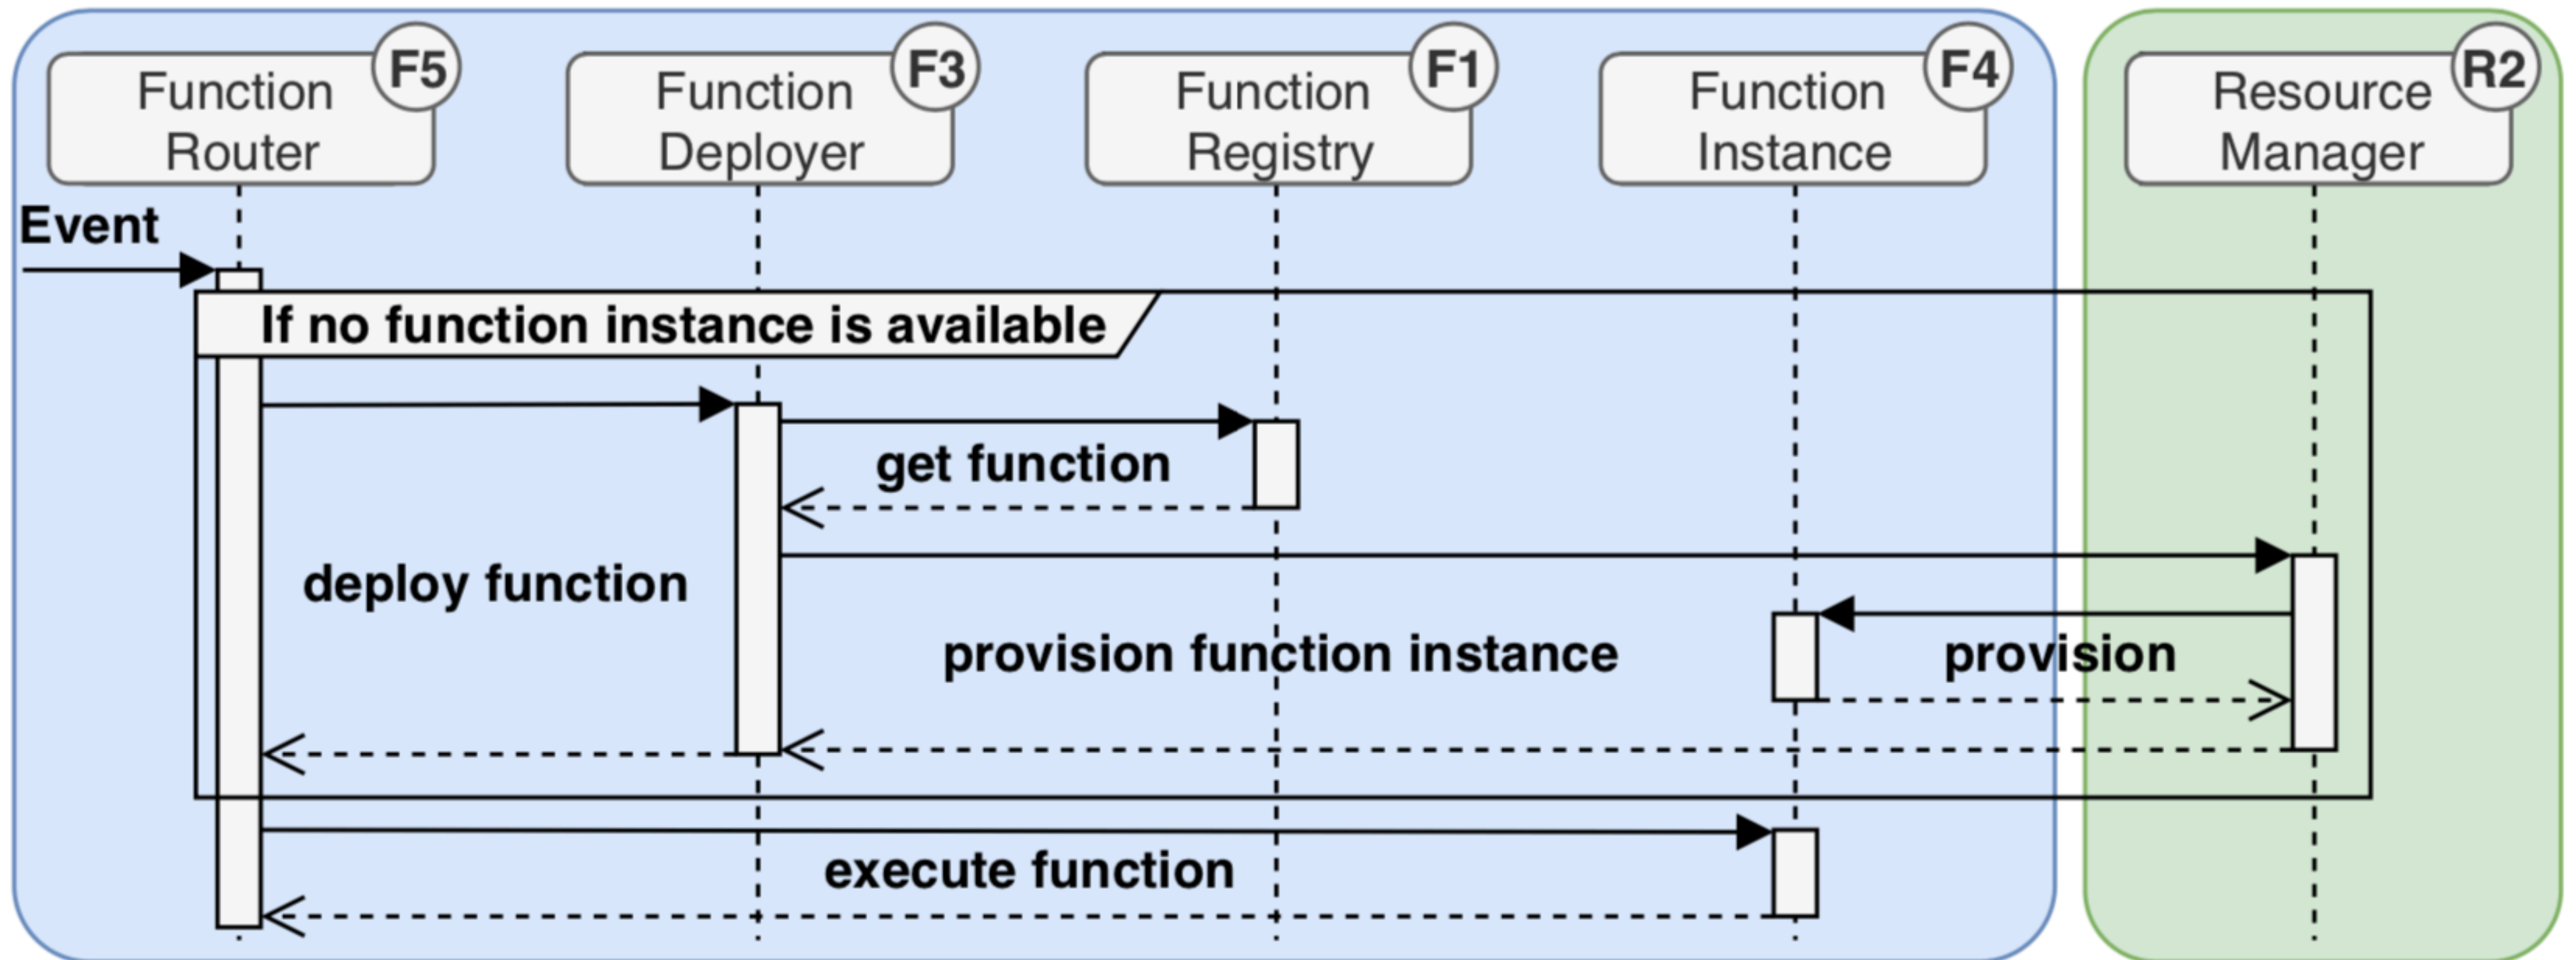


# FaaS platform reference architecture



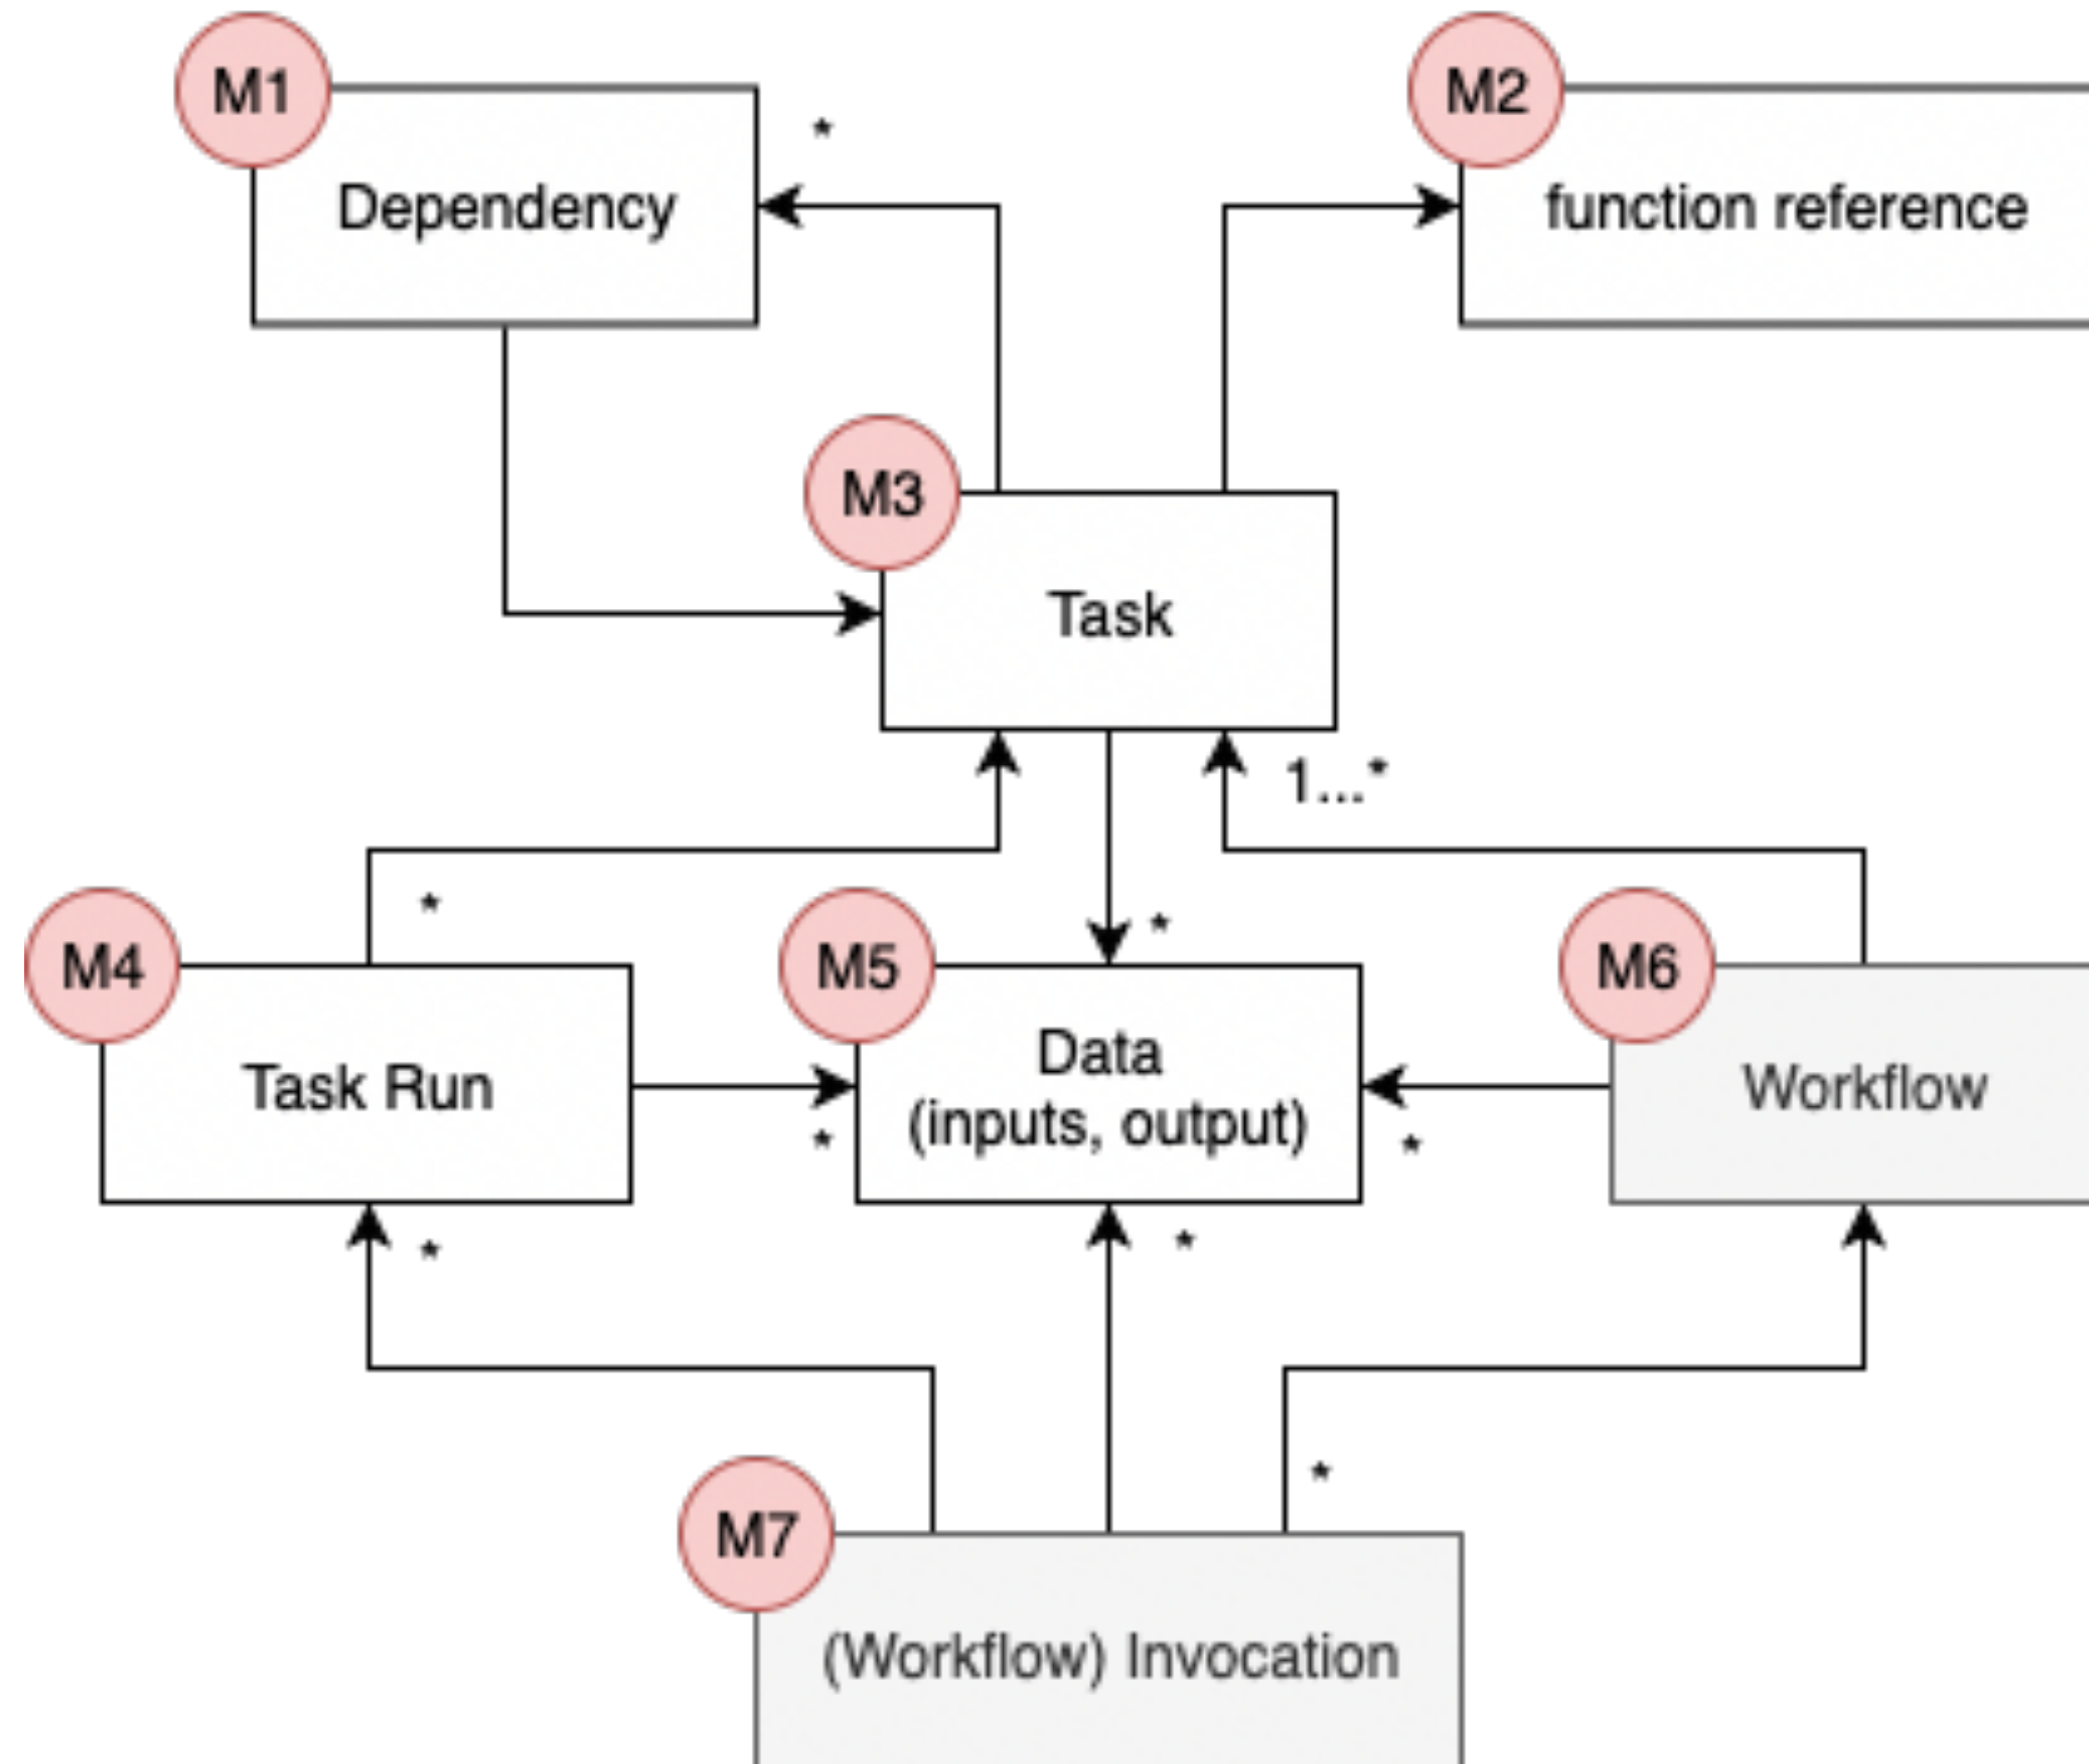


# Serverless function execution



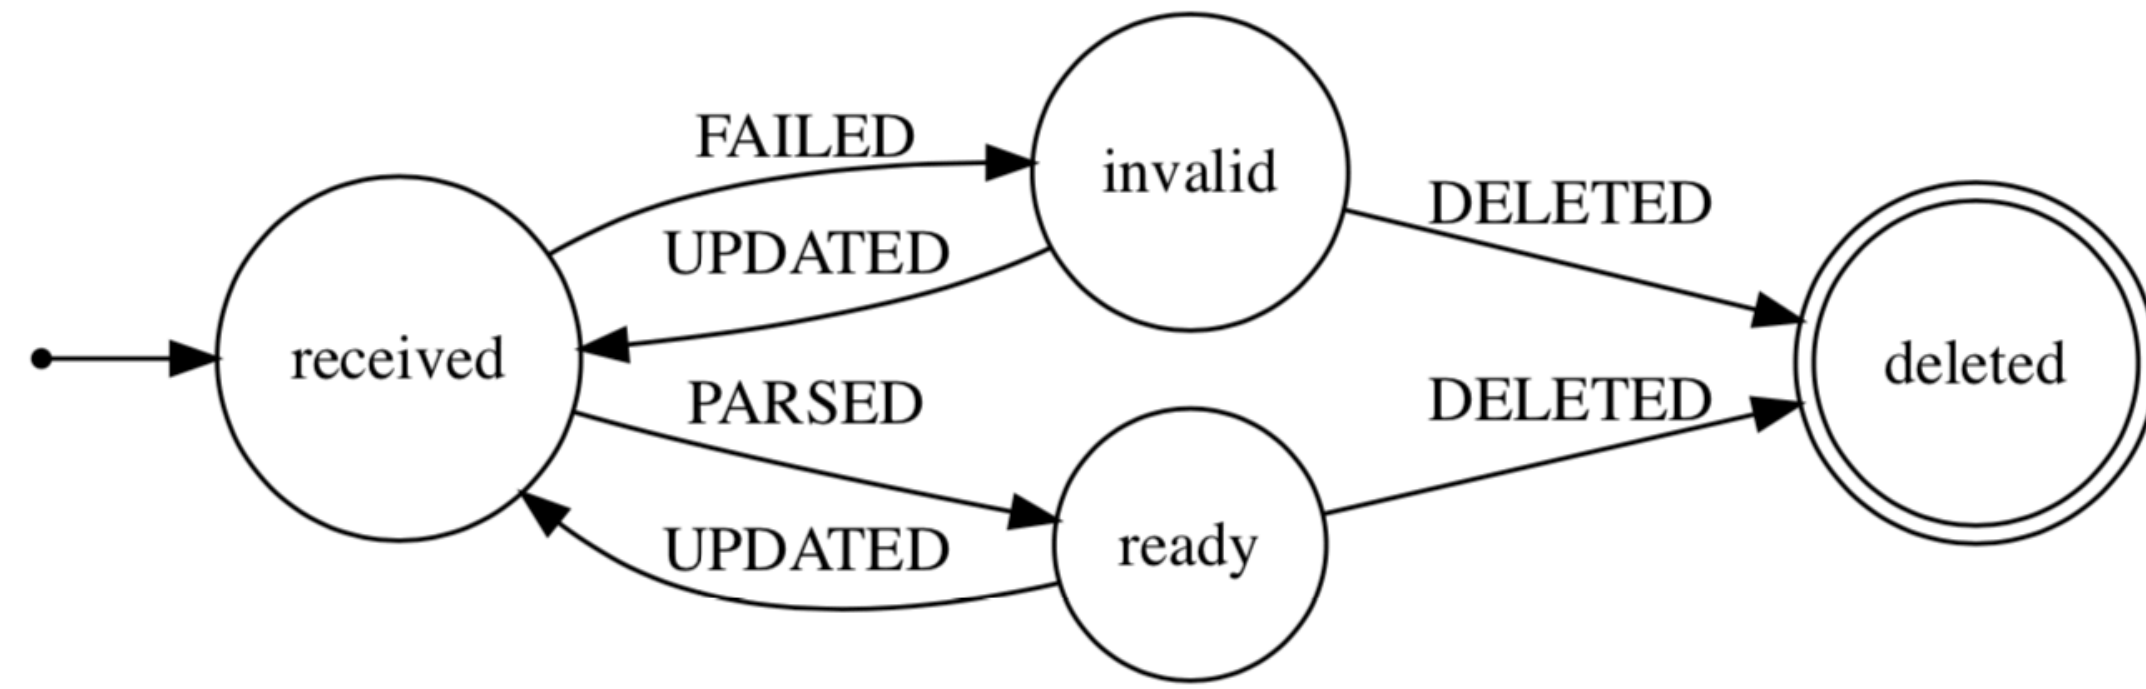


# SWL: data model

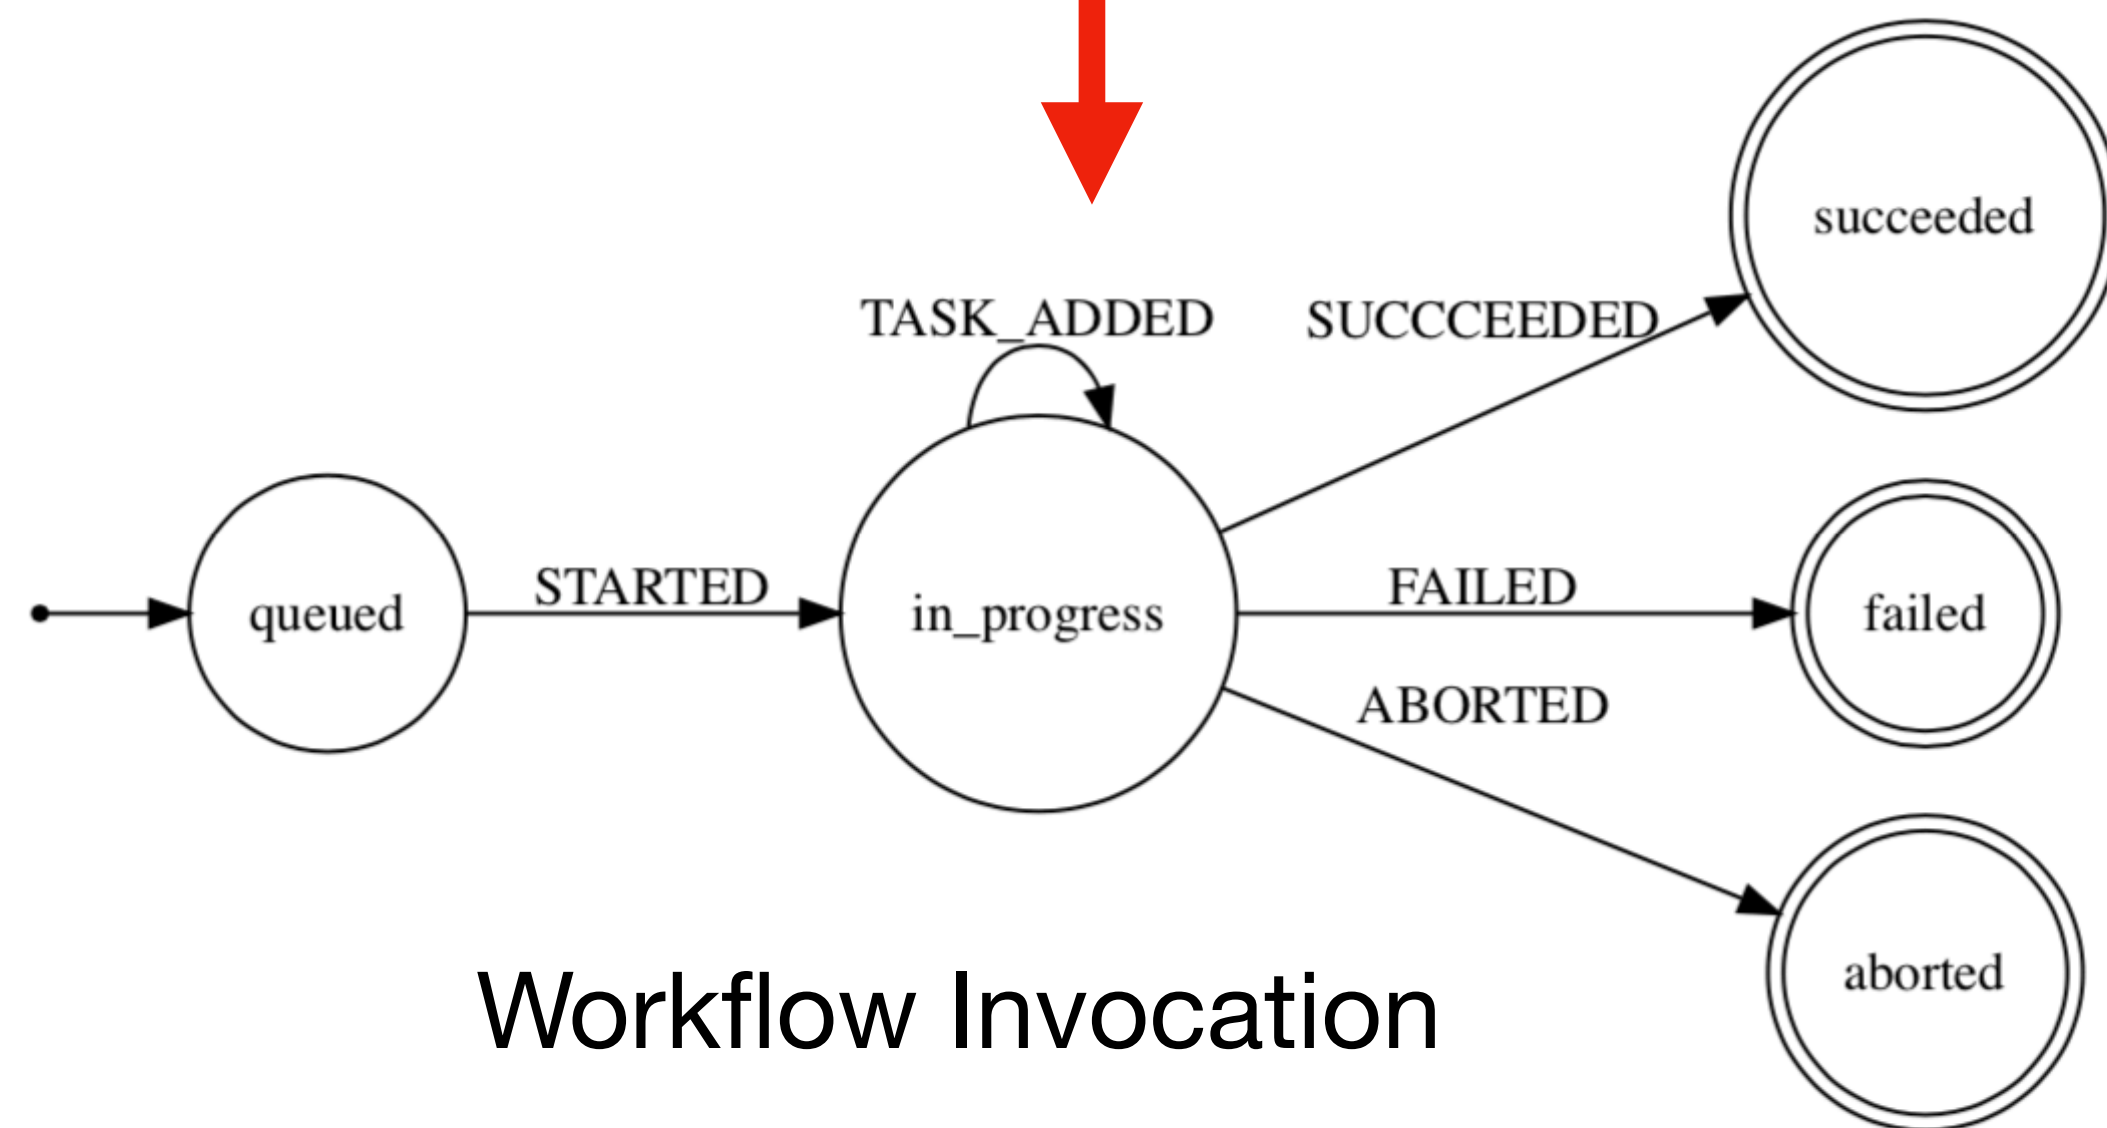




