Addressing Performance Challenges in Serverless Computing

Serverless Computing

- **Event-Driven**; no ephemeral processes
- **Granular Billing**; Pay per execution
- **(Almost) no operational logic**

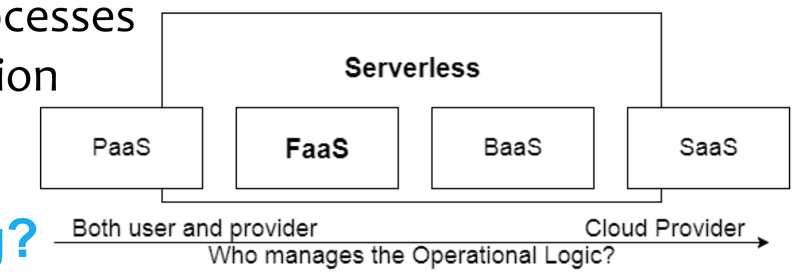

Function-as-a-Service (FaaS)

- A form of serverless computing
- User provides **functions** (source code)
- Provider manages **resources**, **lifecycle**, and **event-driven execution** of functions

Why serverless computing?

- User can focus solely on **business logic**; deferring **operational logic** to cloud provider.
- User only pays for actual usage (not for reserved or unused resources).
- Provider has **more insight and control** over workloads to improve the low cloud resource utilization. [3]

Challenges

- **Community challenges**: Lack of terminology, standardization, portability, and benchmarks. [1]
- Many challenges in **(system) operations**, **software engineering**, and **performance engineering**.
  - Performance challenges: overhead, workload prediction, better scheduling policies. [2]

Approach

- **SPEC RG Cloud:** independent, multi-country, multi-institution, multi-disciplinary (DistribSys, SwEng, PerfEng)
  - Target community and interdisciplinary challenges.
- **@Large Research:** approach specific challenges using our expertise and background (DistribSys, grids, workflows, graph processing).

Preliminary Results

fission workflows

- **Workflow-based** cloud function composition.
- Built on top of **production-grade infrastructure**.
- **Pluggable architecture** to support multiple FaaS platforms and clouds.
- Introduces the **Serverless Workflow Language.**
- Reduces function deployment latency using **prewarming.**

FaaS Reference Architecture

- Extensive **survey** (~100 systems) of current cloud landscape.
- Identified **use cases** and common components.
- FaaS reference architecture, based on initial mapping:

