Function as a Service

General Principles, Container Virtualization, OpenFaaS, OpenWhisk – Summer Term 2019

Henry-Norbert Cocos
cocos@stud.fra-uas.de

Computer Science
Faculty of Computer Science and Engineering
Frankfurt University of Applied Sciences

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Contents

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Container Virtualization

Container Virtualization capsulates Applications in virtual environments!

This technology has a long going history!
(`chroot` was first implemented in 1979 [1])

Containers are more efficient than Hypervisor-based Virtualization or Paravirtualization!

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Docker

Figure: Docker

Source:
https://www.docker.com/
brand-guidelines

Docker

- Released by dotCloud 2013
- Enables Container Virtualization
- A more advanced form of Application Virtualization
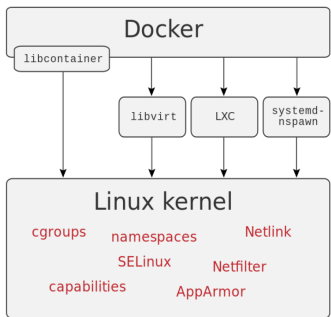- Available for:
  Linux, MacOS, Windows

Docker Architecture

Figure: Docker Architecture

Docker Architecture

- Docker uses the Linux Kernel
- `libcontainer` creates containers
- `libvirt` manages Virtual Environments
- LXC will be replaced by `libcontainer`

Docker Application Architecture I

Figure: Docker Application Architecture

Source: https://docs.docker.com/engine/
docker-overview/#docker-architecture

Applications in Docker [2]

- Client-Server Architecture
- Docker Client `docker`
- Docker Daemon `dockerd`

Docker Objects

- Images
- Containers

Docker Application Architecture II

Figure: Docker Application Architecture

Source: `https://docs.docker.com/engine/docker-overview/#docker-architecture`

Docker Client `docker`

- Manages Docker Daemon/s

Docker Daemon `dockerd`

- Listens to Requests
- Manages Docker Objects (images, containers, etc.)

Docker Application Architecture III

Figure: Docker Application Architecture

Source: https://docs.docker.com/engine/
docker-overview/#docker-architecture

Docker Objects
- **Containers**
  - Runnable Instance
  - Isolated from other
    containers
- **Images**
  - Read-Only File
  - Defines an Application

Docker Benefits

Docker has the following benefits:

- Less resource consumption than OS Virtualization
- Isolation of Applications
- Fast deployment
- Perfect for testing purposes
- Containers can be restarted

Docker Swarm and Kubernetes

The Docker Engine has a build in solution for Cluster deployment and management. The swarm mode enables the control over multiple Docker hosts and is crucial for the scalability of applications [3]. Kubernetes is a different system that enables deployment over multiple hosts.

KFURT
ERSITY
OF APPLIED SCIENCES

Function as a Service

Function as a Service (FaaS) has emerged as a new technology in Cloud Computing!

FaaS reduces administration tasks and brings the focus back to the Source Code! [4]

FaaS enables more effective event-driven applications!
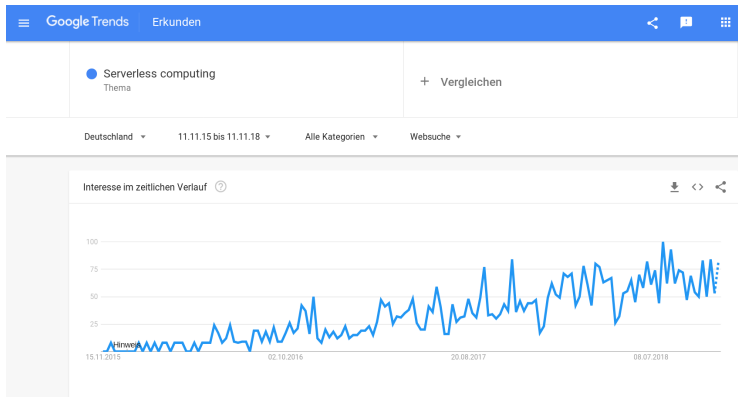
Function as a Service

Figure: Google Trends for Serverless Computing

Function as a Service

(a)   (b)

(c)   (d)

Figure: Popular FaaS Offerings:
(a) AWS Lambda [5]
(b) Google Cloud Functions [6]
(c) IBM Cloud Functions [7]
(d) Apache OpenWhisk [8]

Function as a Service (FaaS)

- Event-driven
- Scalable
- Fast deployment of code
- Payment per invocation

Amazon Alexa

Alexa Skills are executed in AWS Lambda!

