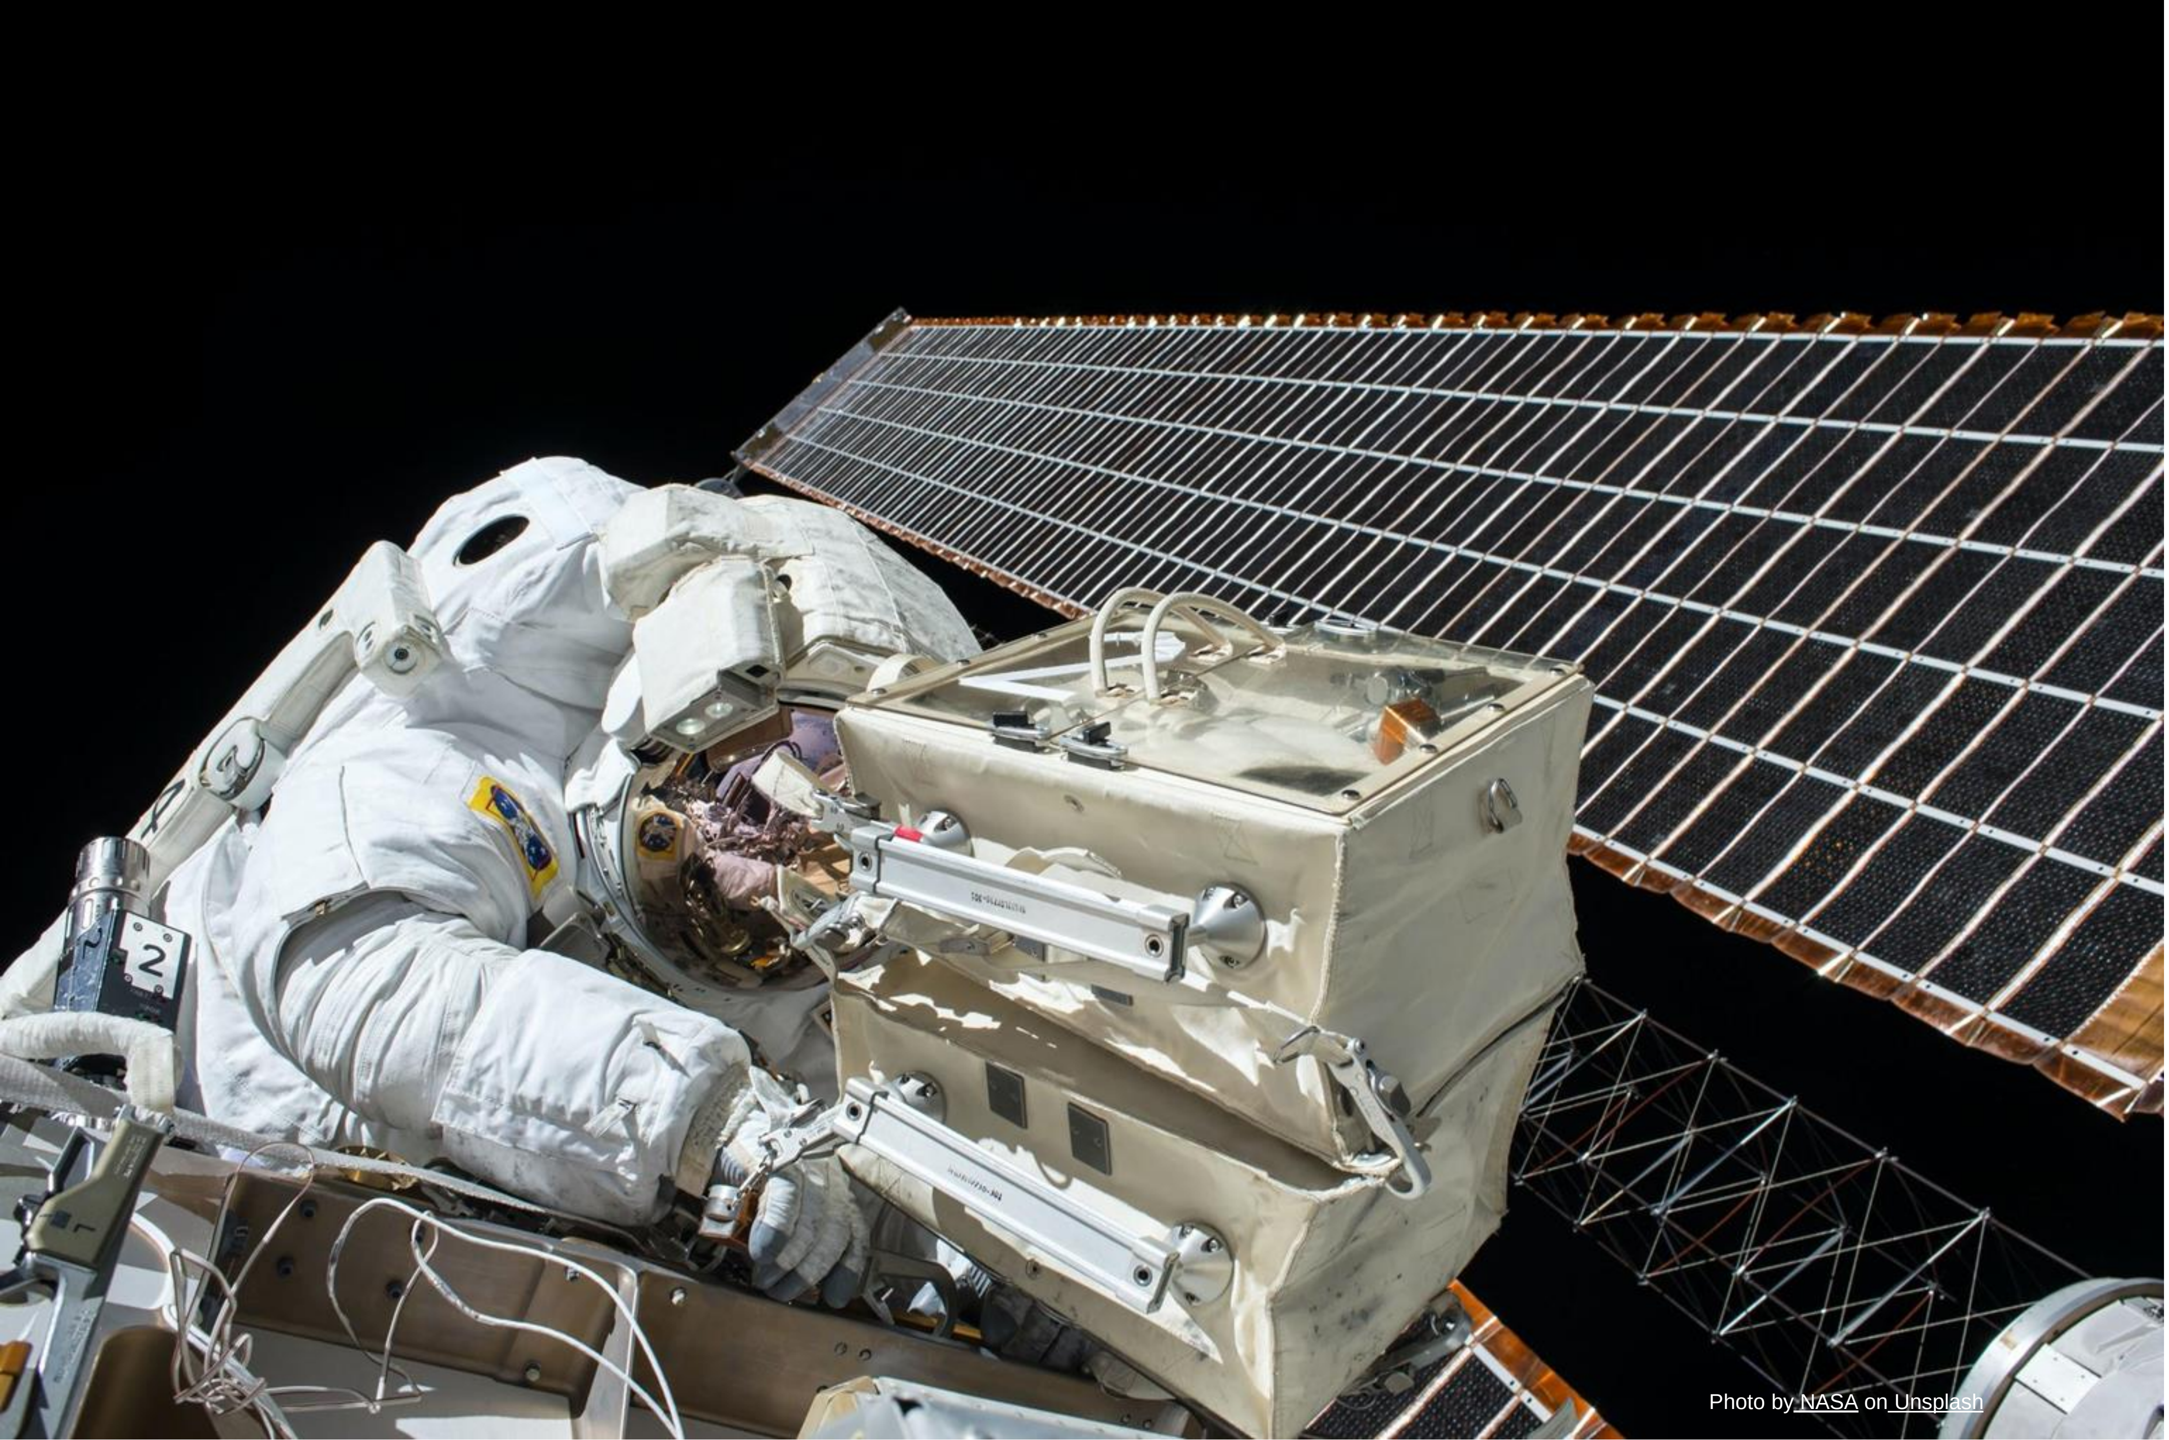

Photo by NASA on Unsplash

CAMUNDA
CON
2020.2

Chaos Engineering Meets Zeebe

@ChristopherZell

Definition: Chaos Engineering

Chaos Engineering is the discipline of experimenting on a system in order to build confidence in the system's capability to withstand turbulent conditions in production.

