Roboť

Challenges for Serverless Native Cloud Applications

Third International Workshop on Serverless Computing

Ben Kehoe

Cloud Robotics Research Scientist AWS Community Hero @ben11kehoe

2018-07-02



About me

- Cloud robotics: Connecting robots to the internet to help them do more and better things
- Cloud Robotics Research Scientist at iRobot (since Feb 2015)
- PhD in Mechanical Engineering from UC Berkeley, Dec 2014
 - "Brass"



















Roomba® 690 Roomba® 890 Roomba® 960 Roomba® 980
Then Now





Why serverless?

The reasons for serverless

What:

- Service-full + ephemeral compute
 - Not always F, not always aaS
- Resources billed \rightarrow resources used
- Smaller, more abstract control plane

Why:

- Lower cost
- Lower operations burden
- Faster time to market

Focus on business value

Serverless architecture

- Call graph \rightarrow component graph
- Distributed system thinking
 - Traditionally occurs at system boundaries
 - Serverless: must be treated systematically
- Build robust-by-design systems

Before serverless

aka the dark ages

```
def foo(input):
   quux = bar(input.baz)
   internalState.quux = quux
```

```
def bar(input):
```

do work

```
return result
```


Logging

R

Lifecycle event

Deployment

Deployment

- Red/black imposes requirements on clients
- Blue/green is the direction providers are headed
- Existing paradigm:
 - Blue/green controller is part of your component graph
 - Update component graph in-place
 - Controller manages roll-out

Function/code versions must be first-class citizens in infrastructure

Summing up

- Red/black is actually pretty easy for serverless, but harder on the service consumers
- Blue/green component graph gets complicated
- Existing deployment tools don't represent intermediate states
- Native support for blue/green component graphs is critical

Questions?