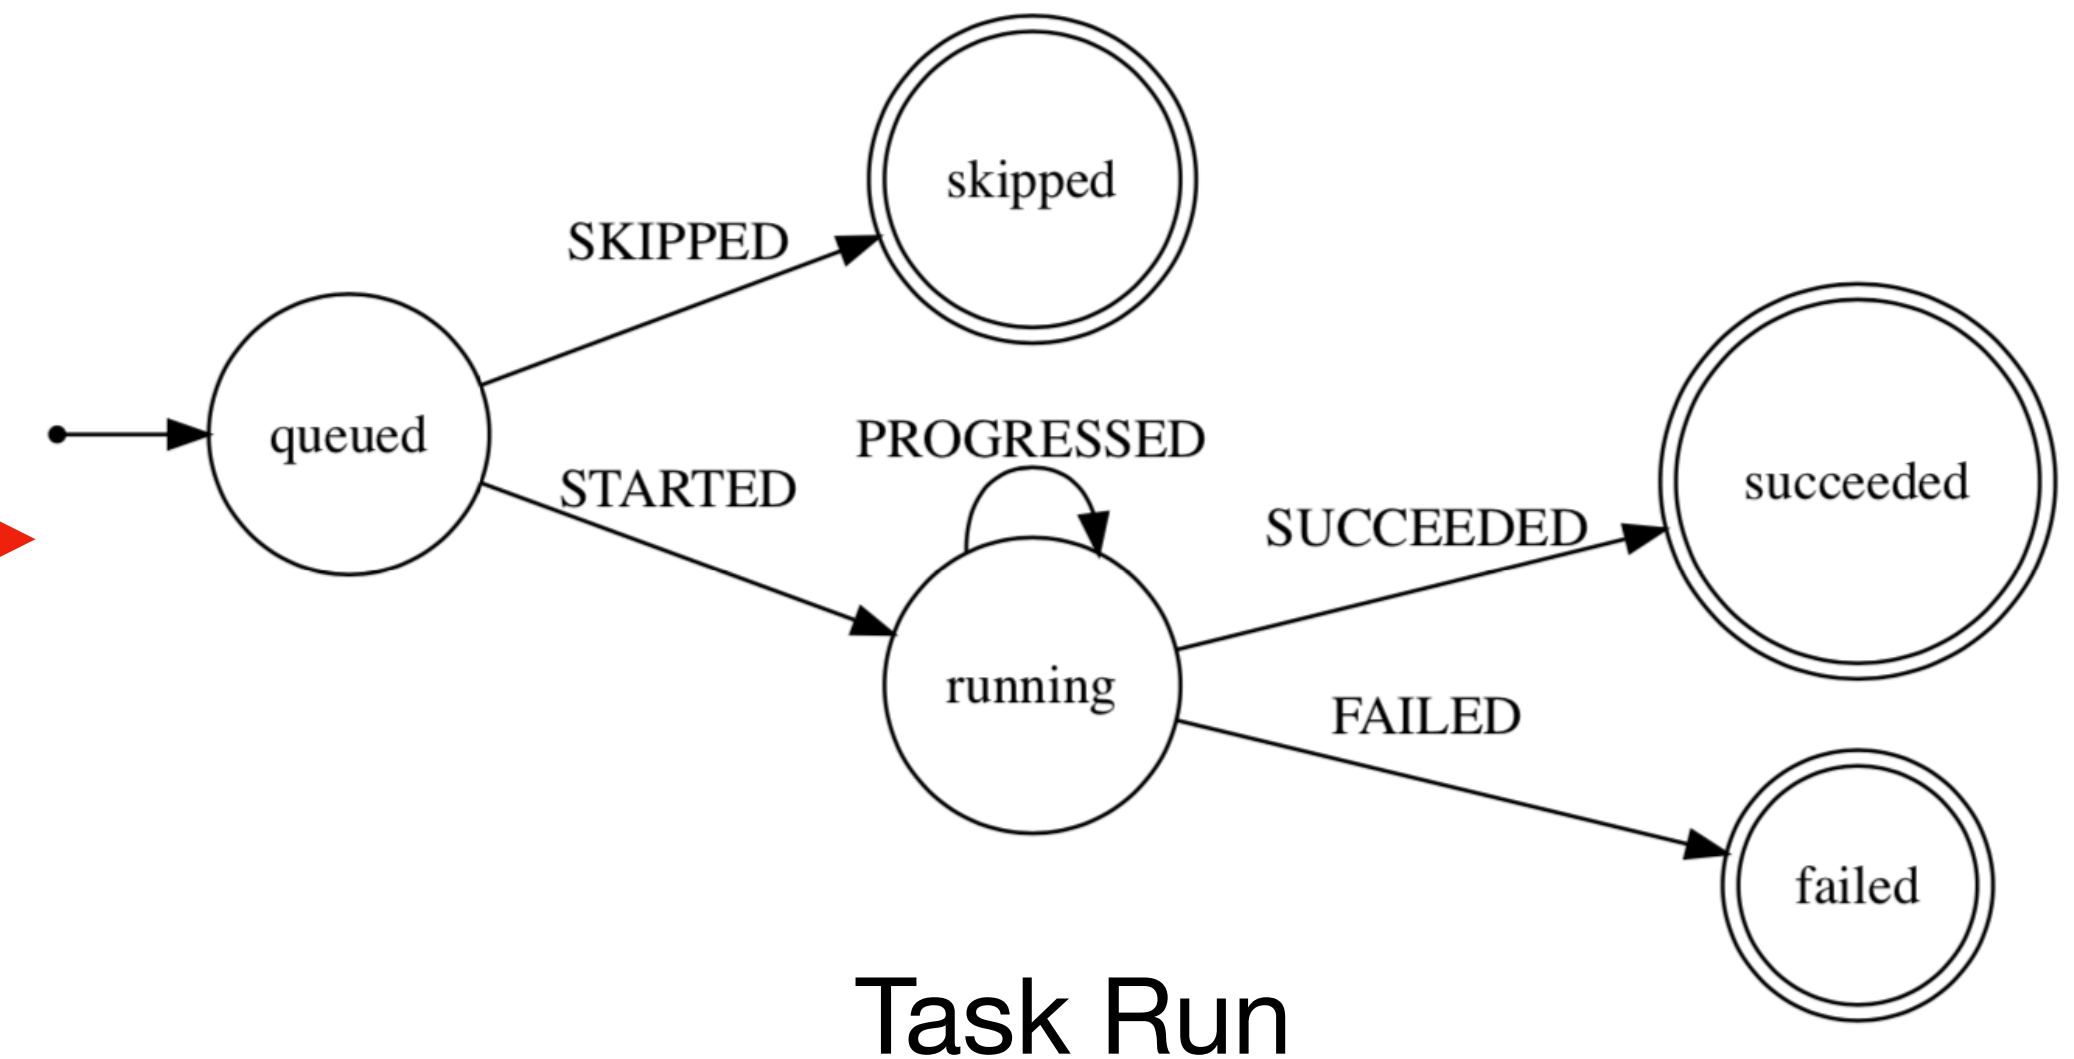
# SWL: execution model



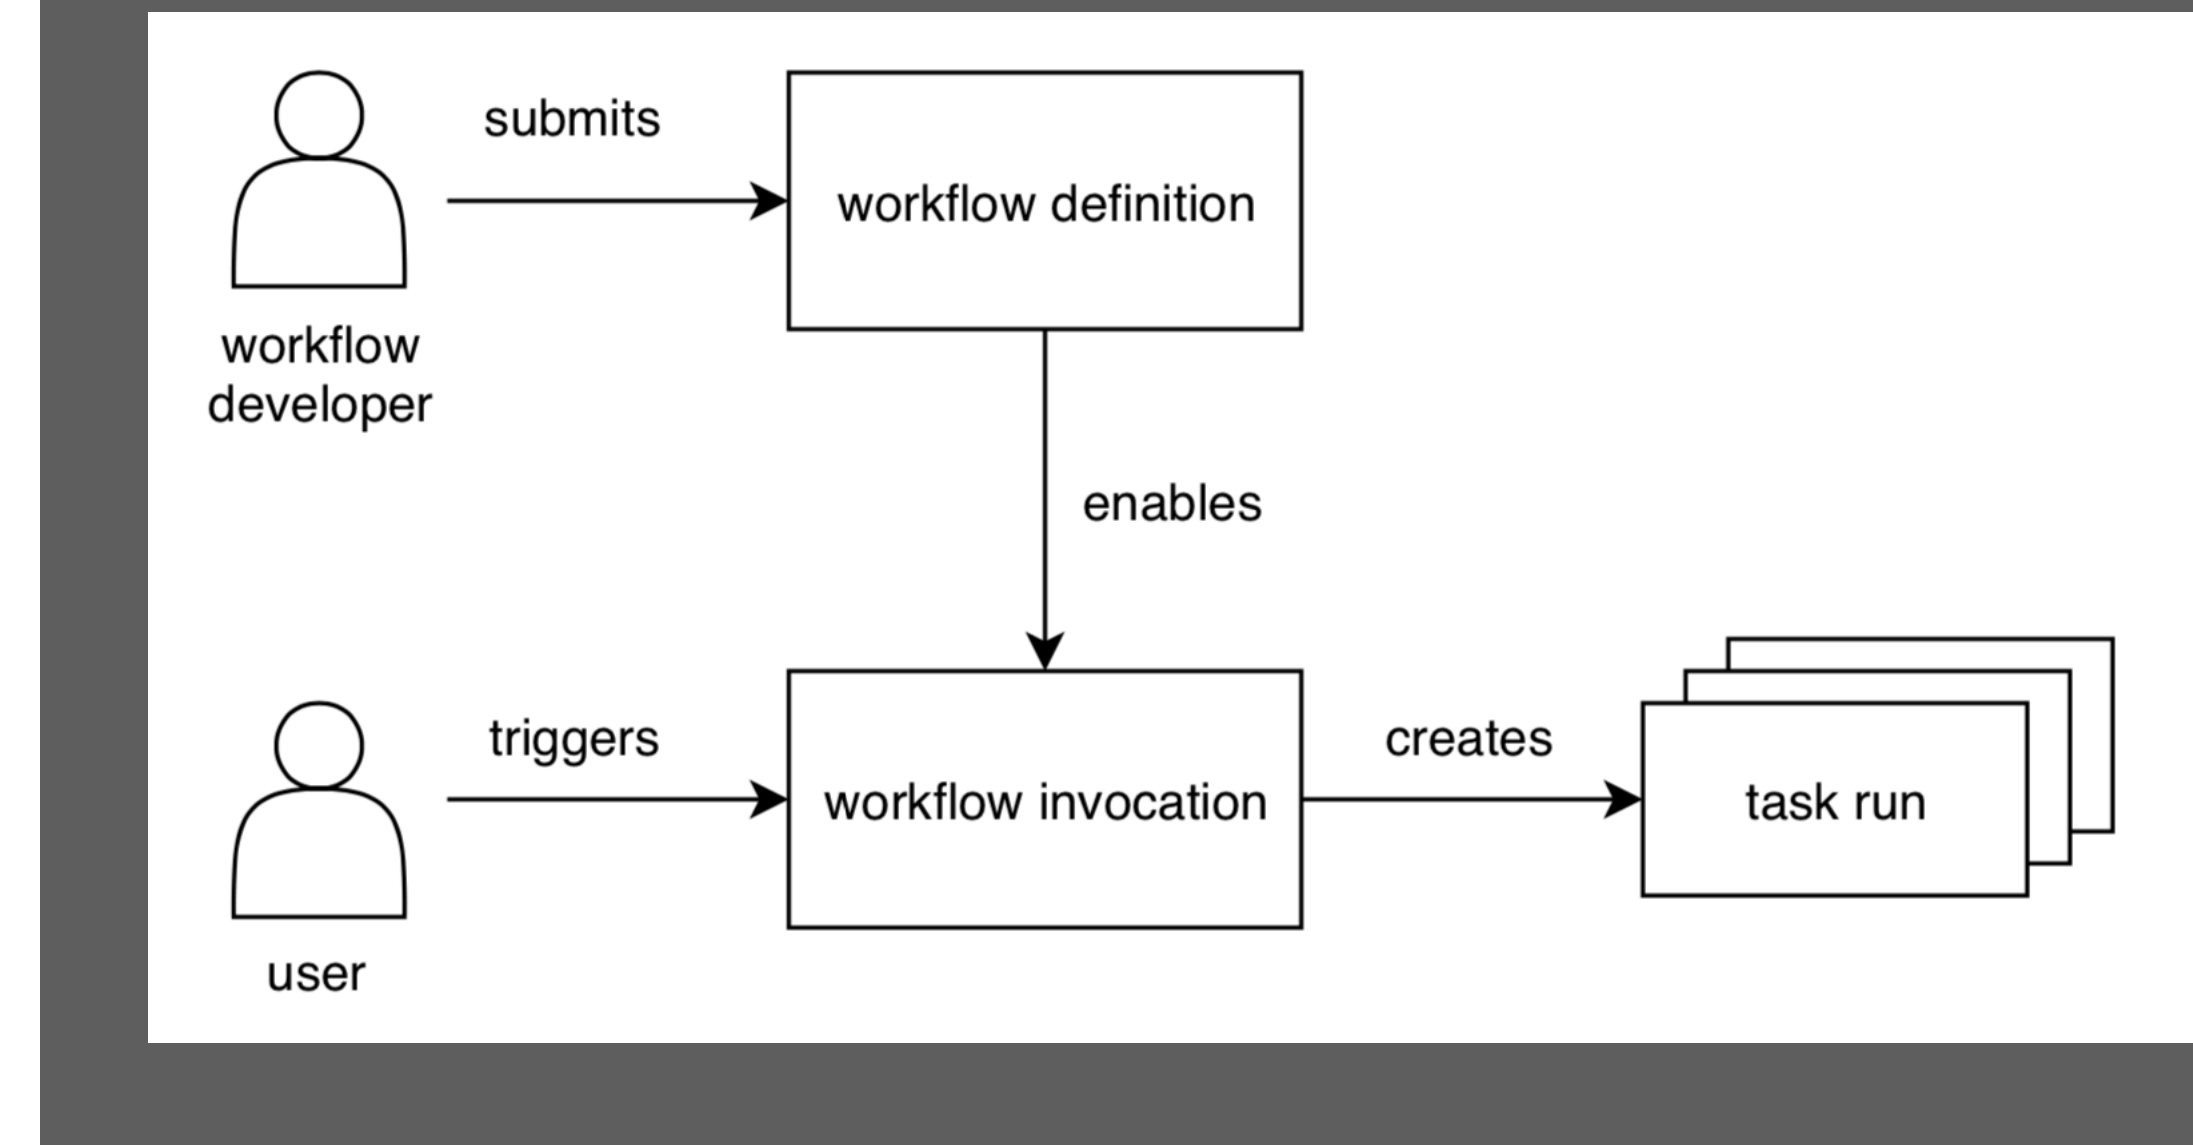
Workflow Definition



Workflow Invocation

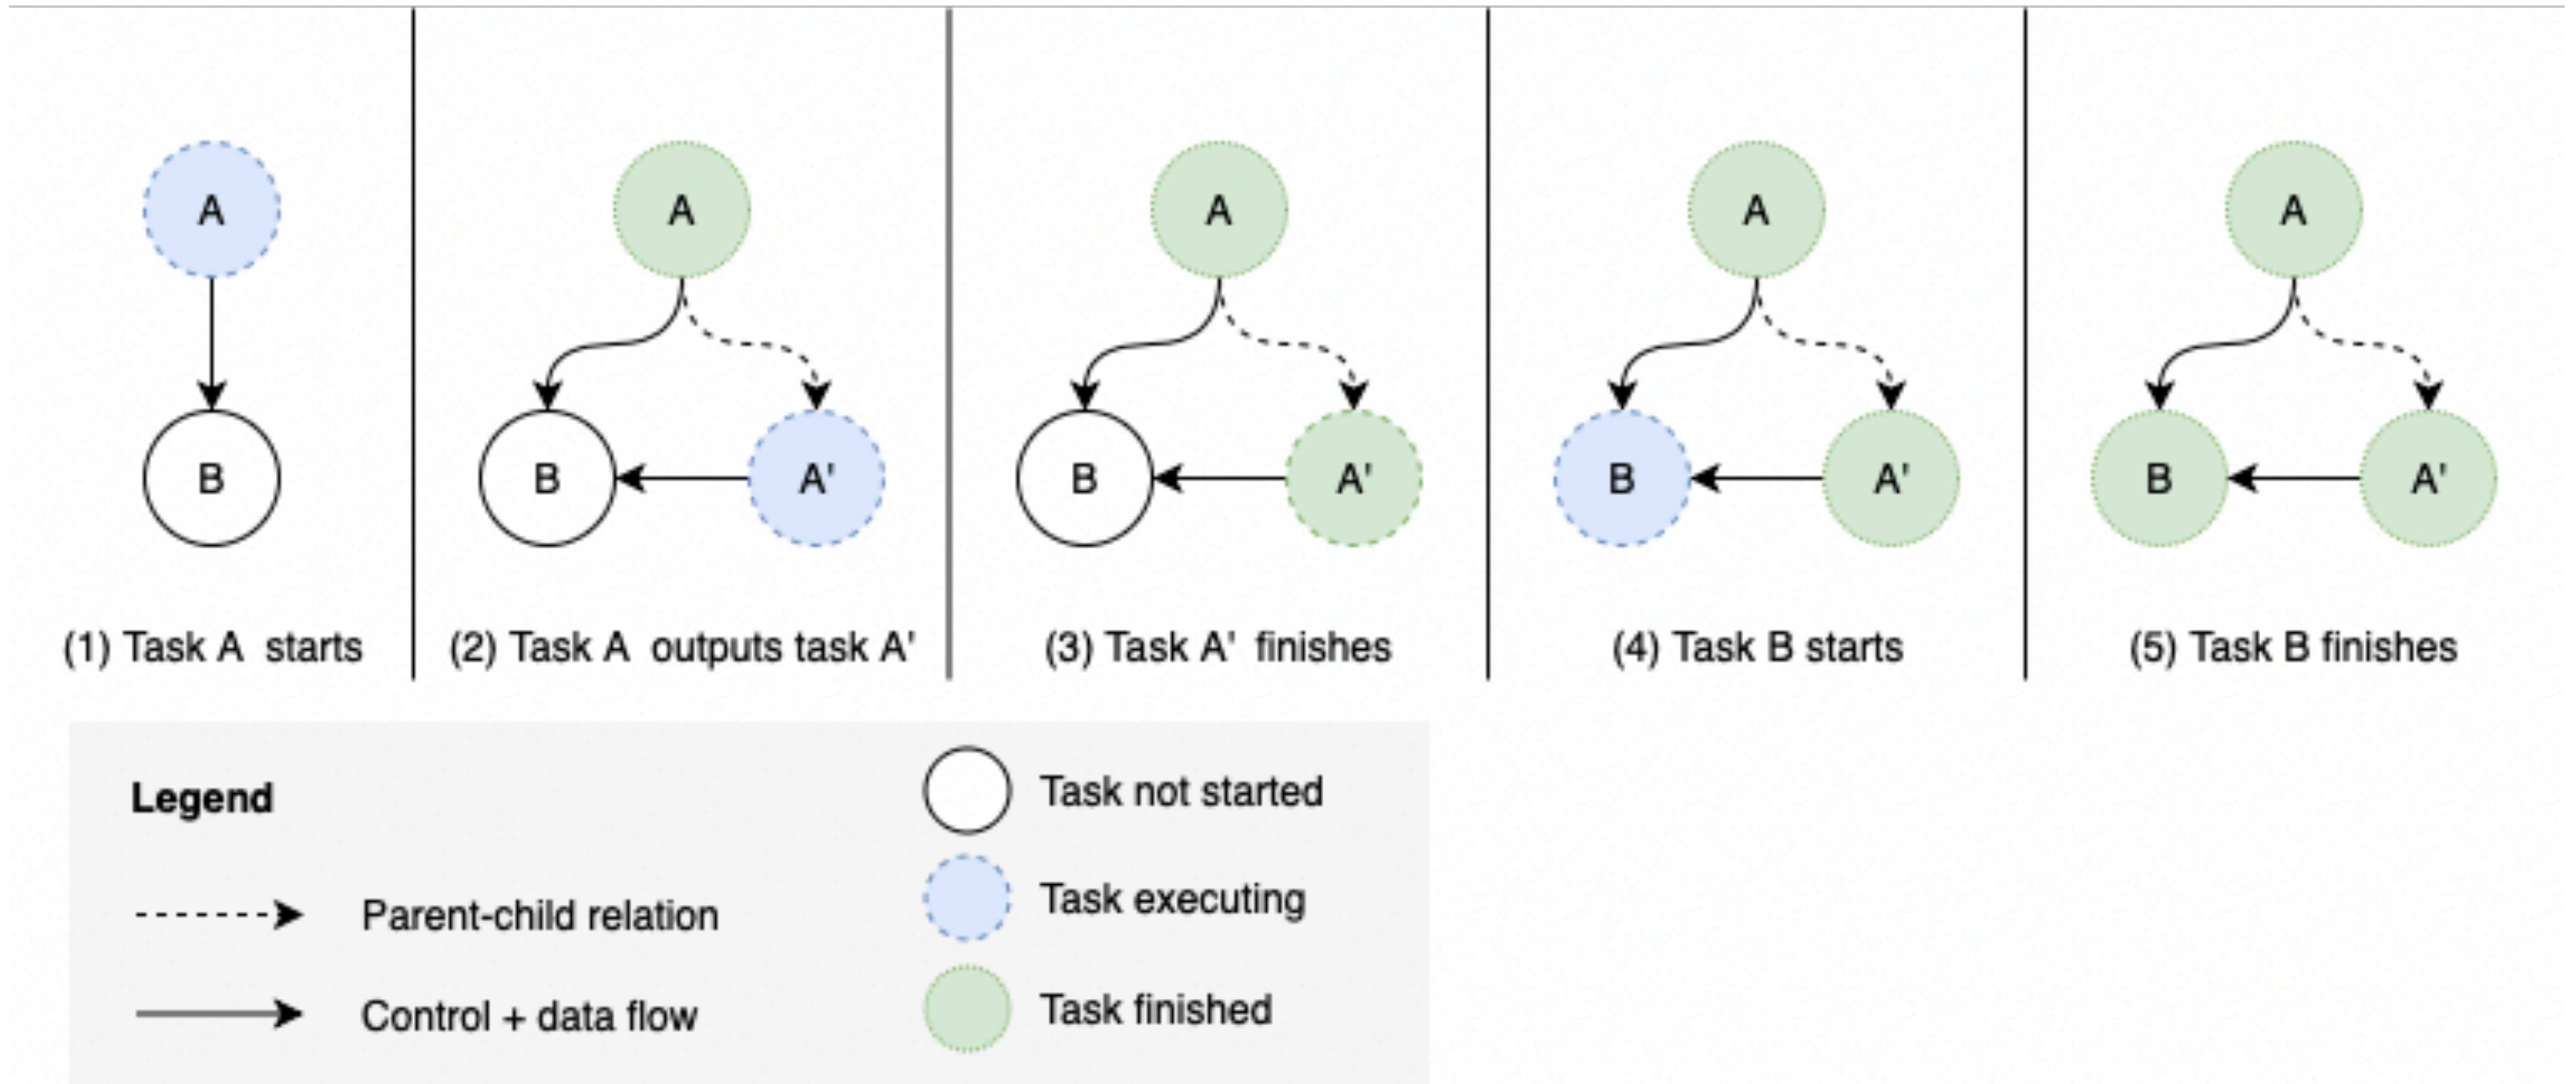


Task Run



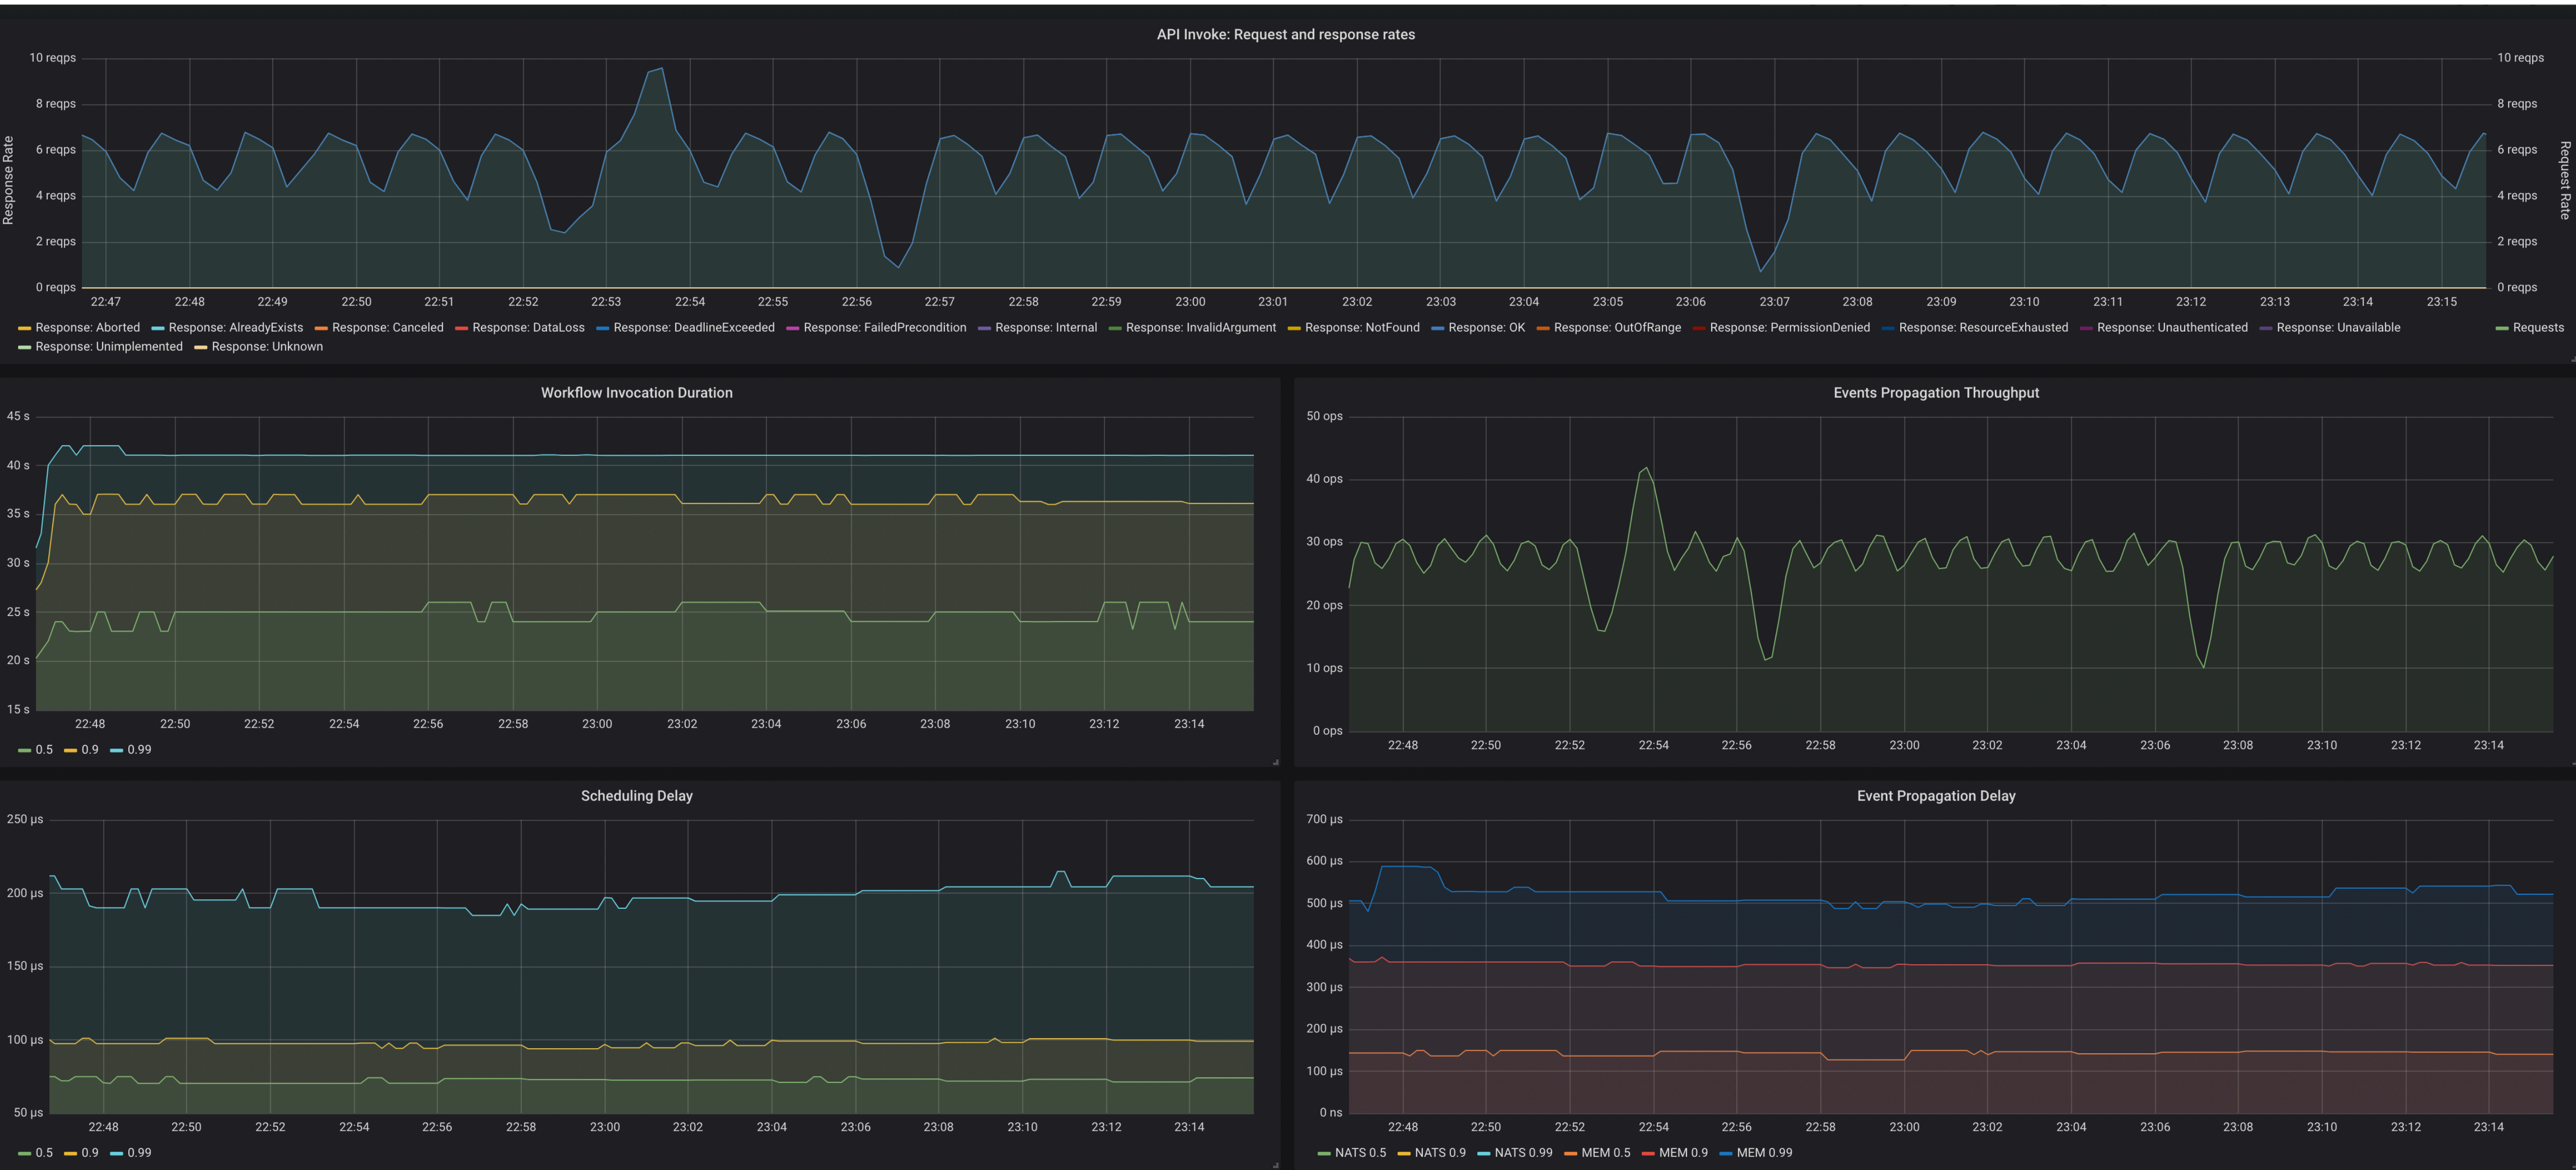


# SWL: dynamic workflow support






# Fission Workflows