principlesofchaos.org

Chaos Engineering

Principles of Chaos define experiment as:

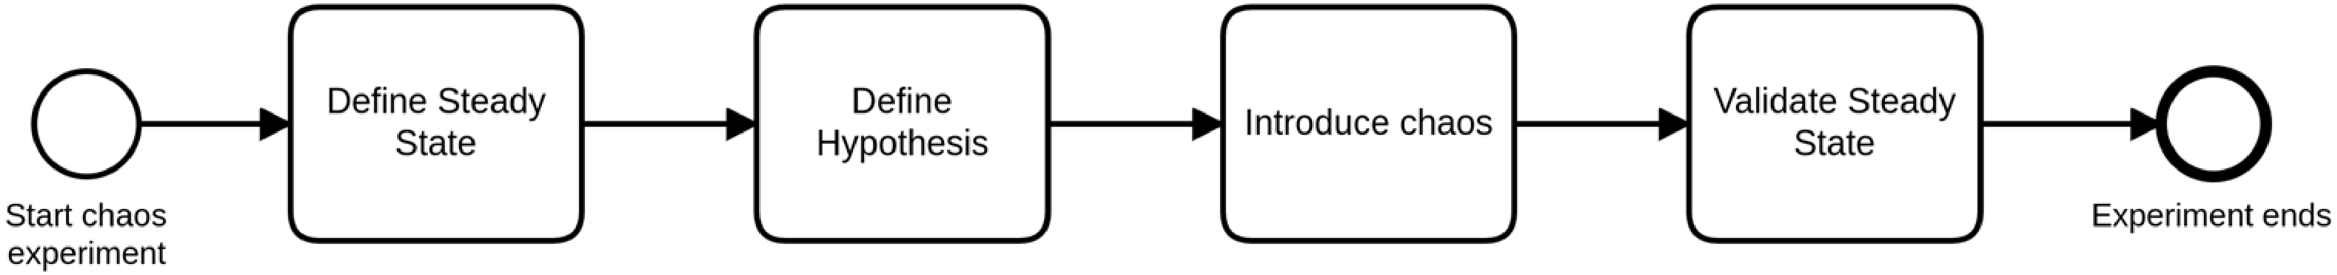

Chaos Engineering

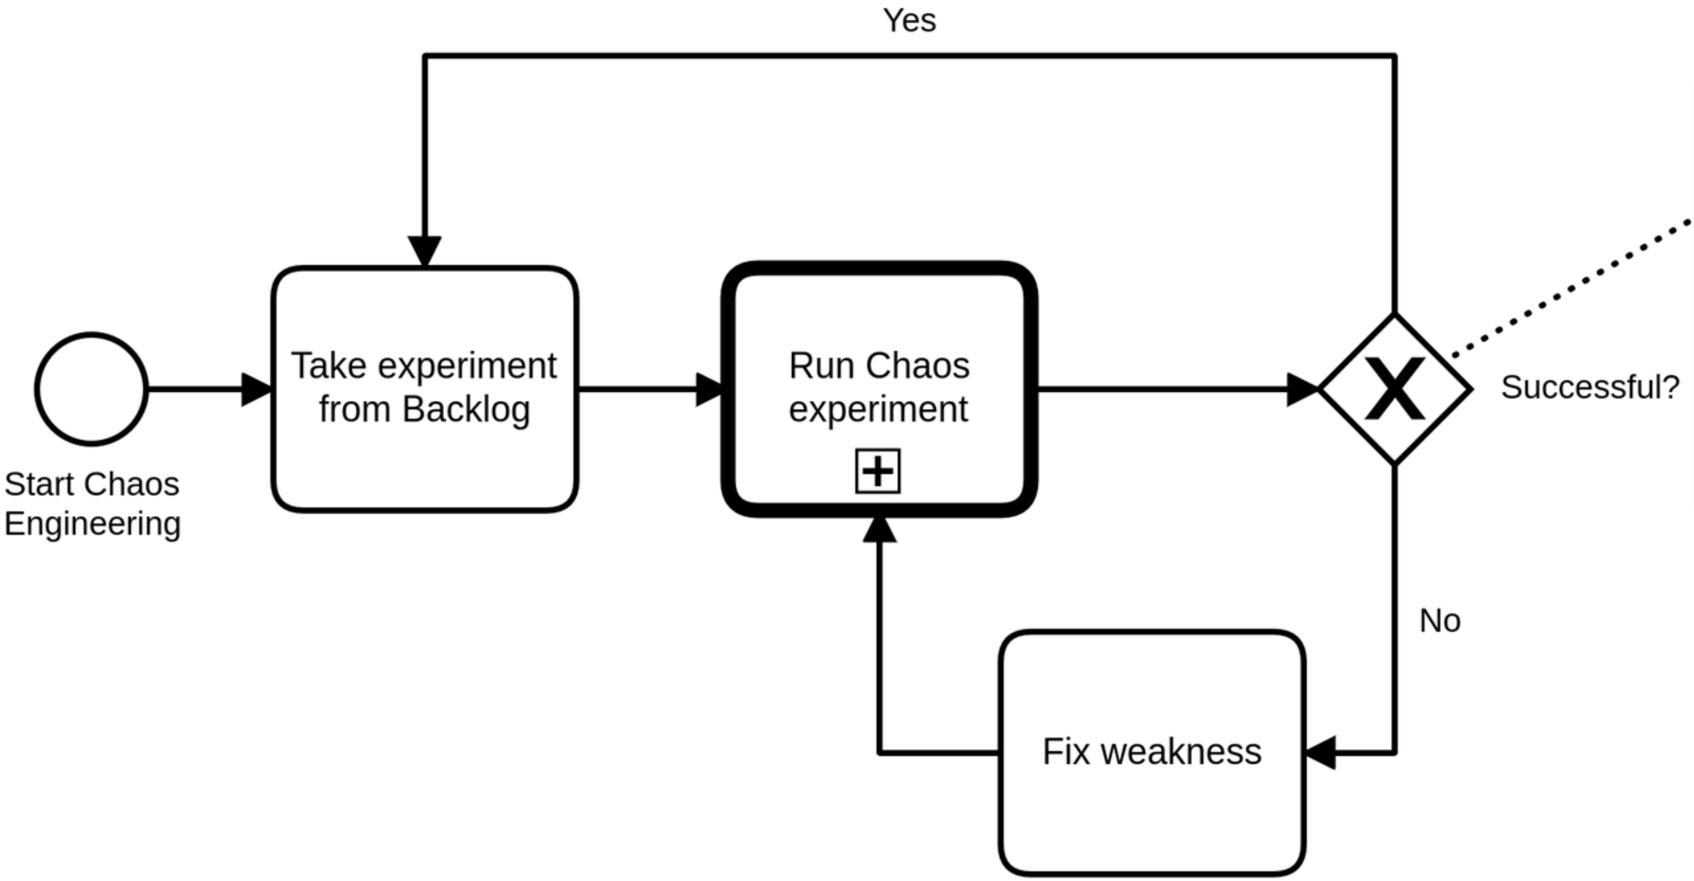

If not successful, we found a weakness we need to fix. Otherwise we build up confidence in our system