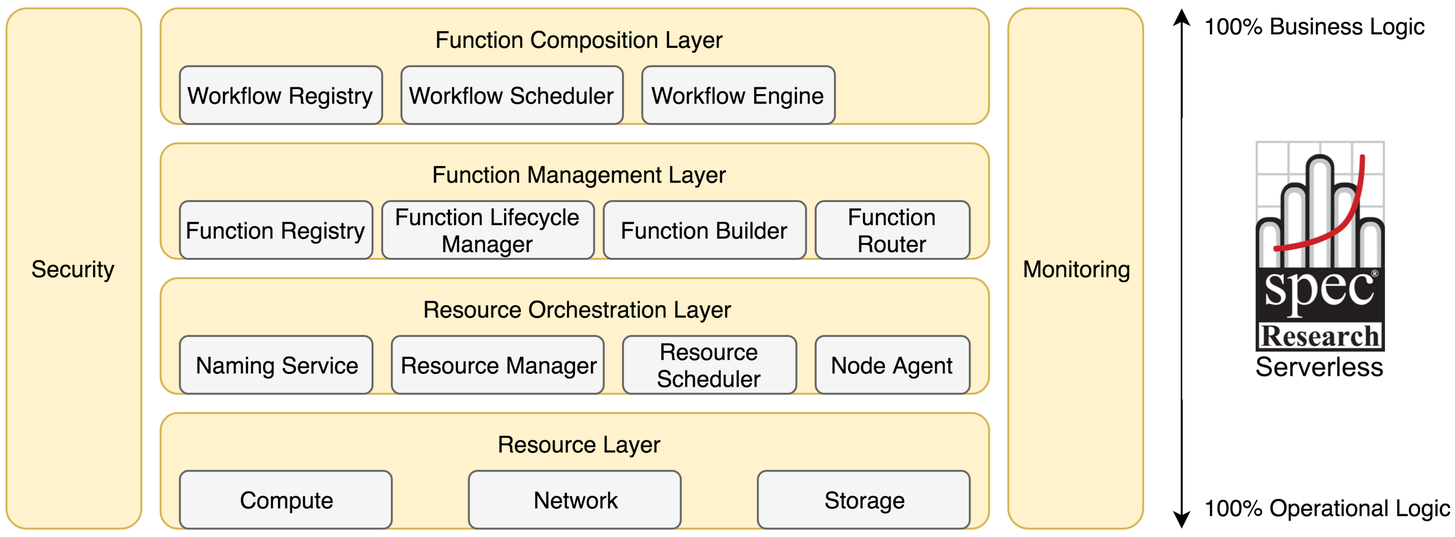

Roadmap

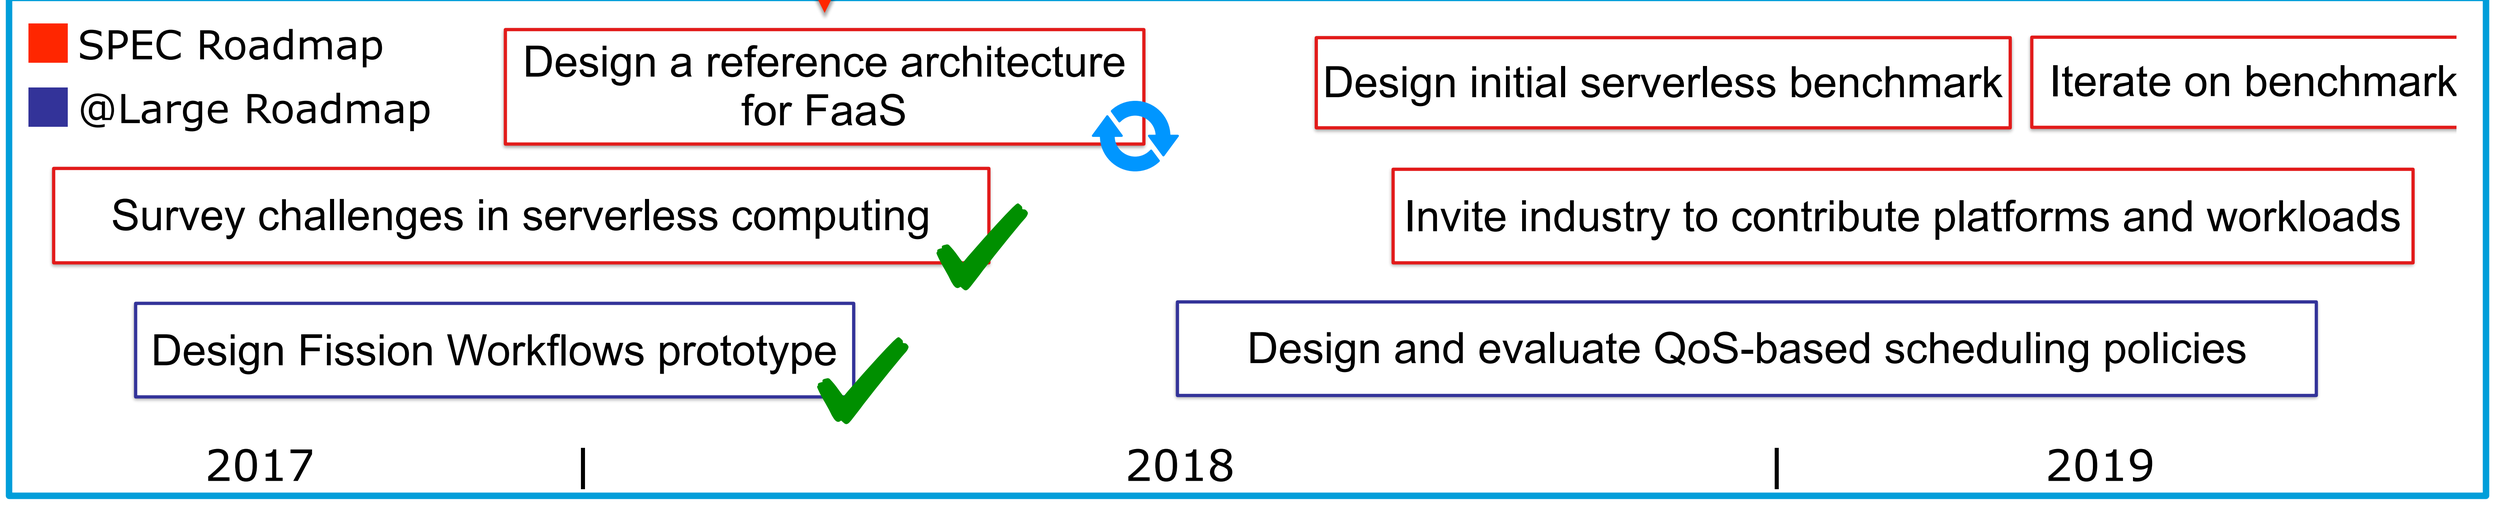

[1] van Eyk, Erwin, et al. "The SPEC cloud group's research vision on FaaS and serverless architectures." *Proceedings of the 2nd International Workshop on Serverless Computing*. ACM, 2017.
[2] van Eyk, Erwin, et al. "A SPEC RG Cloud Group's Vision on the Performance Challenges of FaaS Cloud Architectures" *Proceedings of the 9th International Conference on Performance Engineering, ACM, 2018.*
[3] Cortez, Eli, et al. "Resource Central: Understanding and Predicting Workloads for Improved Resource Management in Large Cloud Platforms." *Proceedings of the 26th Symposium on Operating Systems Principles*. ACM, 2017.

Interested in the next step in cloud computing? Join us!

Erwin van Eyk, Alexandru Iosup
E.vanEyk@atlarge-research.com

@Large Research
Massivizing Computer Systems

spec Research Serverless

TU Delft Delft University of Technology
Julius-Maximilians-Universität Würzburg
SAP
PLATFORM9
VU University Amsterdam
IBM
ESPOL
You?