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Public FaaS offerings – AWS Lambda

Figure: AWS Lambda [5]

AWS Lambda

- Released in 2014
- Fully automated administration
- Automated Scaling
- Built in fault tolerance
- Support for multiple languages: Java, Node.js, C# and Python

Public FaaS offerings – IBM Cloud Functions

Figure: IBM Cloud Functions [7]

IBM Cloud Functions

- Released in 2016
- Event-driven Architecture
- Automated Scaling
- Apache OpenWhisk is basis of IBM Cloud Functions (No Vendor Lock-in!)
- Support for multiple languages: JavaScript, Python, Ruby, ... [1]

---

[1]More on that in Section 4

FaaS Generic Architecture I

Figure: Generic FaaS Architecture [9]

FRANKFURT
UNIVERSITY
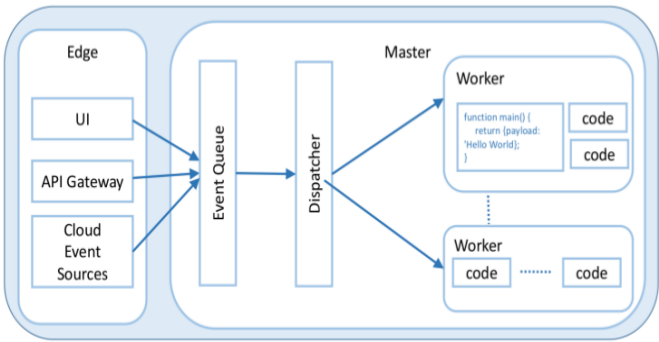OF APPLIED SCIENCES

FaaS Generic Architecture II

Edge

- **UI** – An UI for the management of functions
- **API Gateway** – The general API for the implemented functions

Event Queue/Dispatcher

- **Event Queue** – Manages the triggered Events
- **Dispatcher** – Manages the scaling of invocations

Worker

- **Worker Processes/Containers** – Execute the function invocations

Interesting Paper

Figure 10 and the explanation of the architecture are taken from the paper of Baldini et.al. [9]

KFURT
ERSITY
IENCES

Function as a Service

The Service consists of:

Scalability – Reaction to large number of Requests

Environment – Running the code on a Platform

Virtualization – Capsulation of running code

Figure: Evolution of Virtualization [10]

OpenFaaS

Figure: OpenFaaS

Source:
https://github.com/openfaas

OpenFaaS

- Open Source Platform
- Functions can be deployed and scaled
- Event-driven
- Lightweight
- Support for multiple languages: C#, Node.js, Python, Ruby

OpenFaaS Architecture I

Functions as a Service

Figure: OpenFaaS Architecture [11]

OpenFaaS Architecture [11]

- **Gateway API**
  - Provides a Route to the functions
  - UI for the management of functions
  - Scales functions through Docker
- **Function Watchdog**
  - Functions are added as Docker Images
  - Entrypoint for HTTP Requests
  - In → `STDIN`
    Out → `STDOUT`

OpenFaaS Architecture II

Functions as a Service

Figure: OpenFaaS Architecture [11]

OpenFaaS Architecture [11]

- **Prometheus**
  - Collects Metrics
  - Function Metrics can be inspected
  - Can be accessed through Web-UI
- **Docker**
  - Isolates Functions in Docker Images
  - Docker Swarm distributes functions
  - Kubernetes can be used to orchestrate Docker Instances

OpenFaaS Benefits

OpenFaaS has the following benefits:

- Open Source
- Low resource consumption
- Deployment of functions
- Autoscaling
- Build in Monitoring and Metrics (Prometheus)

OpenFaaS on Raspberry Pi

OpenFaaS together with Docker Swarm have a low resource consumption. Therefore OpenFaaS has been installed on a cluster of 6 Raspberry Pis. Further evaluation of the service on Raspberry Pis has to be made. More information about installation on Raspberry Pi [12].

KFURT
ERSITY
OF APPLIED SCIENCES

Installing OpenFaas

In order to work with OpenFaas 3 packages need to be installed:

- **Docker**
- **OpenFaas Framework**
- **OpenFaas CLI**

Installing Docker

Install Docker:

```
$ curl −sSL https://get.docker.com | sh
```

Add `docker` **User to** `<USER>` **User Group:**

```
$ usermod <USER> −aG docker <USER>
```

Initialize Docker Swarm on Master Node:

```
$ docker swarm init
```

Command on slaves to join workers to docker swarm cluster:

```
$ docker swarm join −−token <TOKEN>
```

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Installing OpenFaas

Download OpenFaaS from github:

```
$ git clone https://github.com/alexellis/faas/
```

Changing into directory and deploy OpenFaaS [2]:

```
$ cd faas && ./deploy_stack.armhf.sh
```