Chaos Engineering Experiment

Define Steady state:

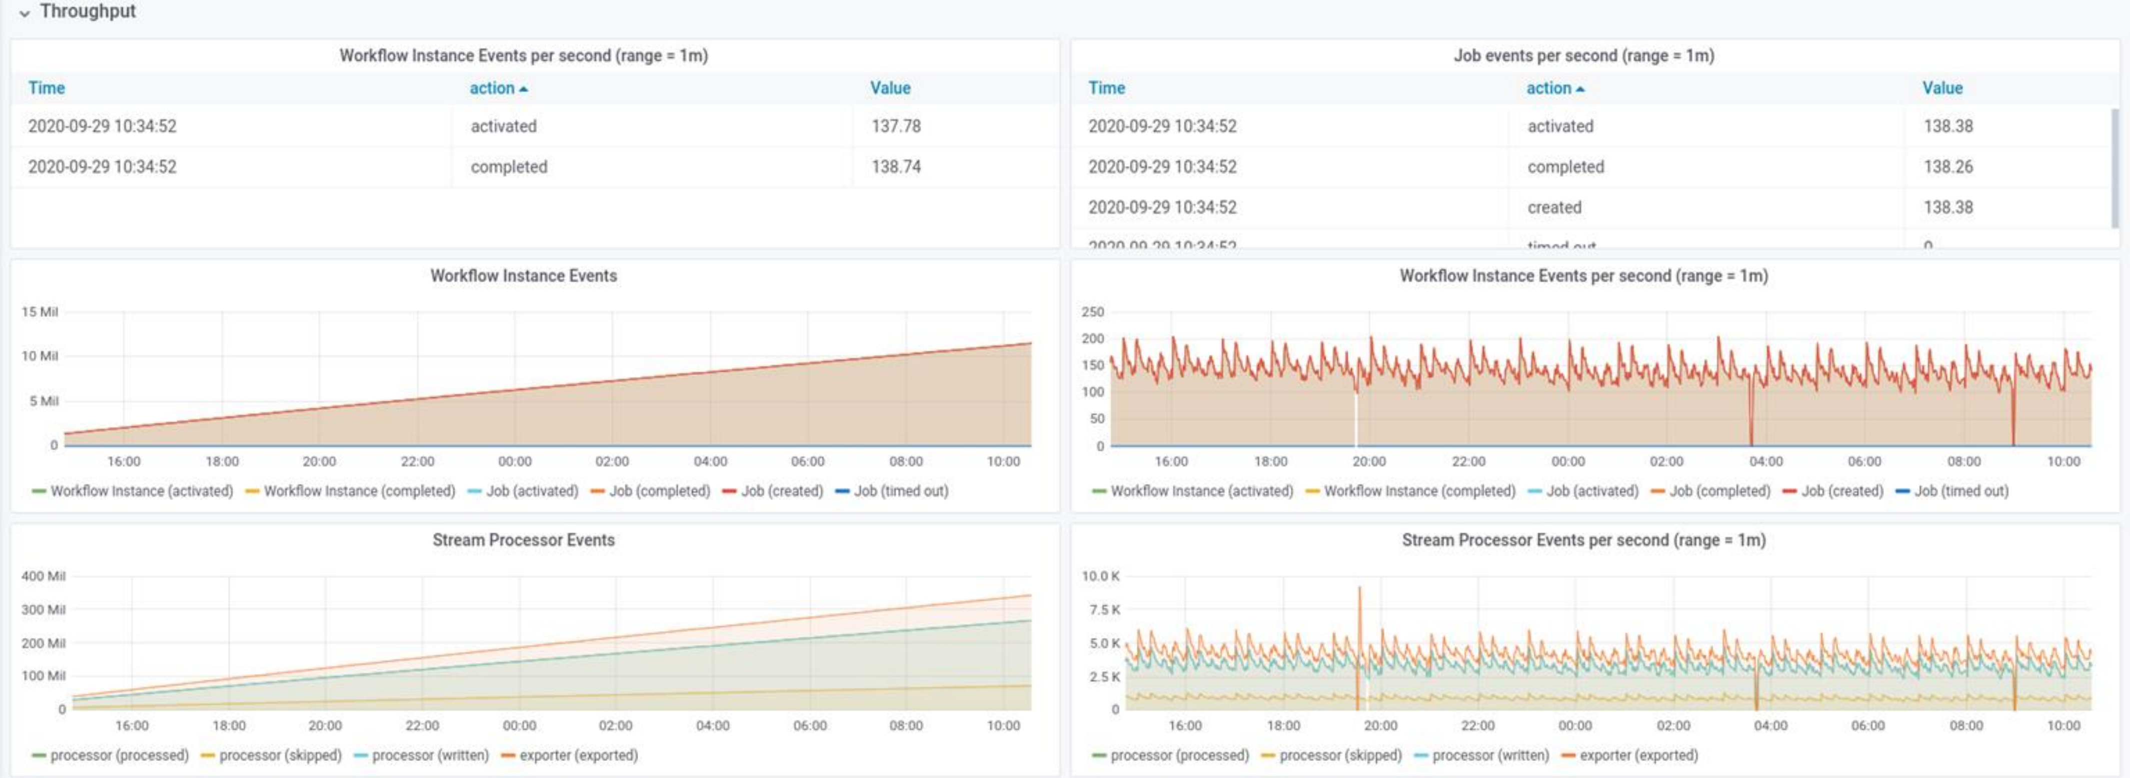

Workflow Instance Events per second (range = 1m)

Time	action ▲	Value
2020-09-29 10:34:52	activated	137.78
2020-09-29 10:34:52	completed	138.74

Job events per second (range = 1m)

Time	action ▲	Value
2020-09-29 10:34:52	activated	138.38
2020-09-29 10:34:52	completed	138.26
2020-09-29 10:34:52	created	138.38
2020-09-29 10:34:52	timed out	0

Chaos Engineering Experiment

Define Hypothesis:

- Throughput keeps stable, even if we restart a follower
- Throughput recovers after leader change

Advanced Principles:

- Hypothesis based on measurable output
- Vary Real World events

Advanced Principles

Run experiments in production:

- As a OSS project we have not **one** production setup / use case
- We needed decided for one and try that out
- In March 2019 we started with our first k8 setup

Advanced Principles

Run experiments in production:

- As a OSS project we have not **one** production setup / use case
- We needed decided for one and try that out
- In March 2019 we started with our first k8 setup

Default Benchmark Setup:

- 5 Nodes, 8 Partitions, Replication factor 3
- 150 workflow instances per second
- Workers to complete the jobs

Stability

Prerequisite of chaos engineering is to have a stable and resilient system.

We learned that during our first experiments.

https://www.thetoyshop.com/games-jigsaws/family-games/Jenga/p/091876

First chaos experiments

1. Start cluster
2. Wait for stable throughput
3. Hypothesis
4. Restart follower
5. Verify throughput

Automating Chaos Experiment’s

- We wanted to automate our manual experiments
- We evaluated several tools #2997
- Found chaos toolkit:

Chaos Toolkit aims to be the simplest and easiest way to explore building your own Chaos Engineering Experiments.