monitoring support




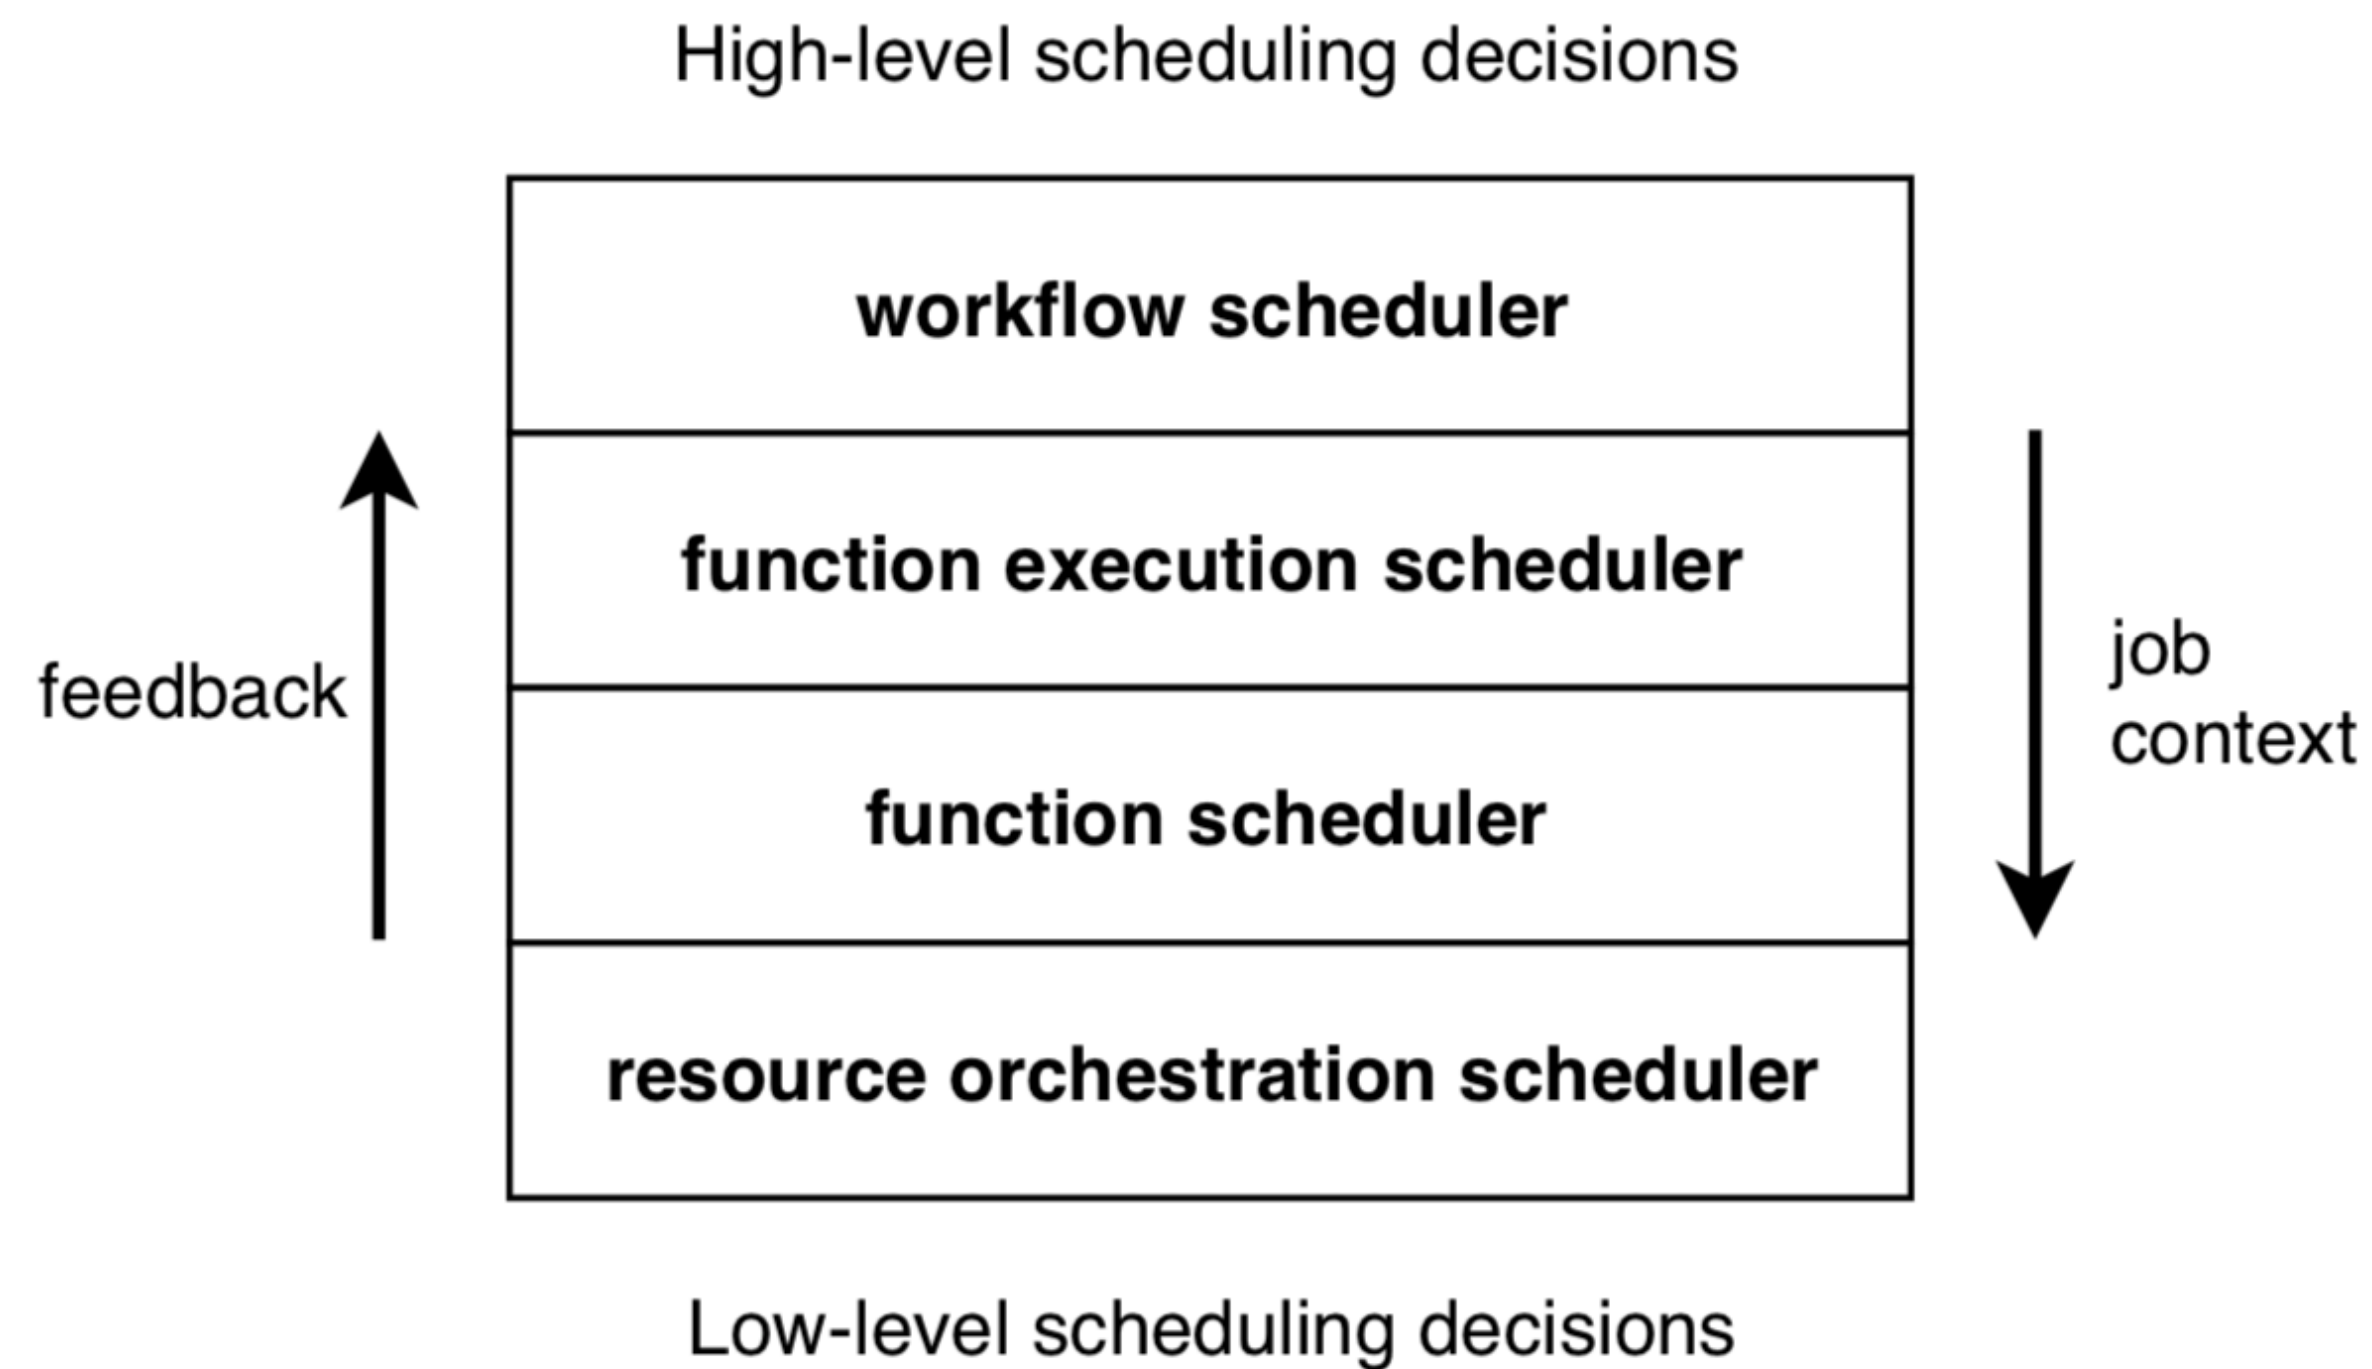


# Fission Workflows tracing support



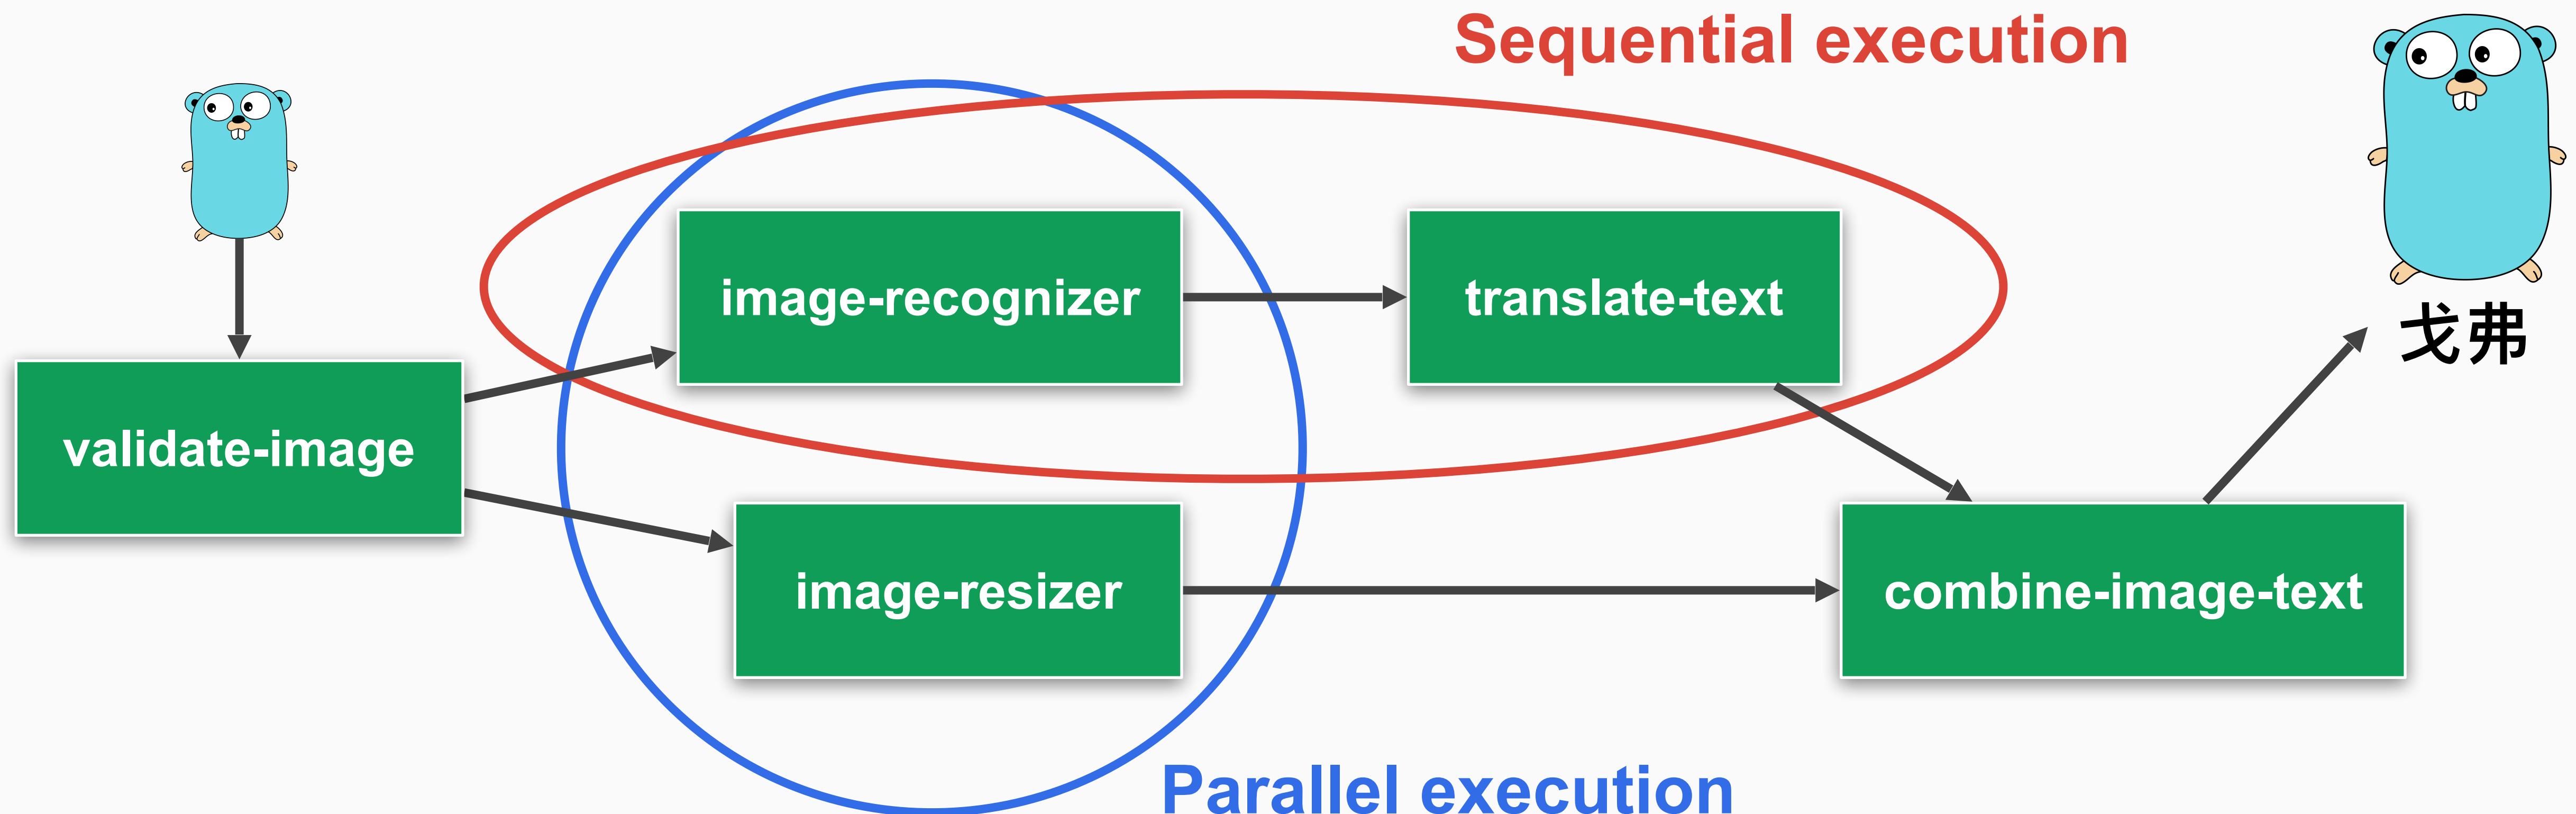


# Serverless scheduling architecture and policies






- Connect existing functions into complex function compositions
- Workflow engine takes care of the plumbing and provides fully monitorable, fault-tolerant function compositions with low overhead.





# Why serverless computing?

## Operational Model



Minimal operational logic

“Infinite” autoscaling

Built-in tooling: monitoring, tracing, health checking, etc.