Install OpenFaaS CLI:

```
$ curl −sSL cli.openfaas.com | sudo sh
```

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

---

[2]The script `deploy_stack.armhf.sh` is necessary for the ARM platform

Creating Functions in OpenFaas

**Now that Docker and OpenFaas have been installed
deployment of functions can begin!**

Creating an application using Minio in OpenFaaS

Application Flow [3]:

- The Application downloads an image and stores it in a Bucket
- The image is loaded from the Bucket and then converted to Black/White
- In the last step the image is stored in another Bucket
- The Application consists of **OpenFaas** and **Minio (a private object-based storage with S3-API)**

For this Application two Functions are needed!

> **OpenFaaS als leichtgewichtige Basis für eigene Functions as a Service**.
> Henry-Norbert Cocos, Christian Baun. iX 9/2018, S.122-127, ISSN: 0935-9680

---

[3]Source Code and explanation available at:
https://blog.alexellis.io/openfaas-storage-for-your-functions/

Creating Directory for Function

Create a functions directory:

```
$ mkdir functions
```

Change into this directory and issue the following command:

```
$ cd functions && faas-cli new --lang python-armhf
    loadimages
$ faas-cli new --lang python-armhf processimages
```

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Templates for Python Functions

The command from the last slide will create the following files in the functions directory:

- `loadimages/handler.py`
- `loadimages/requirements.txt`
- `loadimages.yml`

- `processimages/handler.py`
- `processimages/requirements.txt`
- `processimages.yml`

Install Minio

Install Minio Client and Server as Docker Containers:

```
$ docker pull minio/mc
$ docker run minio/mc ls play
$ docker pull minio/minio
$ docker run —p 9000:9000 minio/minio server /data
```

Start Minio Server

Start Minio Server and get Credentials:

```
$ docker run −p 9000:9000 minio/minio server /data
...
Endpoint:  http://172.17.0.2:9000
http://127.0.0.1:9000
AccessKey: <ACCESSKEY>
SecretKey: <SECRETKEY>
...
```

Configure the Minio Client

In the next step the Minio Client has to be configured.

Configure the Access:

```
$ ./mc config host add TestService http
    ://192.168.178.21:9000 <ACCESSKEY> <SECRETKEY>
```

Creating the Buckets

The Minio Client is used to create two Buckets.

Creating the Buckets:

```
$ ./mc mb TestService/incoming
$ ./mc mb TestService/processed
```

One Bucket for incoming Images and one for processed Images

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

YAML File of Function `loadimages`

```
provider:
  name: faas
  gateway: http://192.168.178.21:8080

functions:
  loadimages:
    lang: python
    handler: ./loadimages
    image: loadimages
    environment:
      minio_hostname: "192.168.178.21:9000"
      minio_access_key: <ACCESSKEY>
      minio_secret_key: <SECRETKEY>
      write_debug: true
```

Listing 1: File `loadimages.yml`

YAML File of Function `processimages`

```
provider:
  name: faas
  gateway: http://192.168.178.21:8080

functions:
  processimages:
    lang: python
    handler: ./processimages
    image: processimages
    environment:
      minio_hostname: "192.168.178.21:9000"
      minio_access_key: <ACCESSKEY>
      minio_secret_key: <SECRETKEY>
      write_debug: true

  convertbw:
    skip_build: true
    image: functions/resizer:latest
    fprocess: "convert — —colorspace Gray fd:1"
```

Listing 2: File `processimages.yml`

requirements.txt of the Functions

```
minio
requests
```

Listing 3: File requirements.txt

loadimages Function in Python I

```python
1 from minio import Minio
2 import requests
3 import json
4 import uuid
5 import os
6
7 def handle(st):
8   req = json.loads(st)
9
10  mc = Minio(os.environ['minio_hostname'],
11   access_key=os.environ['minio_access_key'],
12   secret_key=os.environ['minio_secret_key'],
13   secure=False)
14
15  names = []
16  for url in req["urls"]:
17   names.append(download_push(url, mc))
18  print(json.dumps(names))
```

Listing 4: File loadimages Part I

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

loadimages Function in Python II

```python
 1 def download_push(url, mc):
 2   # download file
 3   r = requests.get(url)
 4
 5   # write to temporary file
 6   file_name = get_temp_file()
 7   f = open("/tmp/" + file_name, "wb")
 8   f.write(r.content)
 9   f.close()
10
11   # sync to Minio
12   mc.fput_object("incoming", file_name, "/tmp/"+file_name)
13   return file_name
14
15 def get_temp_file():
16   uuid_value = str(uuid.uuid4())
17   return uuid_value
```