ChaosToolkit

Automating Chaos Experiment’s

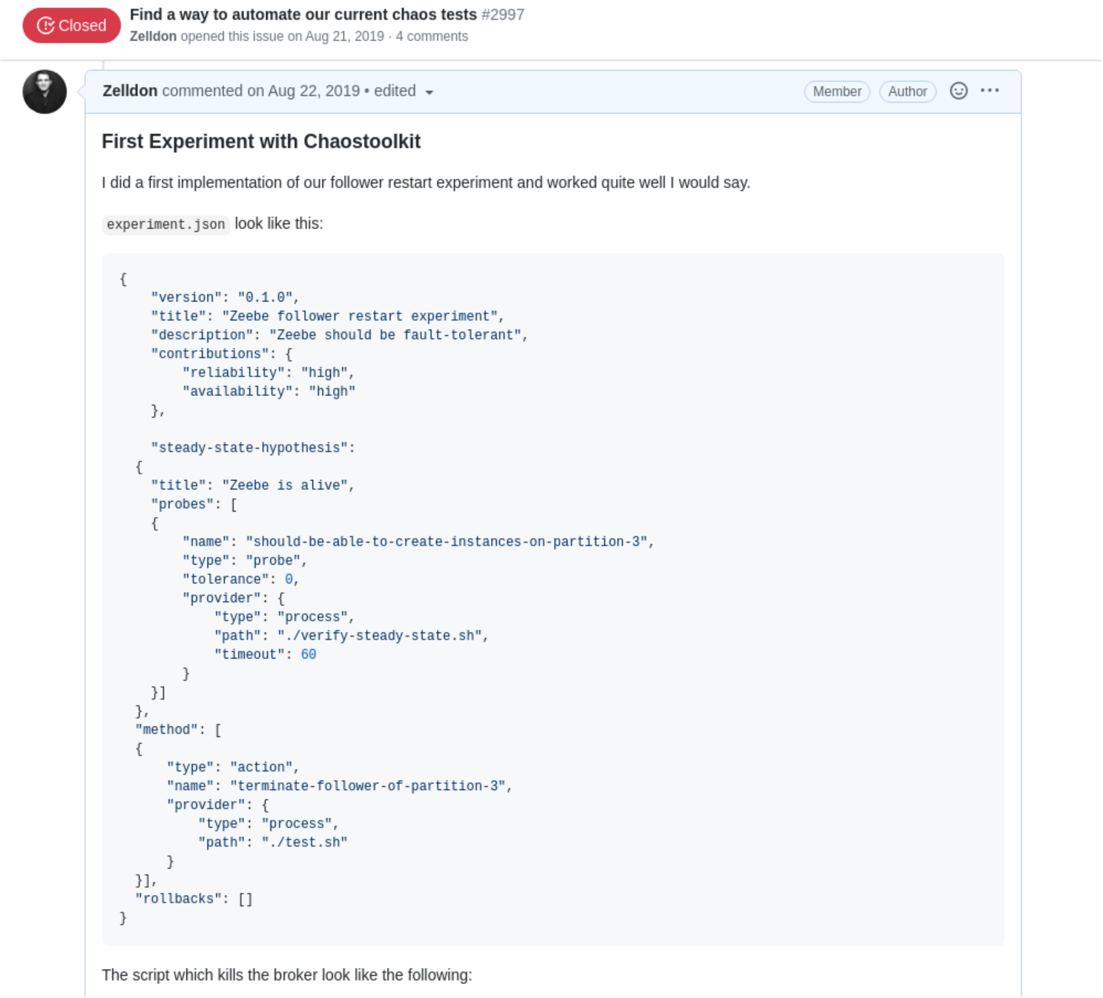

Find a way to automate our current chaos tests #2997

Zelldon opened this issue on Aug 21, 2019 · 4 comments

Zelldon commented on Aug 22, 2019 • edited

First Experiment with Chaostoolkit

I did a first implementation of our follower restart experiment and worked quite well I would say.

`experiment.json` look like this:

```
{
    "version": "0.1.0",
    "title": "Zeebe follower restart experiment",
    "description": "Zeebe should be fault-tolerant",
    "contributions": {
        "reliability": "high",
        "availability": "high"
    },

    "steady-state-hypothesis":
  {
    "title": "Zeebe is alive",
    "probes": [
    {
        "name": "should-be-able-to-create-instances-on-partition-3",
        "type": "probe",
        "tolerance": 0,
        "provider": {
            "type": "process",
            "path": "./verify-steady-state.sh",
            "timeout": 60
        }
    }]
  },
  "method": [
  {
      "type": "action",
      "name": "terminate-follower-of-partition-3",
      "provider": {
          "type": "process",
          "path": "./test.sh"
      }
  }],
  "rollbacks": []
}
```

The script which kills the broker look like the following:

- Uses same terminology as principles of chaos
- Highly customizable
- Allows to define steady state and methods to disrupt the steady state
- Rollbacks allow clean ups

Automating Chaos Experiment’s

```
$ chaos run experiment.json
```

It will do the following:

- Validate experiment
- Validate steady state
- Run methods/actions
- Validate steady state
- Succeed or fail

Automating Chaos Experiment’s

Chaos toolkit makes it simple to run the experiments in our CI, we run them every night.

The Jenkins job creates a Zeebe cluster and run all existing chaos experiments from a repository.

ChaosToolkit ✚ Jenkins

https://chaostoolkit.org/

jenkins.io

Building a Hypothesis Backlog

In order to run more experiments you need more hypotheses.

You need to build a backlog of it.

