Optimizing Latency in Function-as-a-Service with Distributed Promises

Soam Vasani, <u>Erwin van Eyk</u>

Platform9 Systems Inc.

Function-as-a-Service Market estimated to be worth **\$7.72 Billion** by 2021



Gartner 2017 Function-as-a-Service is one

of the top trends in cloud computing.

Serverless Matches Container Adoption



Source: The New Stack Analysis of a February 2017 survey of 500+ IT professionals (https://newrelic.com/content/dam/newrelic/resources/ebooks/cloud-survey-report-ebook.pdf). Static Cloud: Public cloud used to some extent but applications are managed like before. Dynamic Cloud: A significant portion of strategic workloads are run in the public cloud and the enterprise is able to agilely re-allocate resources.

https://www.marketsandmarkets.com/Market-Reports/function-as-a-service-market-127202409.html http://get.cloudability.com/ebook-state-of-cloud-2018.html

THENEWSTACK

Research into Function-as-a-Service





Combined industry and academia effort to address high-level (performance) challenges

Terminology Challenges Reference Architectures Benchmarking **@Large Research** Massivizing Computer Systems

Actively exploring challenges in FaaS

Complex FaaS function composition using **workflows**. On-demand **Graph Processing** with FaaS



Open Source Software as a Service

IaaS OpenStackPaaS KubernetesFaaS Fission (own product)

Serverless Cloud Native Landscape V2.1

See the serverless interactive landscape at s.cncf.io

Greyed logos are not open source



FaaS Function Model





Controller manages function definitions and exposes API. **Poolmgr** maintains a pool of running 'generic' containers to reduce deployment time of functions.

Router directs requests to the 'right' deployed function.



Composing Complex Functions



Challenges

- Dealing with communication logic
- Structure hard to understand
- Implicit dependencies
- Performance overhead of function calls (eg. cold starts)



Event Store persist all events related to the workflows. Controller watches active workflows and execute scheduled actions. Scheduler decides which actions need to be taken based on current state.



Example: Photo Transform Workflow



- Data-intensive, on-demand functions
- Several distinct computation steps for transformation





Photo Transform Workflow





Data (Transfer) and lots of it!





| US West DC | | | | |
|------------|--------------------|-------------|------------------------------------|---|
| | | Within a DC | US West DC \rightarrow Europe DC | |
| | Bandwidth Cost [1] | € 0.00 | € 0.09 / 1 GB | |
| | Latency [2] | ~1 ms | ~158 ms | e |

[1] https://aws.amazon.com/ec2/pricing

[2] Gandhi, A., & Chan, J. (2015). Analyzing the Network for AWS Distributed Cloud Computing. ACM SIGMETRICS Performance Evaluation Review World map: copyright

World map: copyright FreeVectorMaps.com



Naive Approach

Naive Approach: expensive and slow!

Color Task can only start after controller receives the data!

Costs € 0.09 * 4 = € 0.36 per GB

Latency ~158 ms * 4 = ~632 ms





Keep the data - the intermediate results - in the same datacenter.





"Intermediate Result" = Promise

A promise holds a (future) value; API:

- Get() Returns the value or error
- Resolve(value) Makes Get() return VALUE
- Reject(error) Makes Get() throw ERROR

Optionally: watch promise

"Distributed" Promises

- Promises across the context of a single process
- "Promises allow a caller to run in parallel with a call and to pick up the results of the call, including any exceptions it raises."
 - Liskov '88, Distributed Programming in Argus
 - Liskov and Shrira '88, Promises: Linguistic Support for Efficient Async Procedure Calls in Distributed Systems

Promises in the FaaS Function model



Await:

- Promise is resolved \rightarrow function starts executing.
- Promise is rejected \rightarrow function fails without executing





Dual Optimization Problem

Computation: where do we execute the functions?

Data: where do we store the promises?

How do **computation** and **data** placements influence each other?

How to optimize the overall performance in presence of these two factors?

The use of promises is a promising candidate for data-intensive functions in FaaS.







- https://fission.io/ https://fission.io/workflows
- 🥑 @erwinvaneyk
- erwin@platform9.com





Kubernetes is eating the resource orchestration ecosystem.



http://get.cloudability.com/ebook-state-of-cloud-2018.html

"Distributed" Promises

- Promises across the context of a single process
- See: the Argus system
 - Liskov '88, Distributed Programming in Argus
 - Liskov and Shrira '88, Promises: Linguistic Support for Efficient Async Procedure Calls in Distributed Systems