## Cost Model




Pay based on usage

No upfront/periodic costs

Granular billing

## Development Model



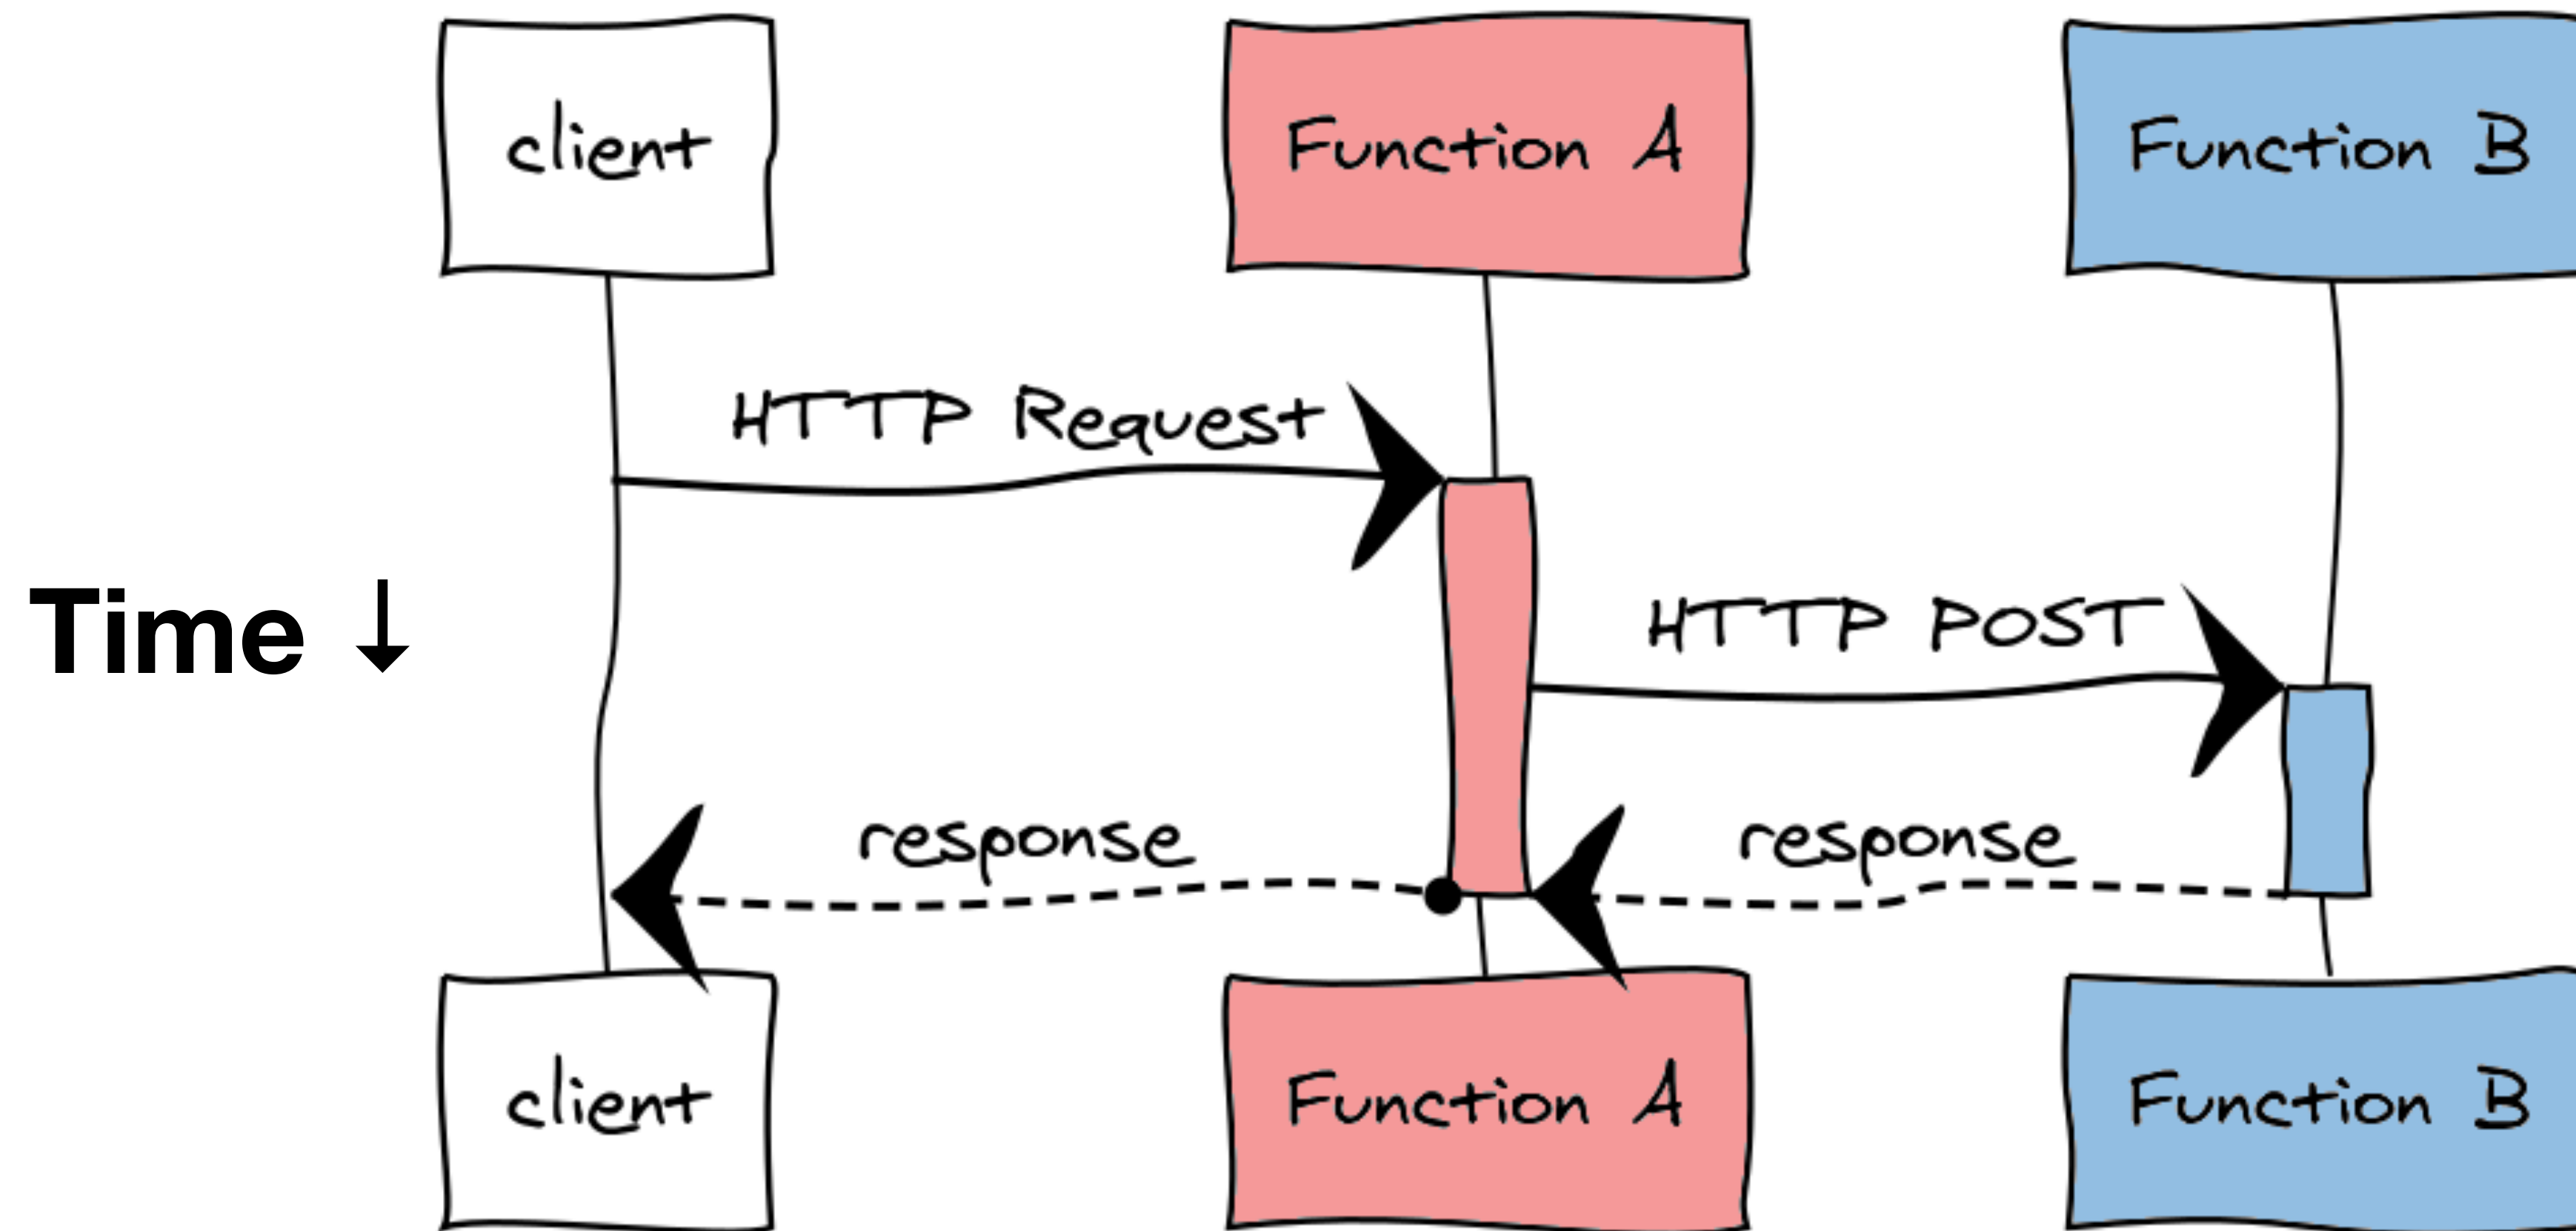
High-level abstractions

Pre-provided integrations

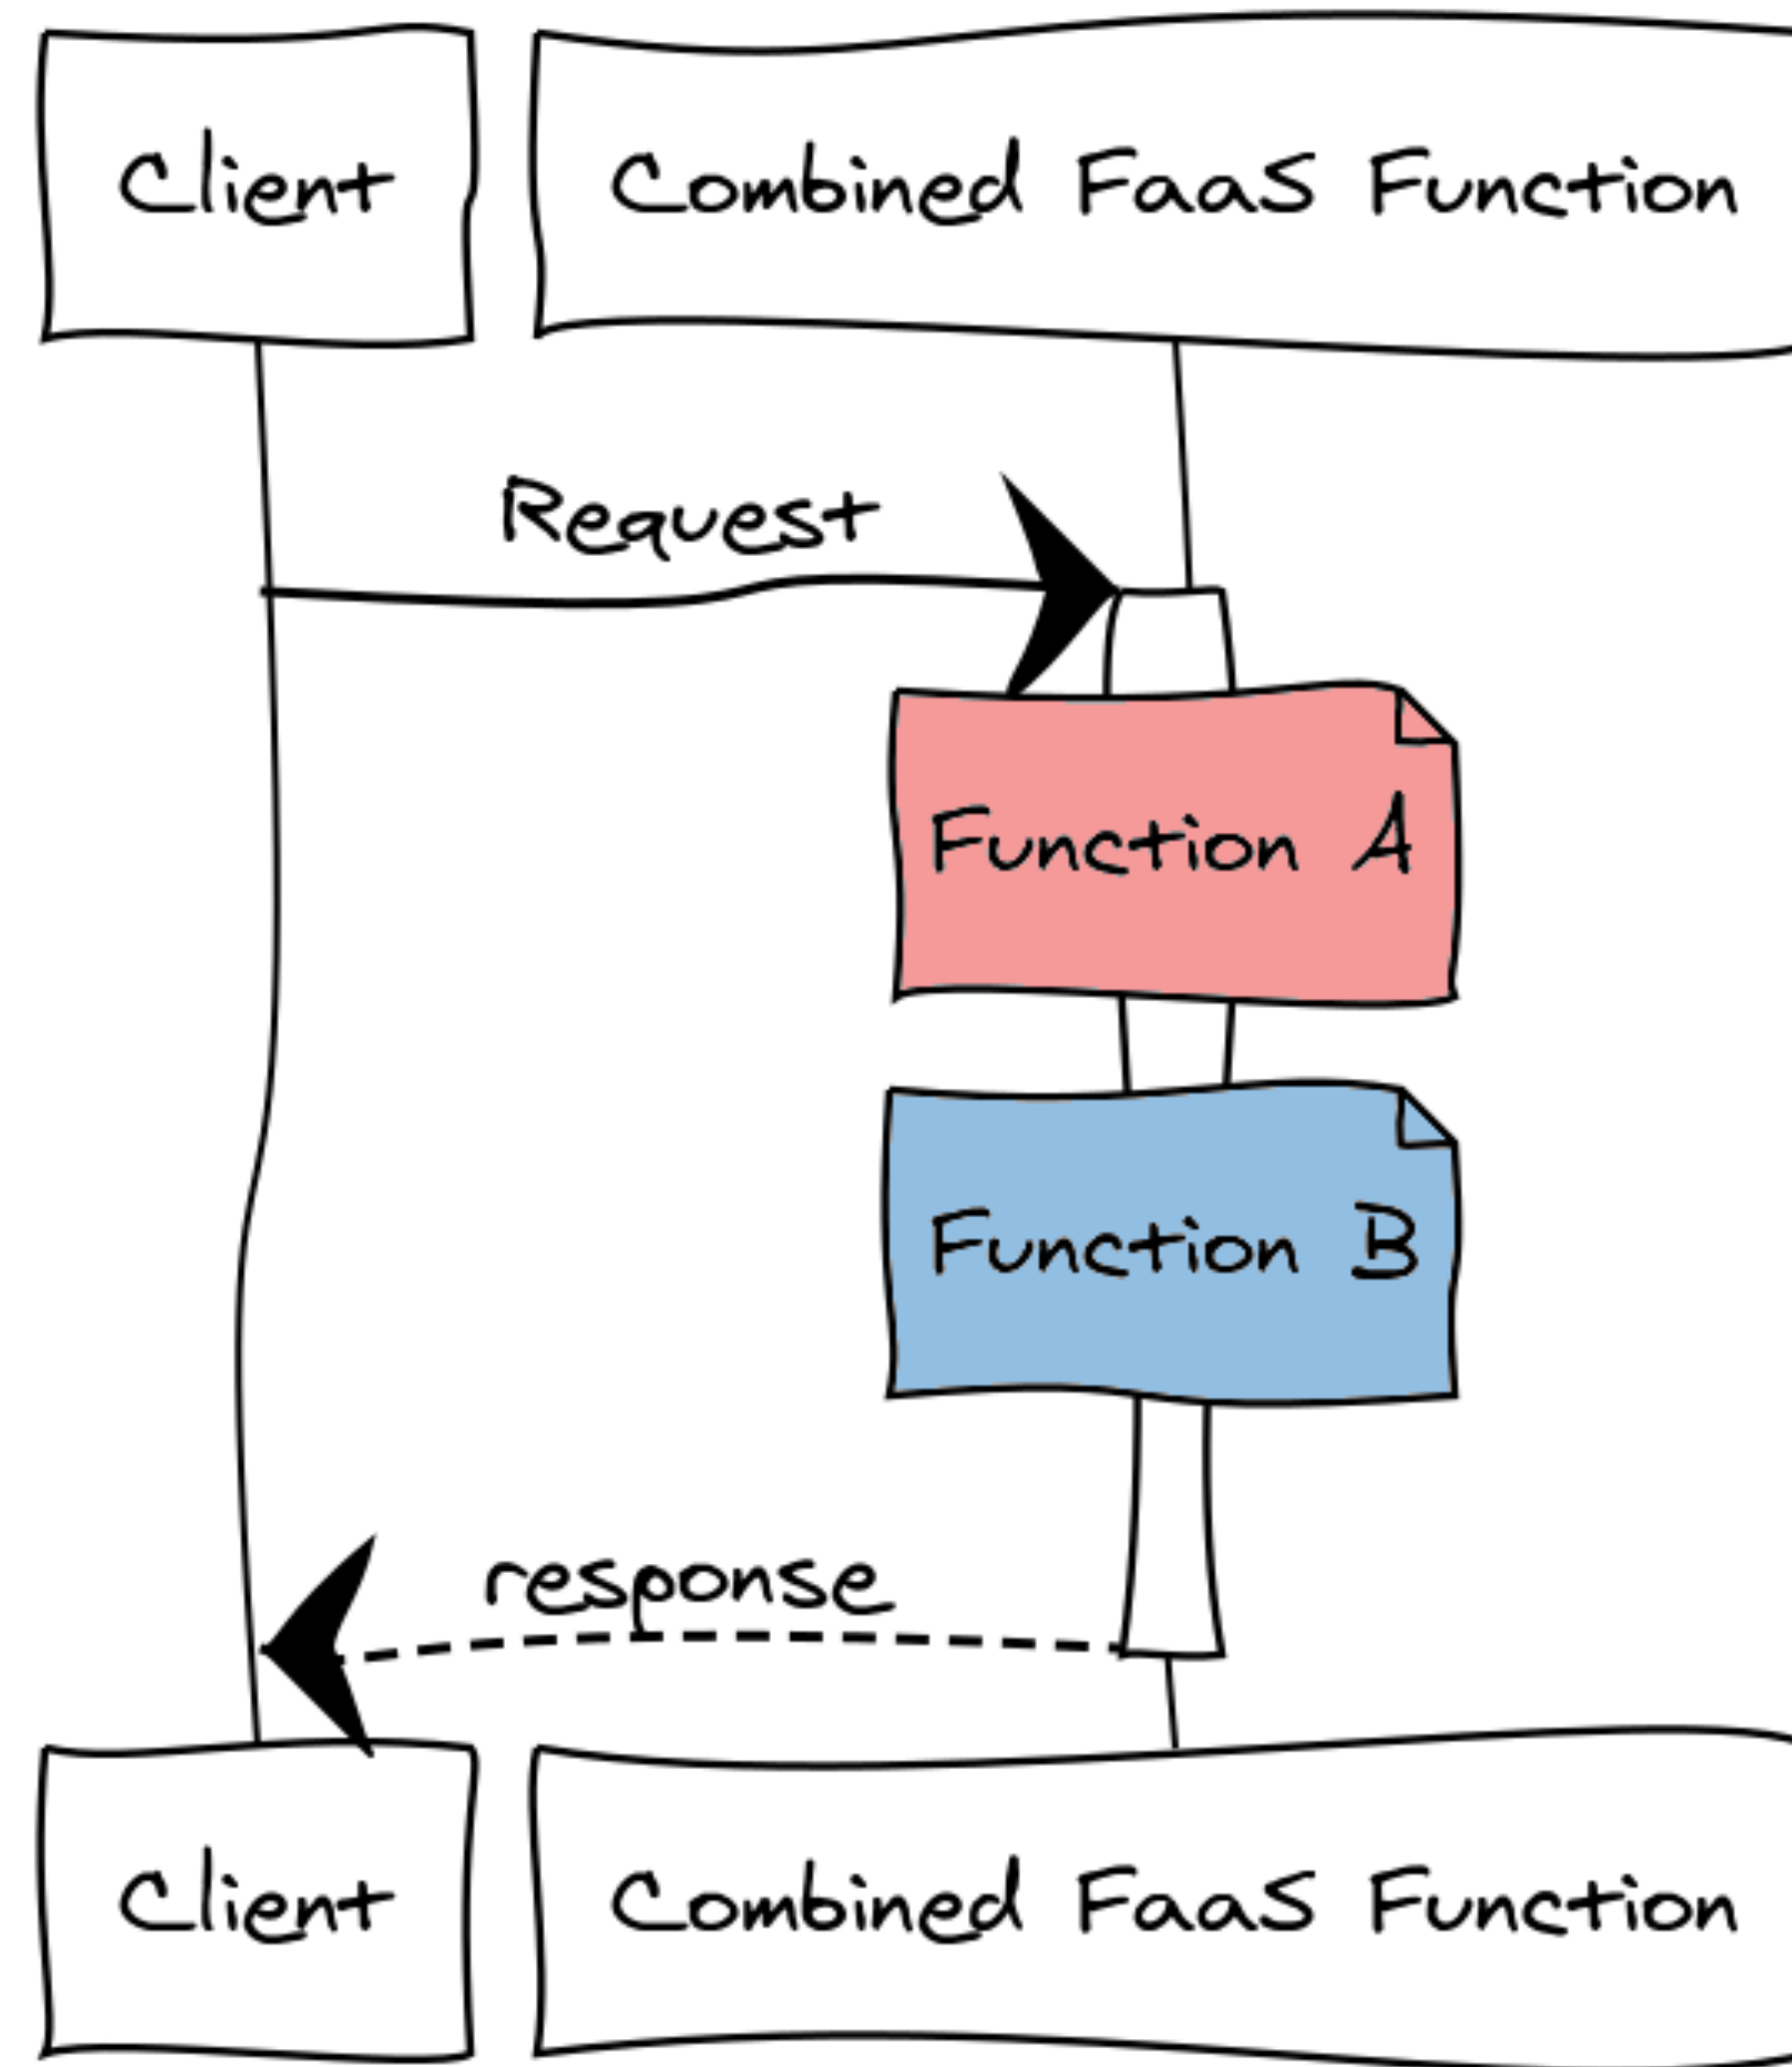
Language-agnostic