Building a Hypothesis Backlog

CAMUNDA CON 2020.2

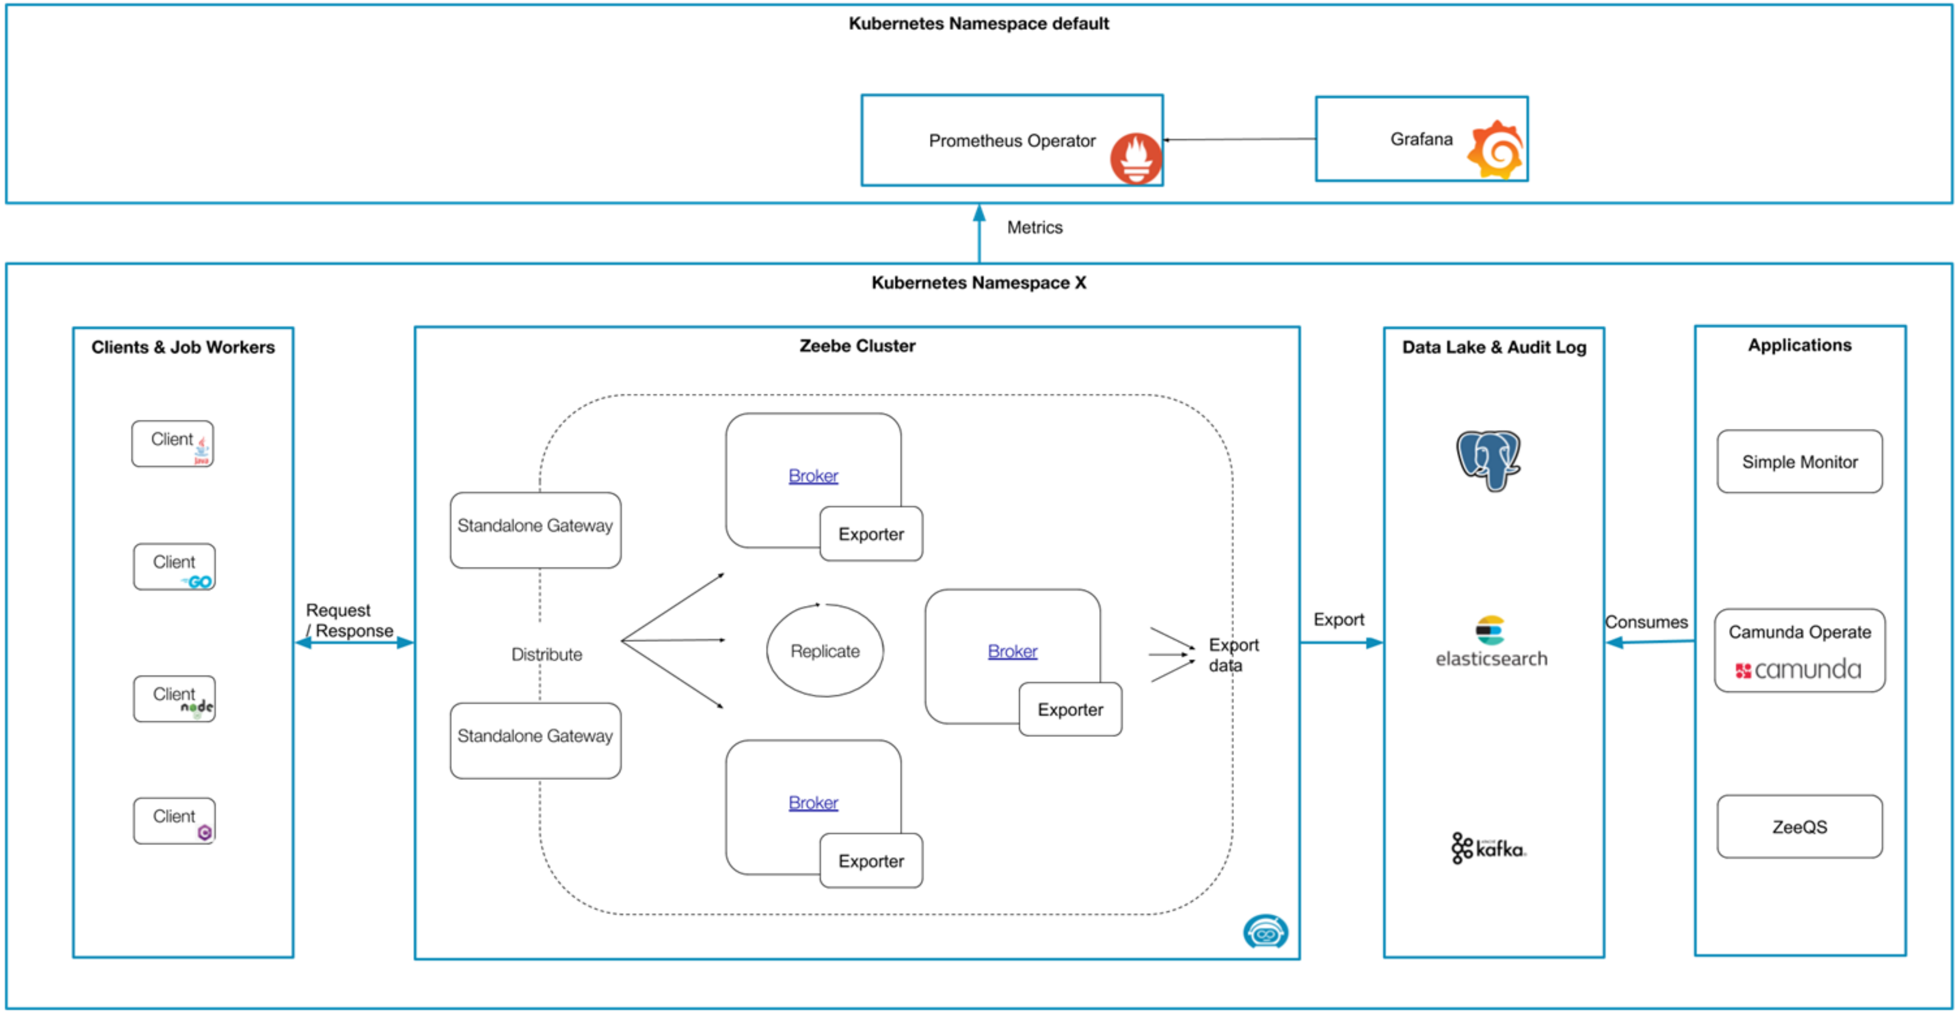

Building a Hypothesis Backlog

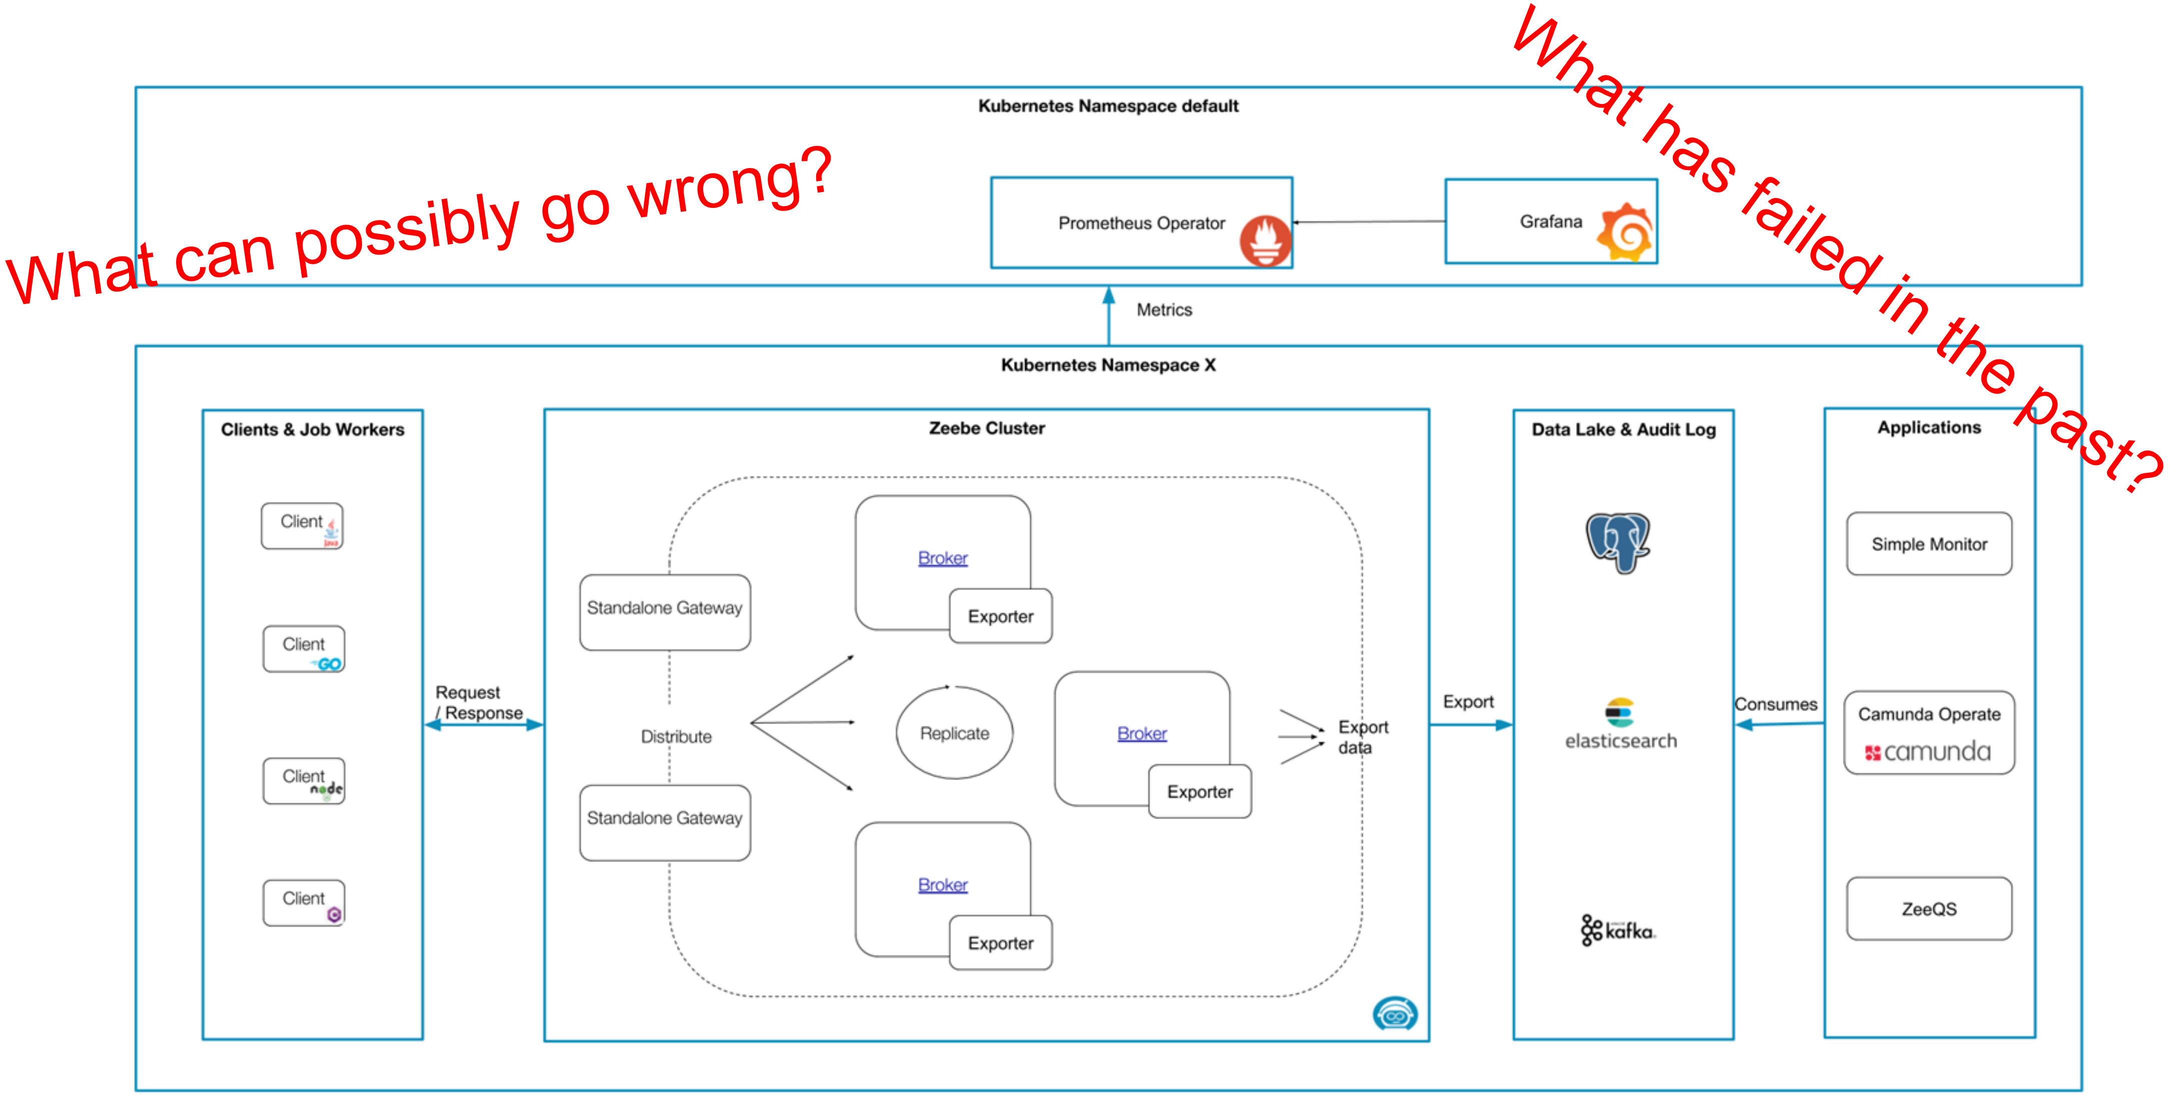

What can possibly go wrong?

What has failed in the past?

Building a Hypothesis Backlog

Translate failure cases

Failure case:

- After a network partition brokers were not able to find each other.

Hypothesis:

- Brokers are able to find each other after a network partition has happened.

Hypothesis Backlog

- New repository zeebe-io/zeebe-chaos for
  - backlog
  - chaos experiments
- Introduced “*Chaos Day*” to convert backlog

Chaos Day

- One day in a week
- Someone of the team and me
- Run new experiments
- Allows to convert backlog
- Helps to:
    - Find weaknesses
    - Spread knowledge
    - Build confidence in the team

Production Environment

Still we needed a real production environment.

Production Environment

- Camunda Cloud runs Zeebe internally
- Multiple different cluster plans
- Since last quarter we running our automated experiments against there setup

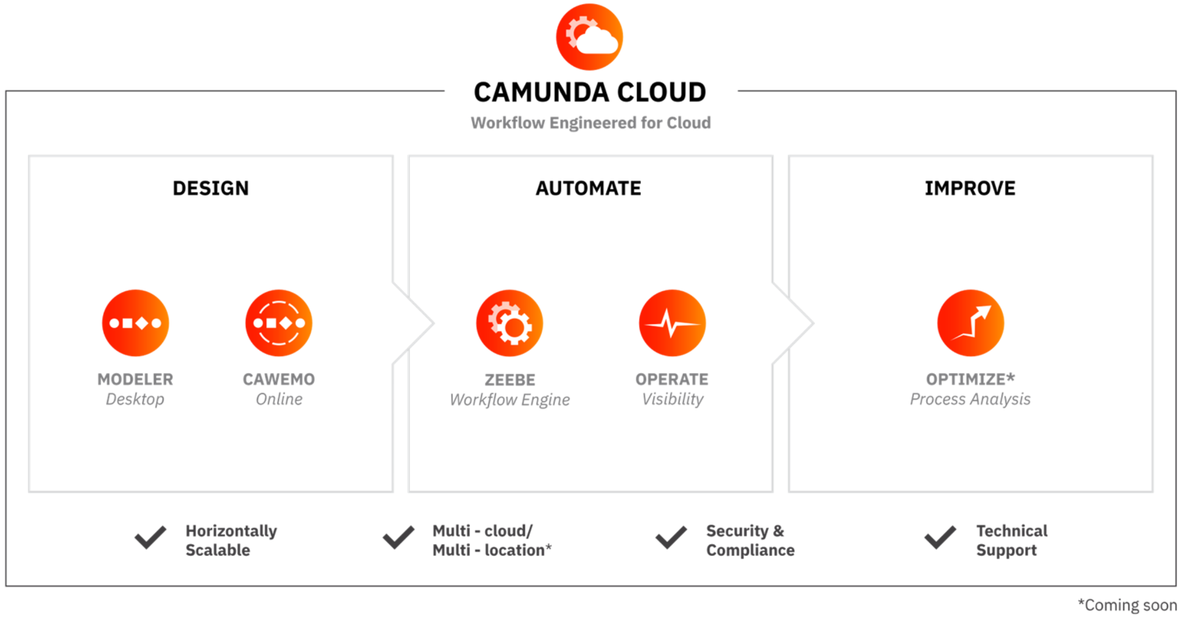

CAMUNDA CLOUD
Workflow Engineered for Cloud

DESIGN

MODELER *Desktop*

CAWEMO *Online*

AUTOMATE

ZEEBE *Workflow Engine*

OPERATE *Visibility*

IMPROVE

OPTIMIZE* *Process Analysis*

✓ Horizontally Scalable

✓ Multi - cloud/ Multi - location*

✓ Security & Compliance

✓ Technical Support

*Coming soon

Chaos Engineering at the Cloud

- Chaos Engineering is about experiment on a system
- There are more parts of our system
  - Alerting, Monitoring
  - Other Systems: Operate, Console etc.
  - Processes, like Incident Process
  - People

Chaos Engineering at the Cloud

- In Camunda cloud:
  - We want to react on incidents
  - We want to have SLA’s
  - We want to be prepared for the chaos

Chaos Engineering at the Cloud

- In Camunda cloud:
  - We want to react on incidents
  - We want to have SLA’s
  - We want to be prepared for the chaos

We found out that GameDay’s might help us here.

GameDay’s

The Book “Learning Chaos Engineering” by Russ Miles gave us good insights and explanation how to plan and run a GameDay, but also to apply other Chaos Engineering practices.