Listing 5: File `loadimages` Part II

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

processimages Function in Python I

```python
from minio import Minio
import requests
import json
import uuid
import os

def handle(st):
  req = json.loads(st)

  mc = Minio(os.environ['minio_hostname'],
    access_key=os.environ['minio_access_key'],
    secret_key=os.environ['minio_secret_key'],
    secure=False)

  names = []
  source_bucket = "incoming"
  dest_bucket = "processed"

  for file_name in req:
    names.append(convert_push(source_bucket, dest_bucket,
      ↪ file_name, mc))

  print(json.dumps(names))
```

Listing 6: File processimages Part I

processimages Function in Python II

```python
def convert_push(source_bucket, dest_bucket, file_name,
    ↪ mc):
  mc.fget_object(source_bucket, file_name, "/tmp/" +
    ↪ file_name)

  f = open("/tmp/" + file_name, "rb")
  input_image = f.read()

  # call function for b/w conversion
  r = requests.post("http://gateway:8080/function/
    ↪ convertbw", input_image)

  # write to temporary file
  dest_file_name = get_temp_file()
  f = open("/tmp/" + dest_file_name, "wb")
  f.write(r.content)
  f.close()

  # sync to Minio
  mc.fput_object(dest_bucket, dest_file_name, "/tmp/"+
    ↪ dest_file_name)

  return dest_file_name

def get_temp_file():
  uuid_value = str(uuid.uuid4())
  return uuid_value
```

Listing 7: File processimages Part II

Building and Deploying the Functions

Build the Functions:

```
$ faas-cli build -f loadimages.yml
$ faas-cli build -f processimages.yml
```

Deploy the Functions:

```
$ faas-cli deploy -f loadimages.yml
$ faas-cli deploy -f processimages.yml
```

Downloading and Converting the images

Download images into the `incoming` Bucket:

```
$ echo '{
"urls": [
"https://images.pexels.com/photos/72161/pexels-photo
    -72161.jpeg?dl&fit=crop&w=640&h=318",
"https://images.pexels.com/photos/382167/pexels-photo
    -382167.jpeg?dl&fit=crop&w=640&h=337"]
}' | faas invoke loadimages
```

Convert the images to grey and store in `processed` Bucket:

```
$ echo '["b0f38ebc-675c-43c1-ada7-8fb95dccee57", "34
    d0ad5d-9a24-4b32-bc3e-25337f6f2f5d"]' | faas invoke
    processimages
```

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Incoming Bucket in Minio

Figure: Incomming Bucket

Processed Bucket in Minio
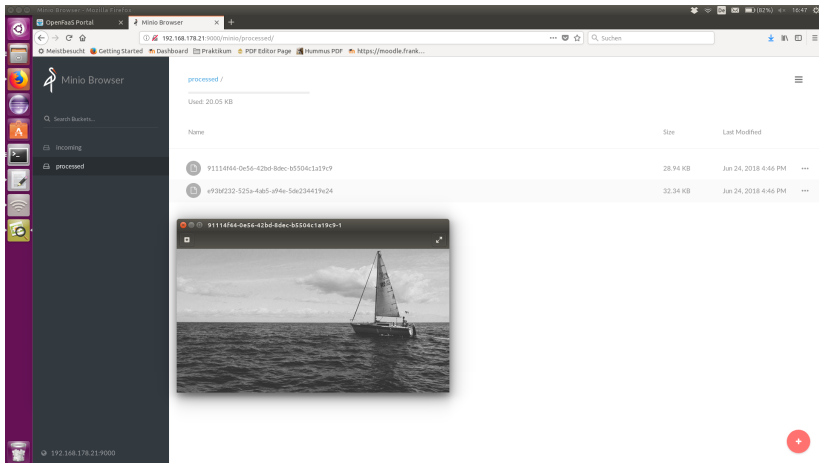
Figure: Processed Bucket

OpenWhisk

Figure: OpenWhisk [8]

OpenWhisk

- Open Source Platform
- Functions can be deployed in a production ready environment
- Support for multiple languages: JavaScript, Python 2, Python 3, PHP, Ruby, Swift
- C, C++, Go programs need to be compiled before upload, Java programs need to be uploaded as JAR-Archives

OpenWhisk Architecture I

Figure: OpenWhisk Architecture
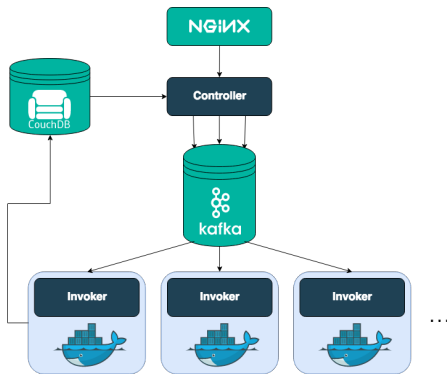
Source:
https://tinyurl.com/y7plrxbw