# Direct composition

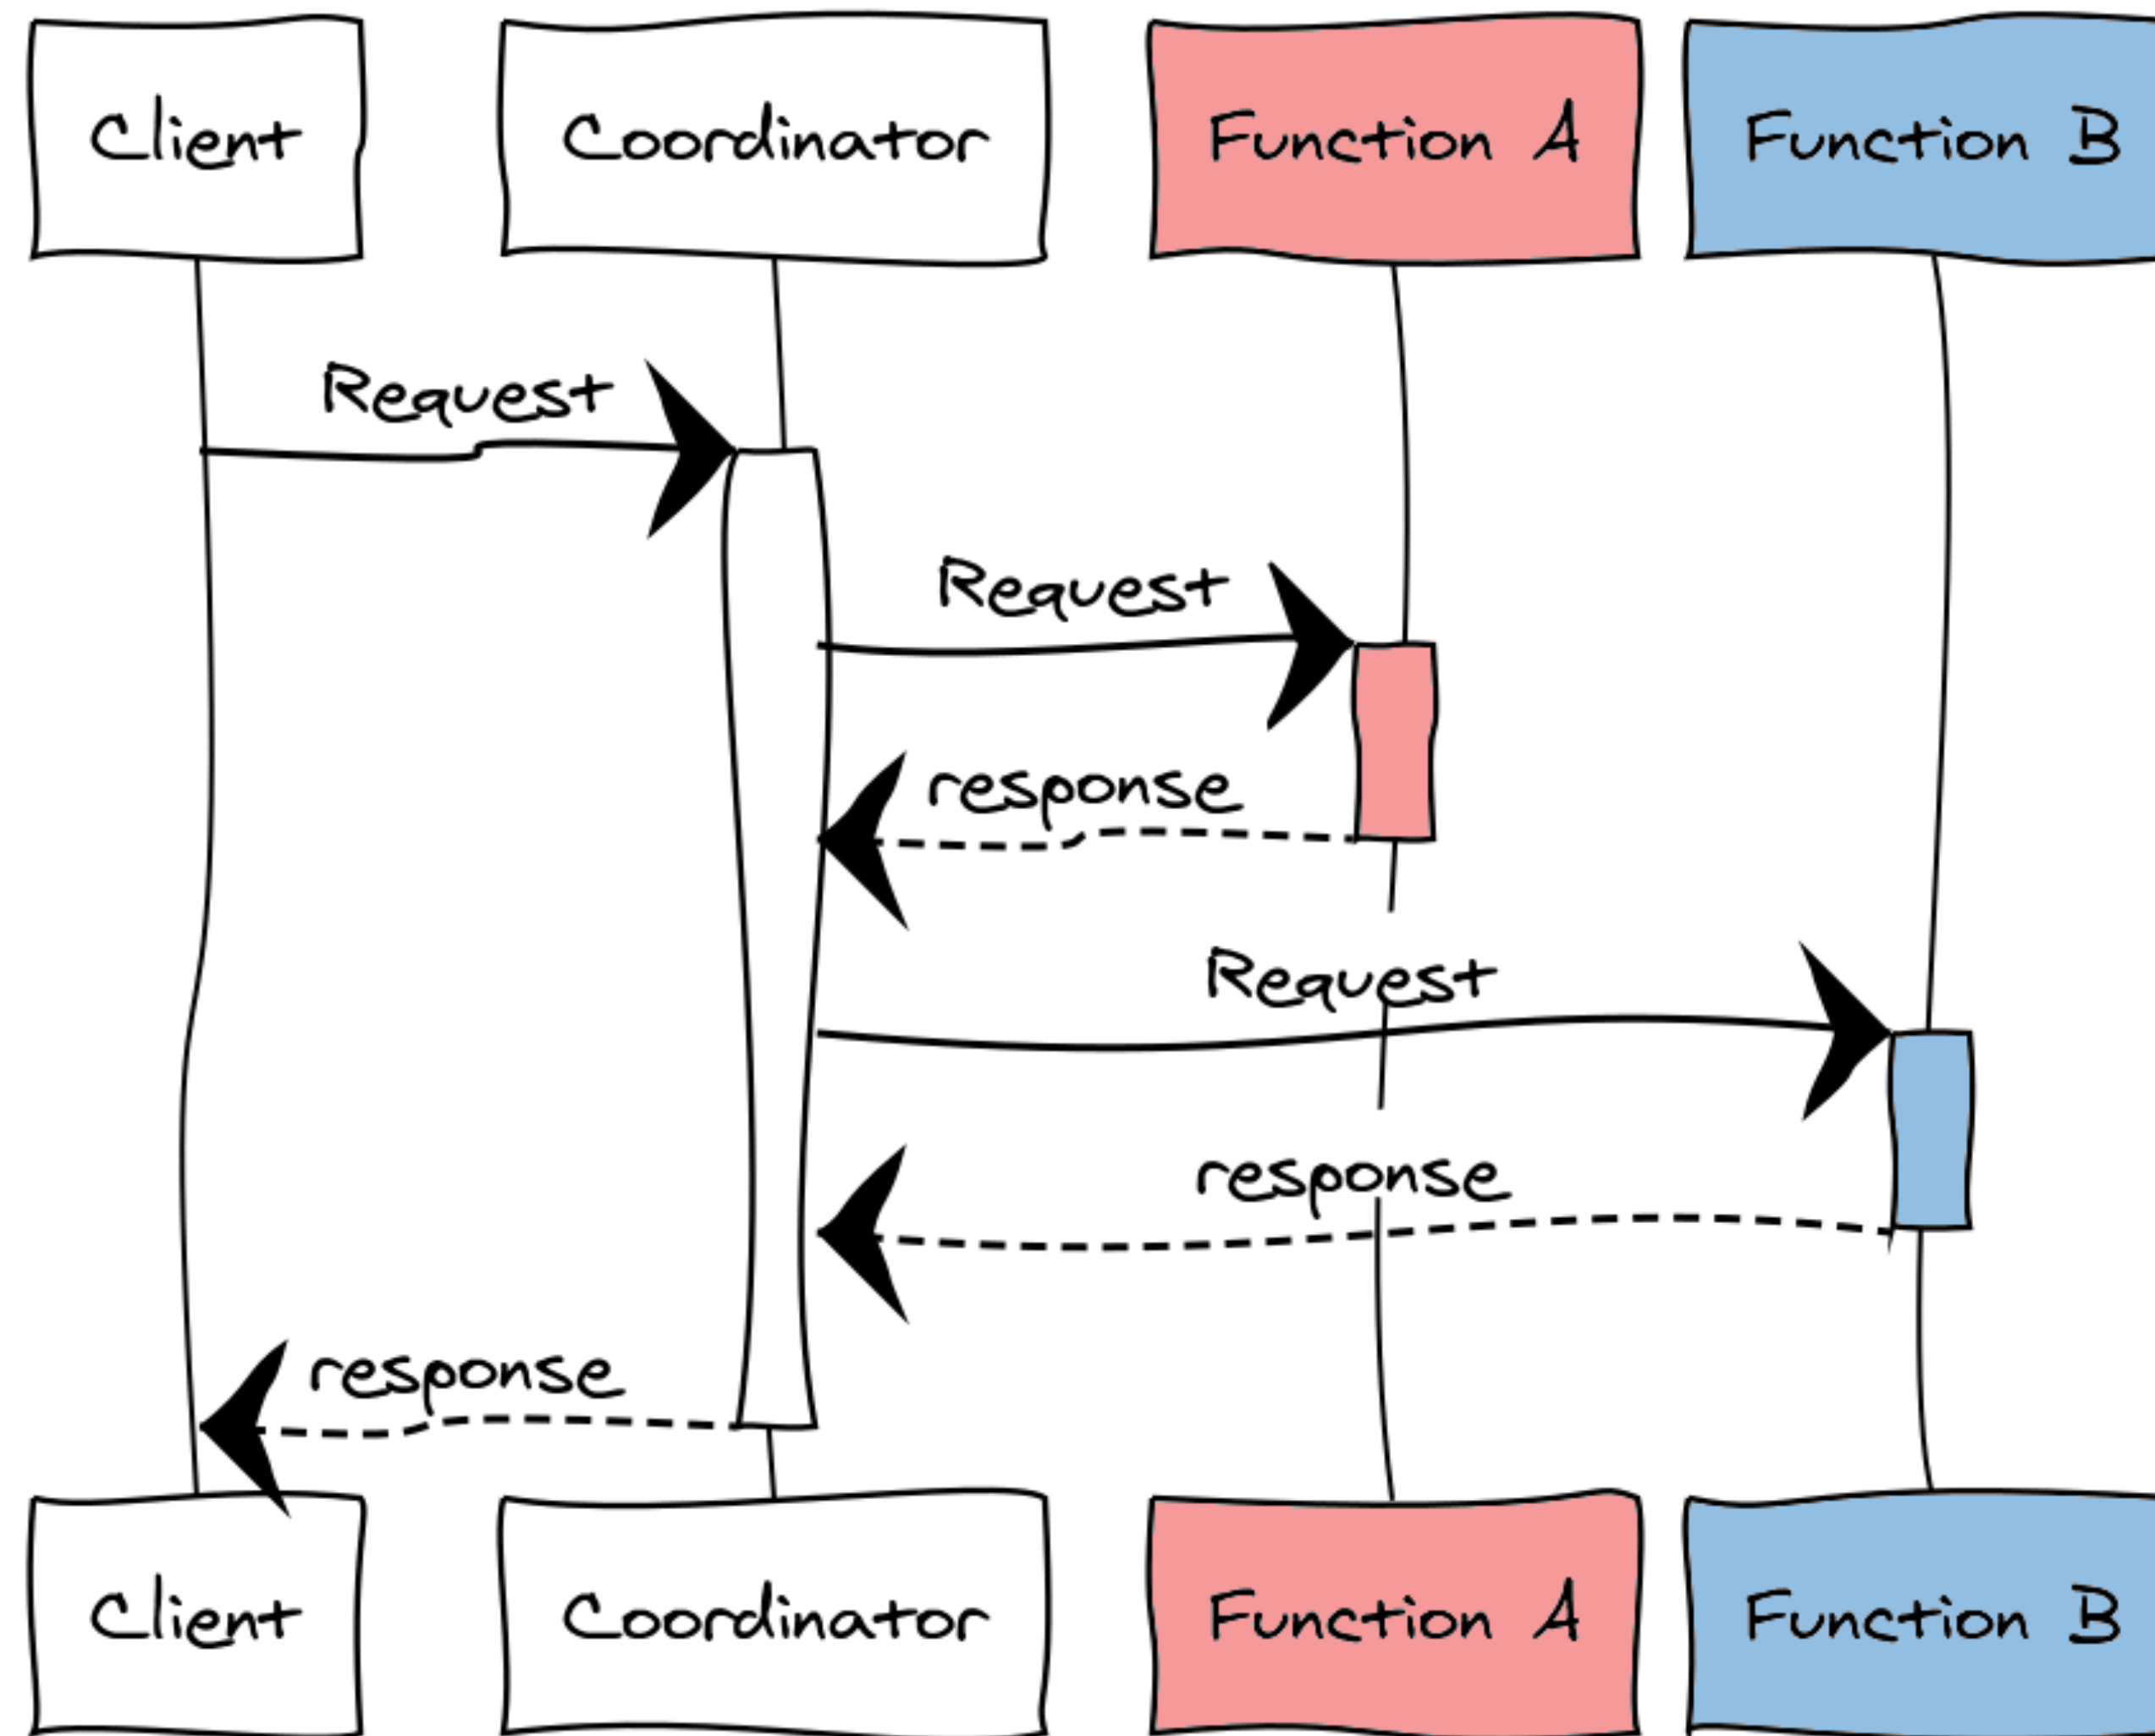


# Compiled composition

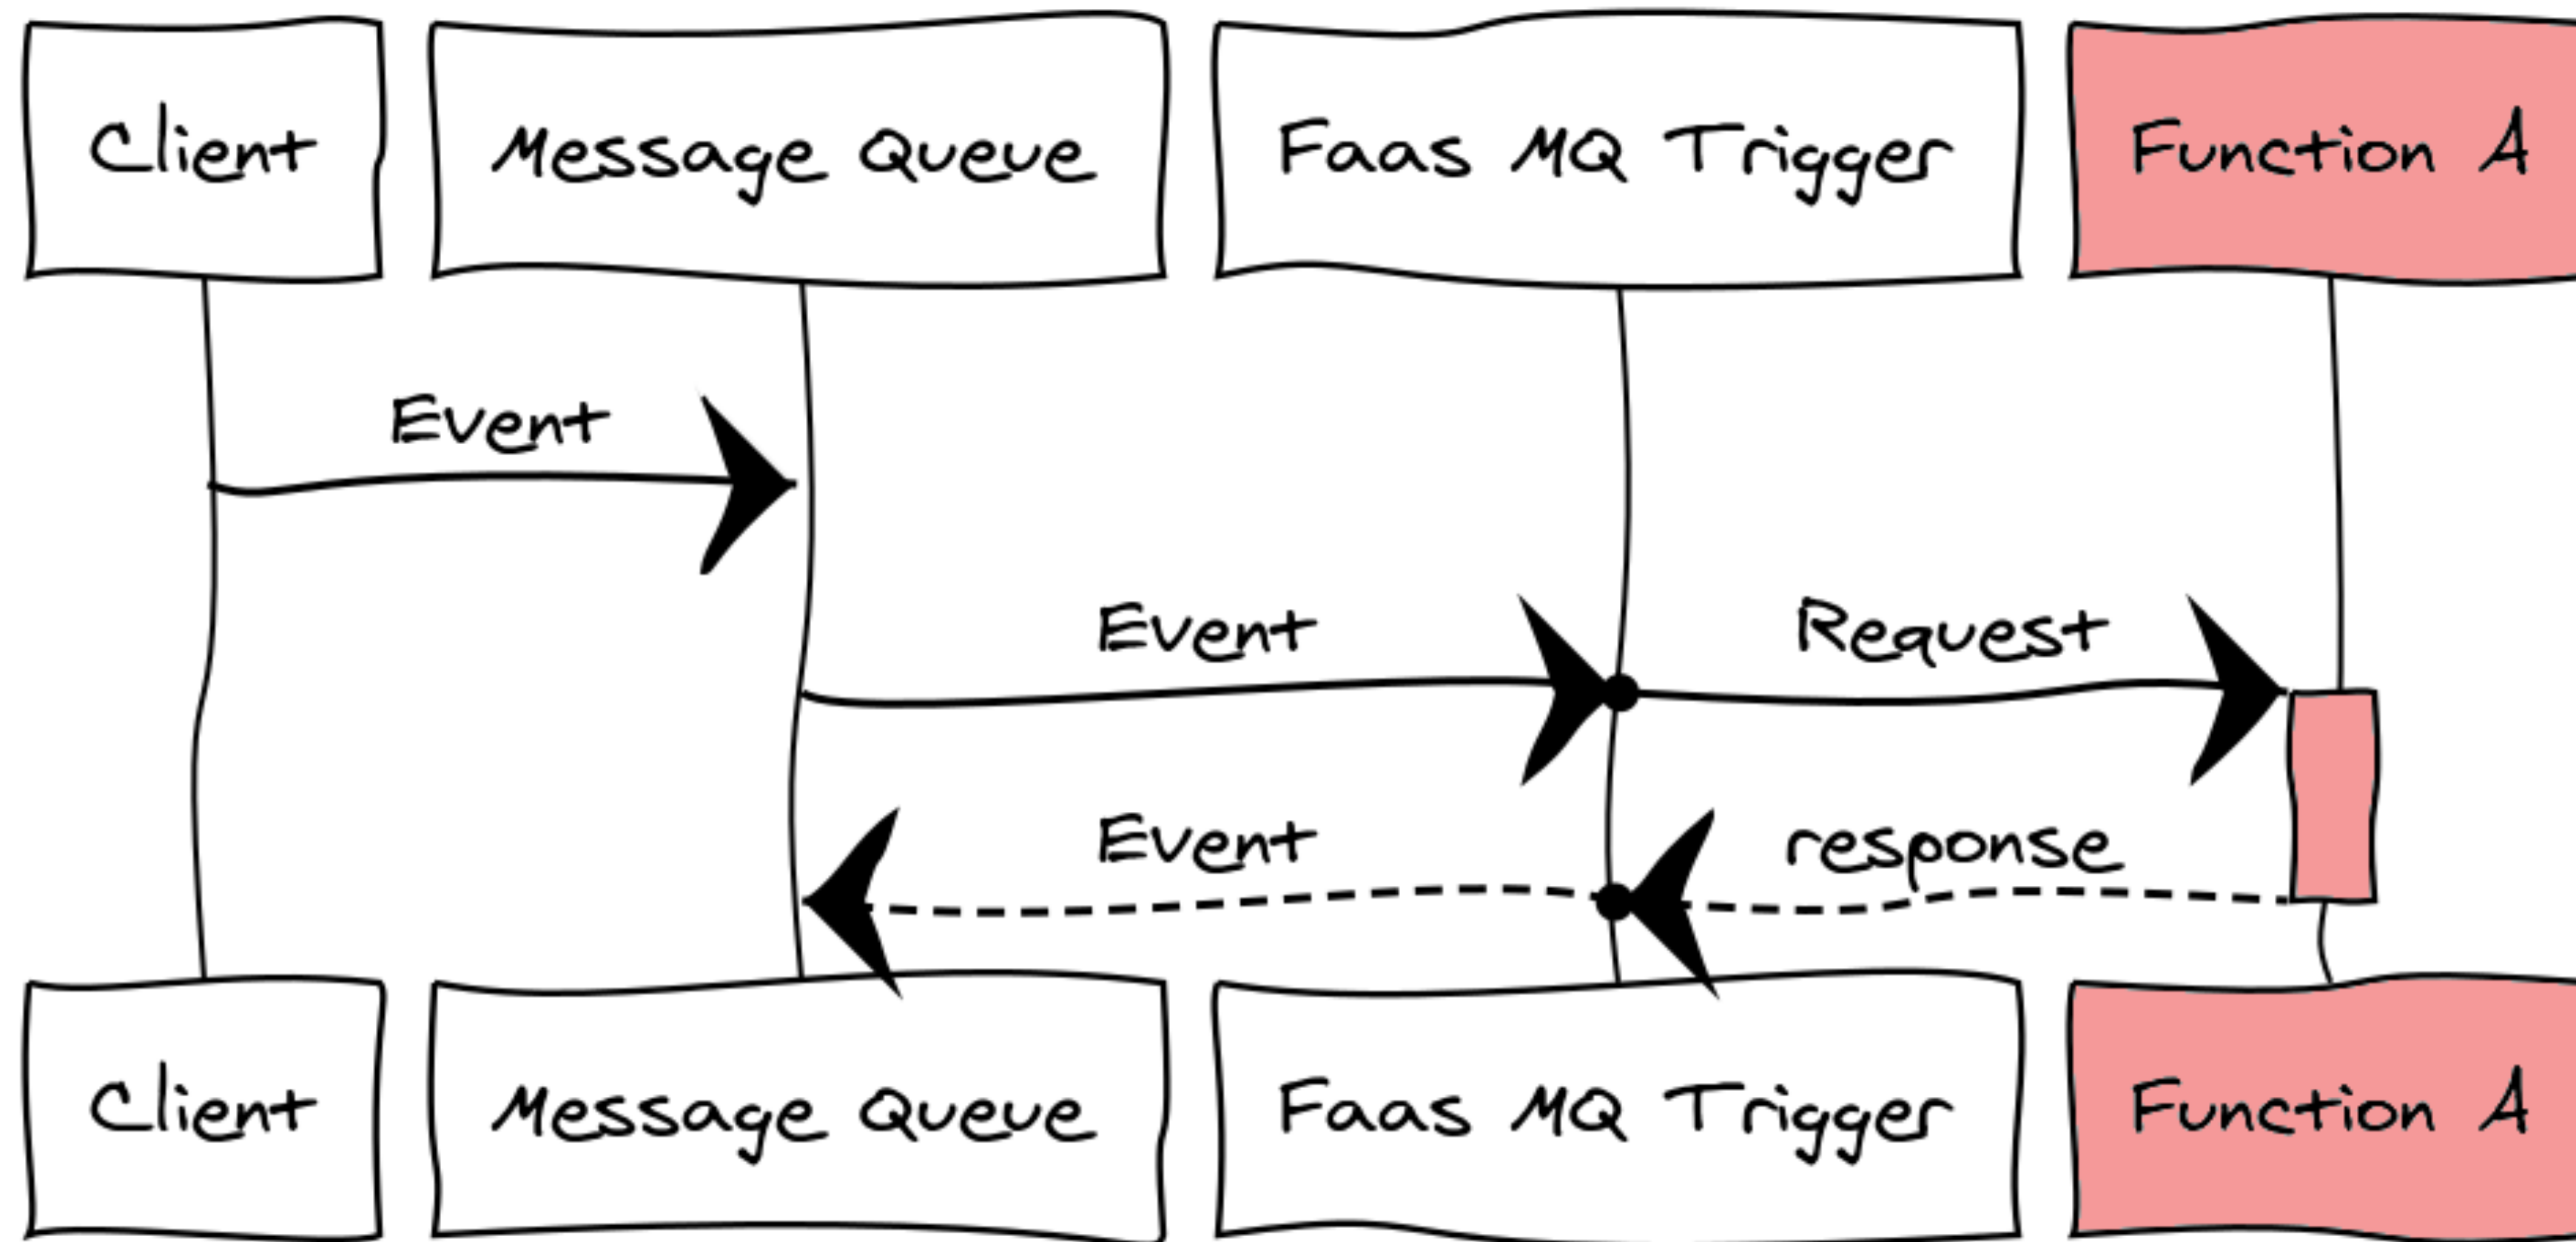




# Coordinator-based composition



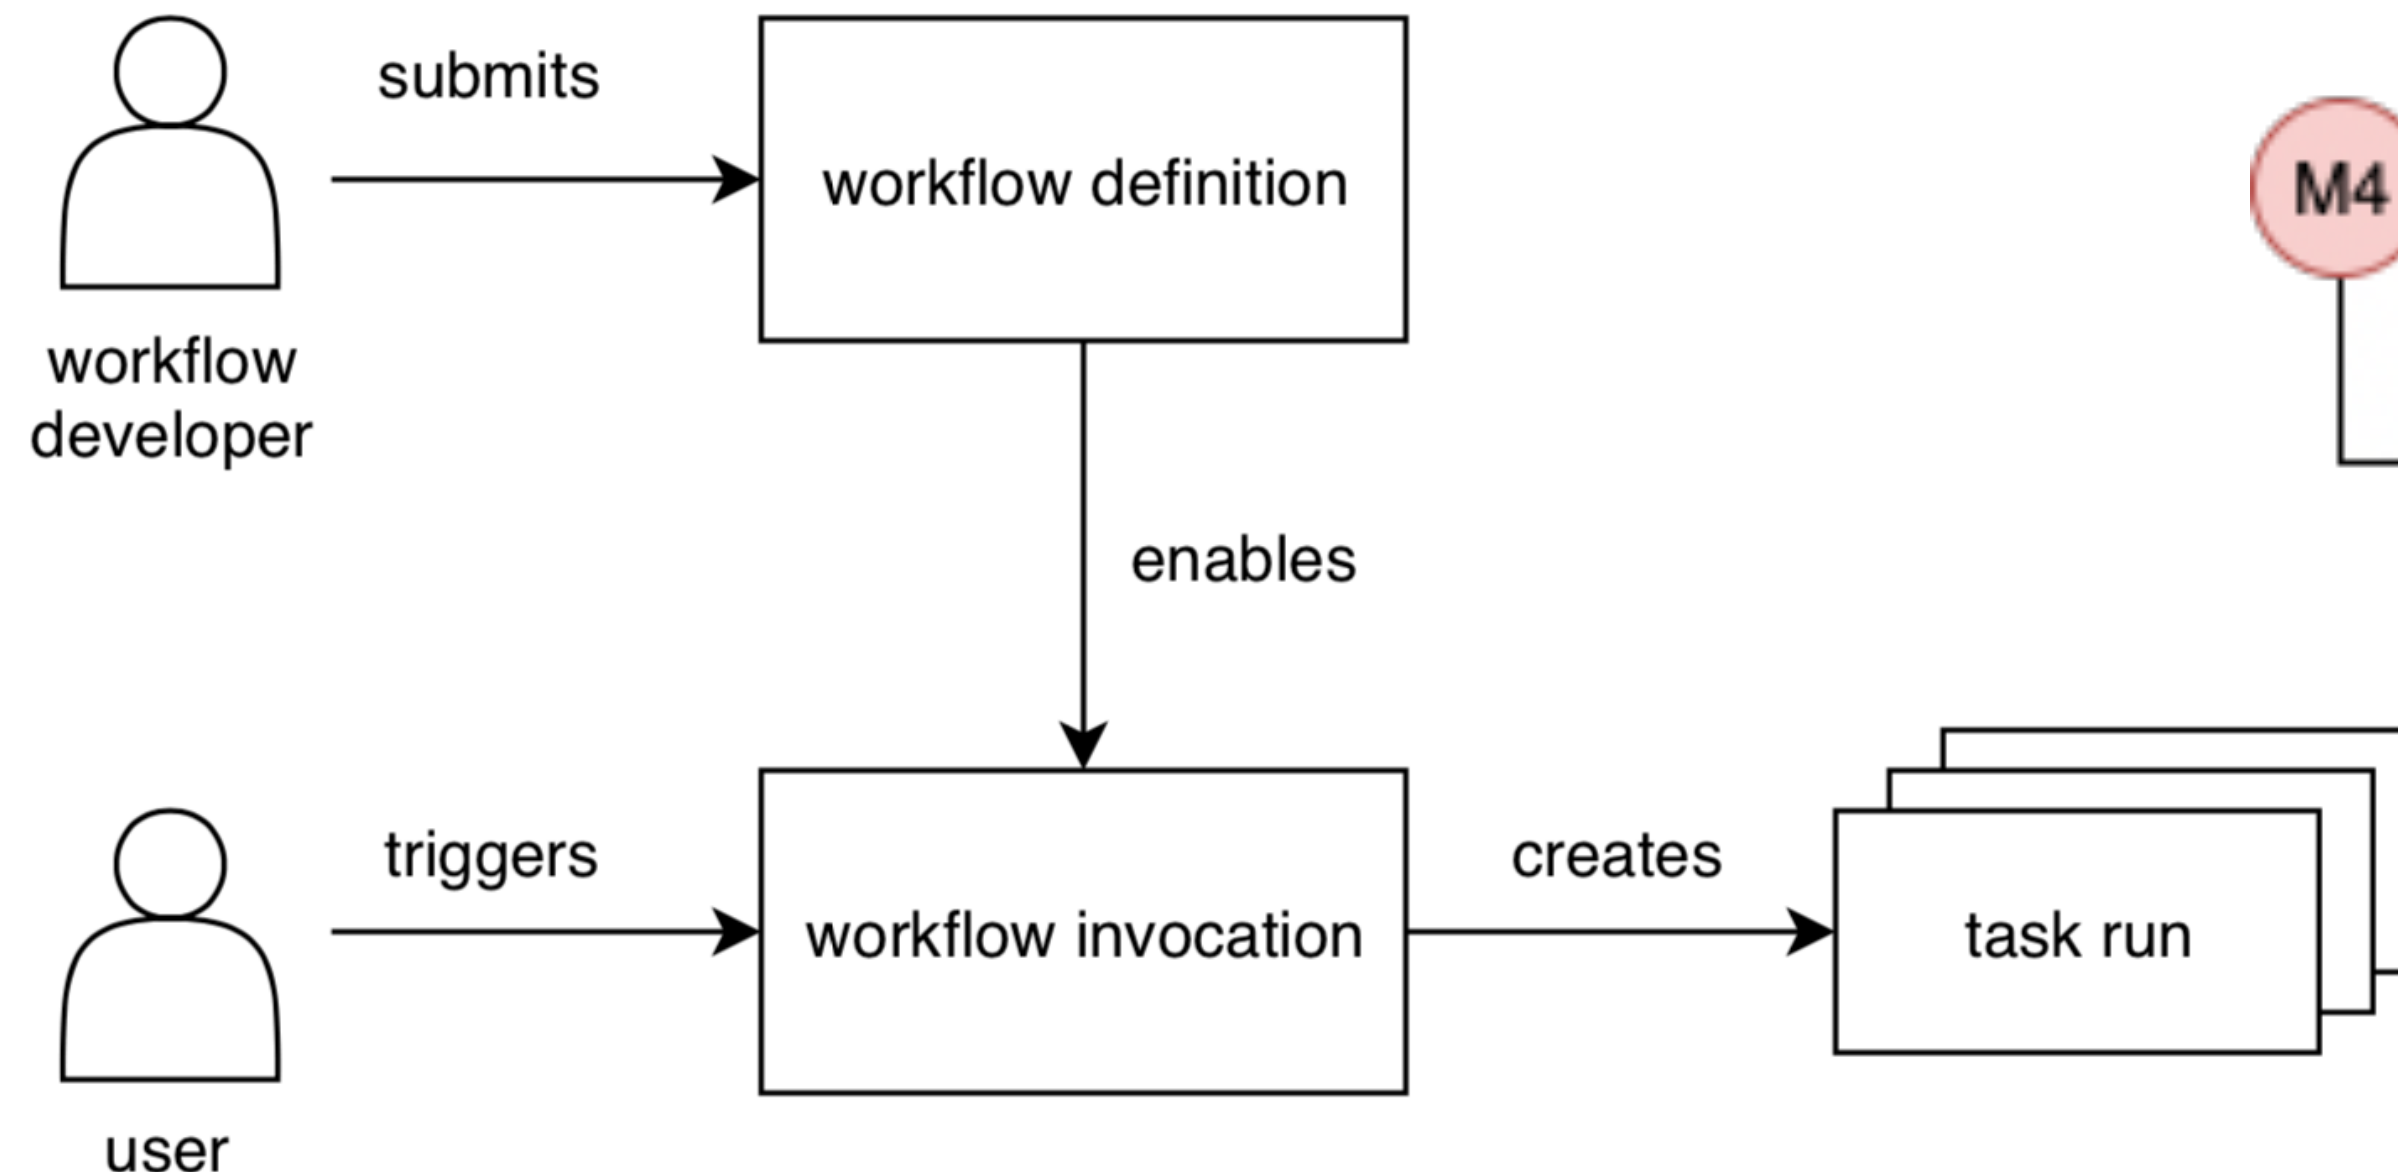
# Event-driven composition



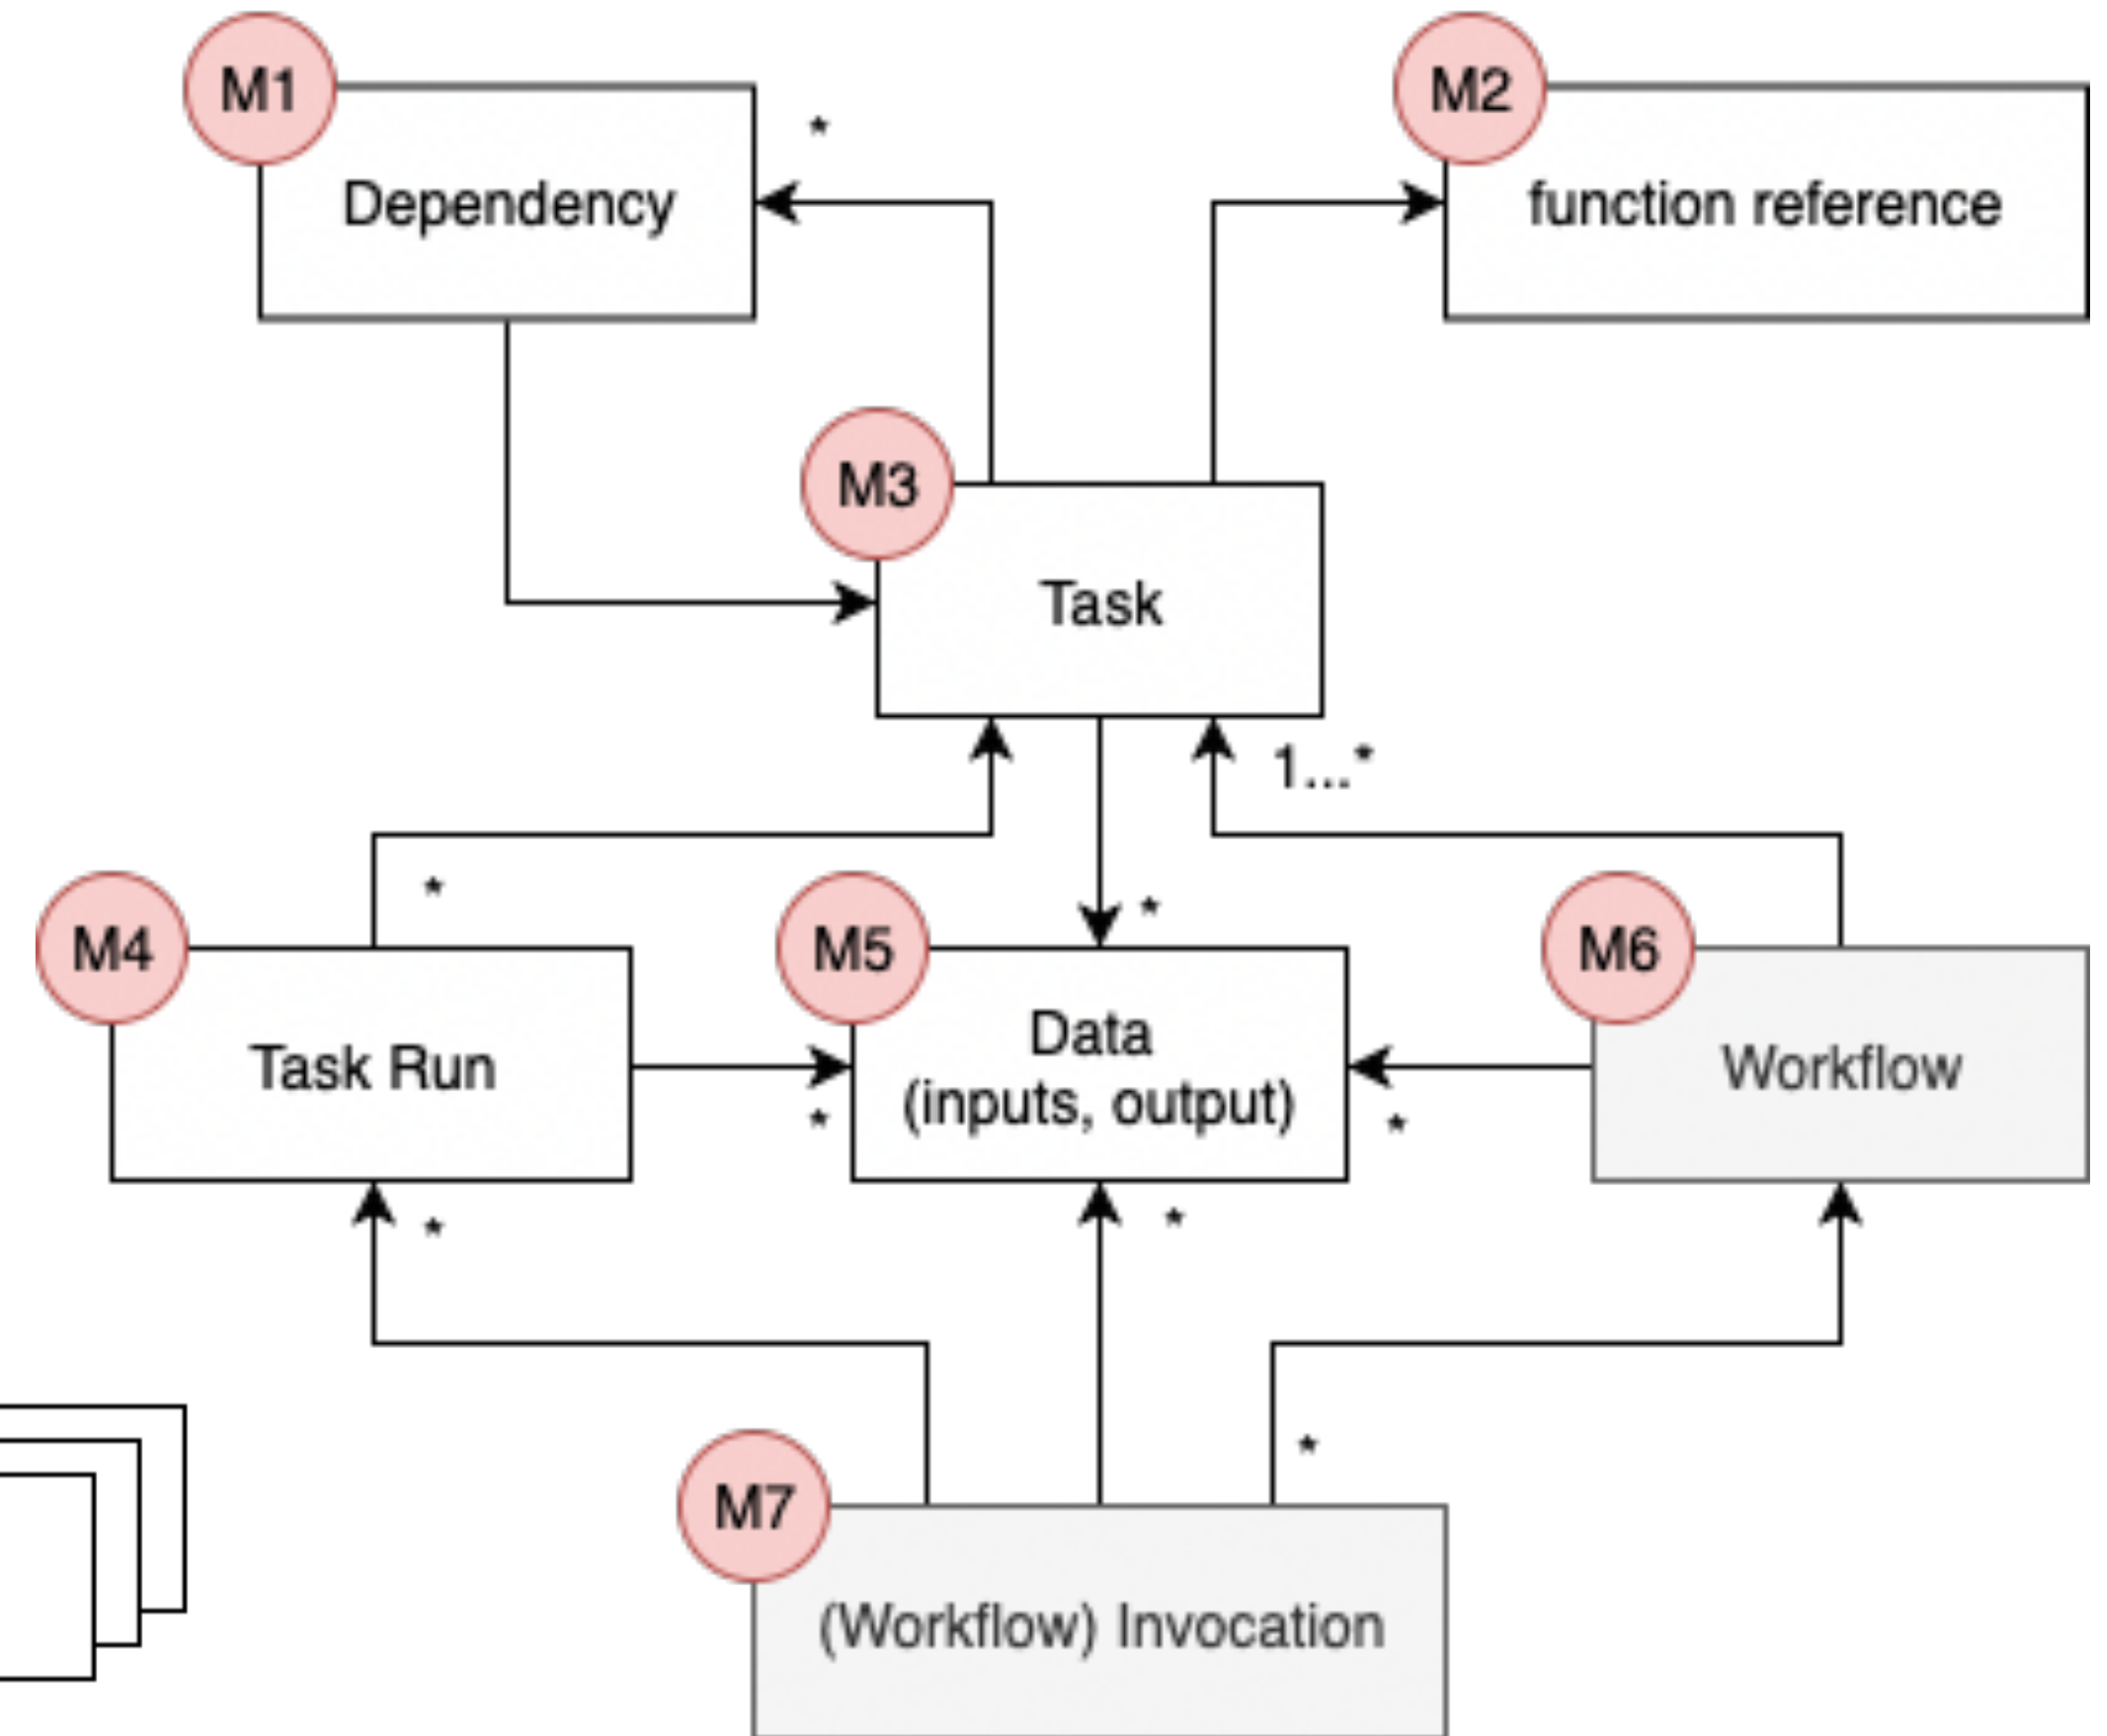


# Serverless Workflow Language (SWL)

- Execution-level workflow language
- Minimal constructs
- Supports complex control flows



**SWL execution model**



**SWL data model**

# SWL-YAML: a reference implementation of SWL

```
1  apiVersion: 1
2  output: WhaleWithFortune
3  tasks:
4    GenerateFortune:
5      run: fortune
6
7    WhaleWithFortune:
8      run: whalesay
9      inputs:
10       body: "{ output('GenerateFortune') }"
11      requires:
12       - GenerateFortune
```

- Uses a declarative format
- Follows syntax of state-of-the-art WMSs
- Supports JavaScript-like expressions