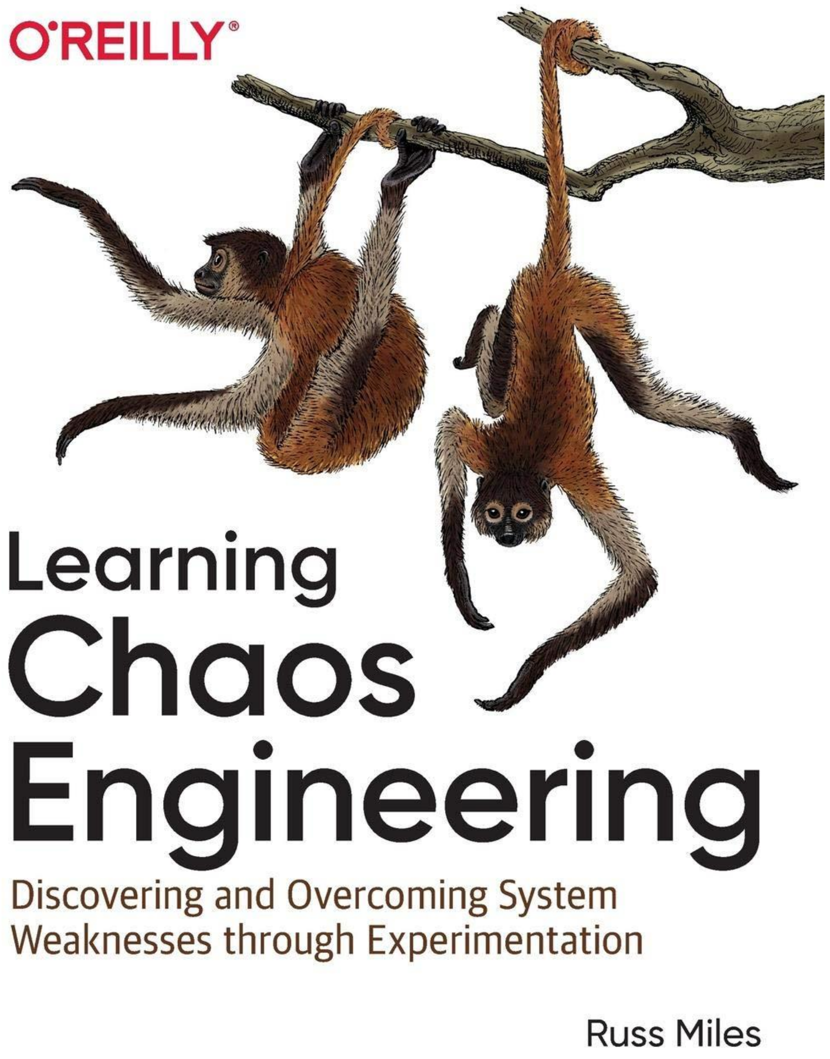
O'REILLY
Learning Chaos Engineering
Discovering and Overcoming System Weaknesses through Experimentation
Russ Miles

GameDay’s

A GameDay is a practice event, [...]
The goal of a Game Day is to practice
how you, your team, and your supporting systems
deal with real-world turbulent conditions.

Learning Chaos Engineering, Russ Miles p. 26

GameDay’s at Camunda

- Like a role-playing game
- Everyone has a role assigned
- An incident happens
- Everyone needs to help to resolve it

GameDay’s at Camunda

- Like a role-playing game
- Everyone has a role assigned
- An incident happens
- Everyone needs to help to resolve it

Consist of three Phases

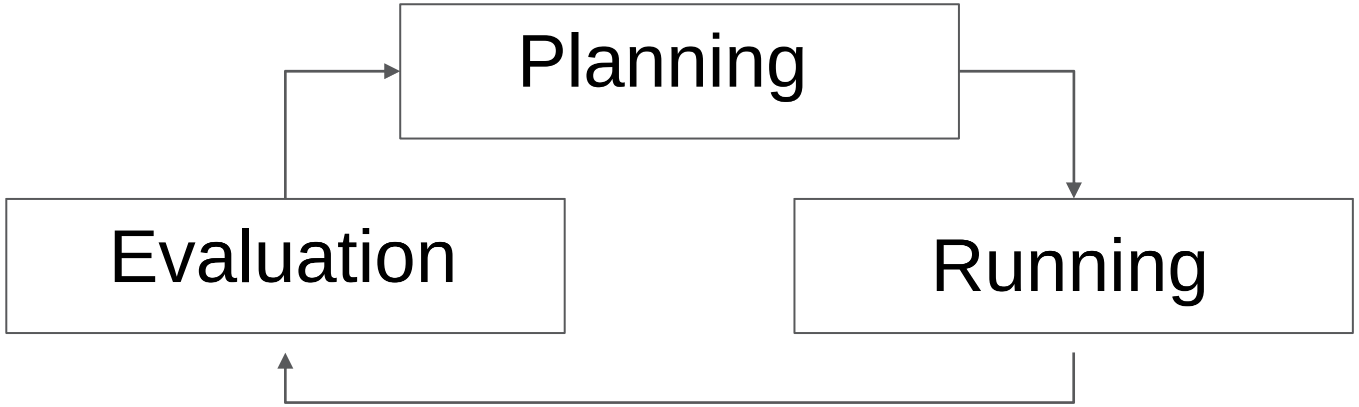

Planning Phase

Role assignment

- Gameday Facilitator
- Safety Monitor
- Customer
- Incident Management roles:
    - Incident Commander (IC)
    - Communications Lead (Comms)
    - Operations Lead (Ops)
- Operate Dev.
- Zeebe Dev.
- Observer

Planning Phase

- *Style*: Inform in advance
- *When:* Create an appointment and block three hours
- *Where:* Currently on dev
- *What*: Define the experiment based on the backlog

Running Phase

- Everyone joins via Zoom
- *Facilitator* gives intro:
    - Make clear what we are doing and why
    - It is not a test, which can have bad consequences
    - All outcomes are good, since we can learn from them
- *Customer* creates incident

Running Phase

- *IC* reacts on support issue
  - Pulls new people if needed
- Gameday ends
  - When incident is resolved
  - Time is up
- Explain the incident
- Appointment for the evaluation

Evaluation Phase

- Incident review
- Evaluate the facilitator’s log
- Discussing about GameDay as well
  - Got a lot of positive feedback, the people like it
- Start with planning the next one

Benefits of GameDay’s

CAMUNDA CON 2020.2

- Finding weaknesses before users do
- Sharing knowledge
  - People learn how the system, processes etc. work
  - Good way to onboard people
- People getting to know each other
  - Makes it easier for the next real incident
  - Good way to come together in this times
- Helps to build confidence in our systems
- It’s fun!

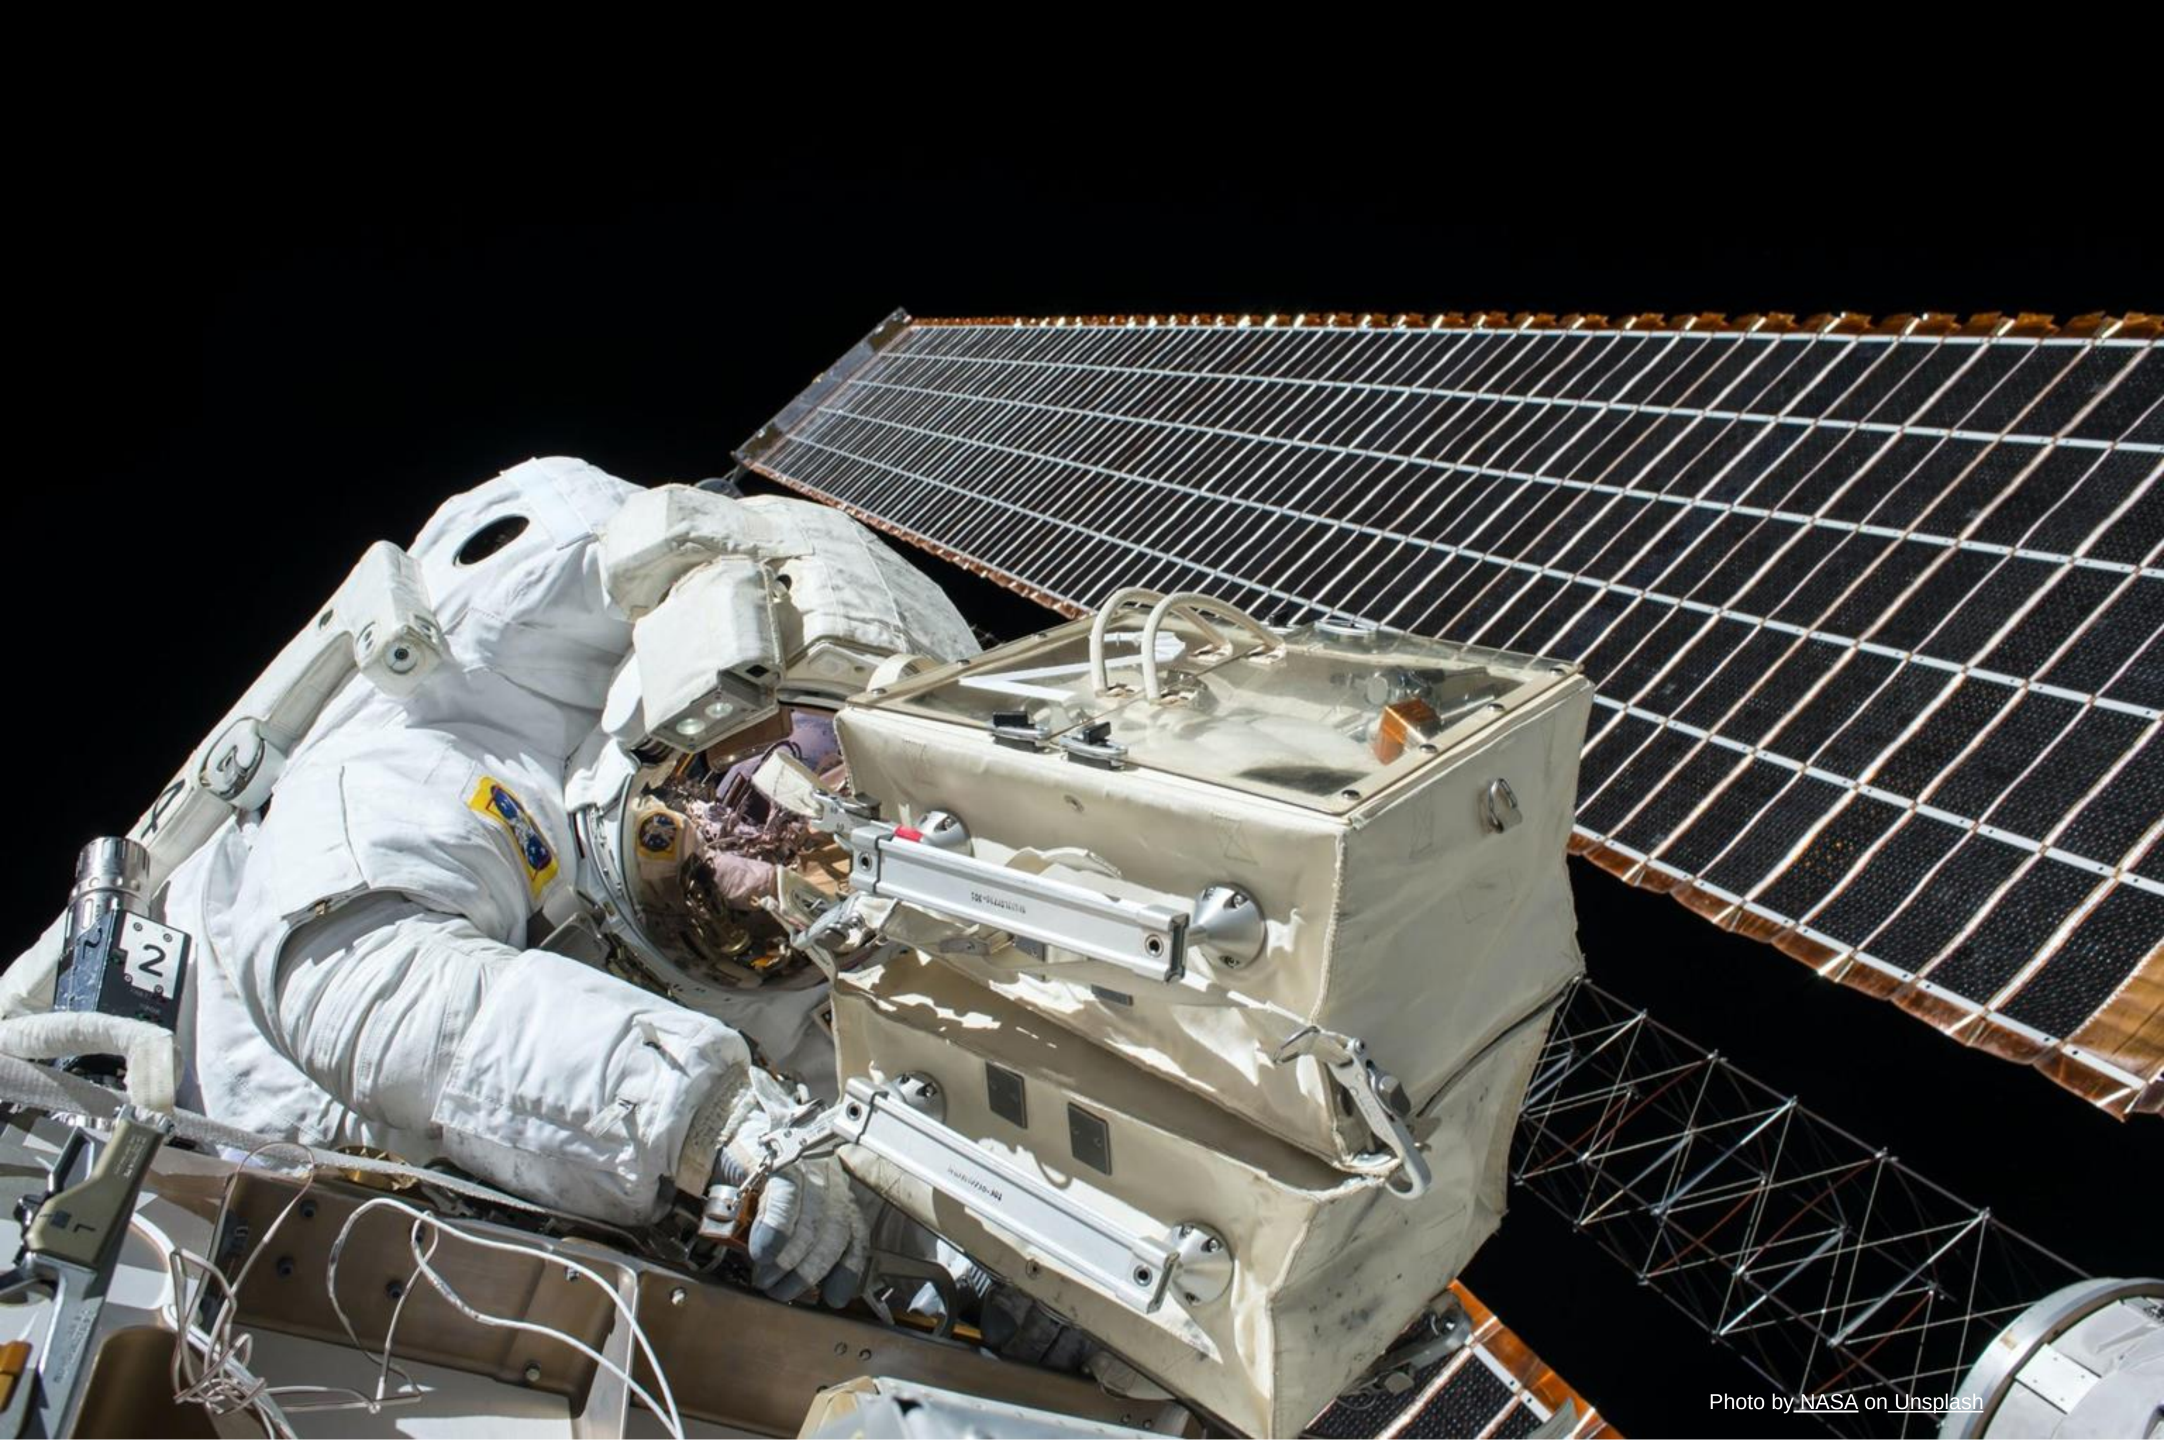

Photo by NASA on Unsplash

CAMUNDA
CON
2020.2

Questions?