OpenWhisk Architecture [8]

Components:

- **Nginx**

- **Controller**

- **Kafka**

- **CouchDB**

- **Invoker**

OpenWhisk Architecture II

Figure: OpenWhisk Architecture
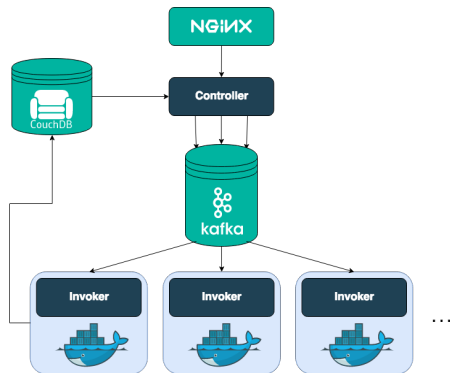
Source:
https://tinyurl.com/y7plrxbw

- **Nginx**
  - Loadbalancer for incoming requests
  - Forwarding requests to the controller

- **Controller**
  - Checks incoming requests
  - Controls the further action

- **Kafka**
  - Publish-Subscribe Messaging Service
  - Queues the requests

OpenWhisk Architecture III

Figure: OpenWhisk Architecture
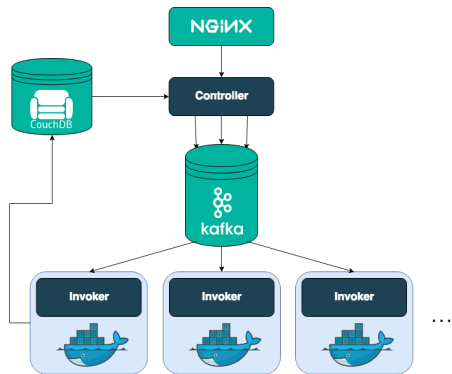
Source:
https://tinyurl.com/y7plrxbw

- **CouchDB**
  - Authentication of requests (permission checking)
  - Stores information on the imported Functions

- **Invoker**
  - Docker Container(s) running the Function
  - Each Invoker can be paused for faster request fullfilment

Installing OpenWhisk

There are 3 ways of installing OpenWhisk:

- In a Docker Container

- As a virtual machine using vagrant and e.g. VirtualBox

- Inside a Kubernetes Cluster

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Installing OpenWhisk – Docker Container

Install OpenWhisk as Docker Containers:

```
$ git clone https://github.com/apache/incubator—
    openwhisk—devtools.git
$ cd incubator—openwhisk—devtools/docker—compose
$ make quick—start
```

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Installing OpenWhisk – vagrant

Install OpenWhisk with vagrant and VirtualBox:

```
$ git clone ––depth=1 https://github.com/apache/
    incubator–openwhisk.git openwhisk
$ cd openwhisk/tools/vagrant
$ ./hello
```

Installing OpenWhisk – Kubernetes

Install OpenWhisk inside a Kubernetes Cluster:

```
$ minikube start --memory 4096 --kubernetes-version v1
    .10.5
$ minikube ssh -- sudo ip link set docker0 promisc on
$ kubectl label nodes --all openwhisk-role=invoker
$ helm init --wait
$ kubectl create clusterrolebinding tiller-cluster-admin
    \
--clusterrole=cluster-admin --serviceaccount=kube-system:
    default
$ git clone https://github.com/apache/incubator-openwhisk
    -deploy-kube
$ helm install ./incubator-openwhisk-deploy-kube/helm/
    openwhisk/ \
--name openwhisk --wait --timeout 900 \
--set whisk.ingress.type=NodePort \
--set whisk.ingress.api_host_name=$(minikube ip) \
--set whisk.ingress.api_host_port=31001 \
--set nginx.httpsNodePort=31001
```

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Creating an application using MongoDB in OpenWhisk

Application Flow [4]

- The Application manages the Stock of a Market
- For this task it stores the data in a database
- The Application recieves parameters for product ID and number of items
- The Application consists of **OpenWhisk** and **MongoDB** (NoSQL) database

Functions as a Service mit OpenWhisk. Henry-Norbert Cocos, Marcus Legendre, Christian Baun. iX 12/2018, S.126-130, ISSN: 0935-9680

---

[4]Source Code available at: `https://github.com/OrangeFoil/openwhisk-examples/tree/master/inventory`

Creating a function in OpenWhisk

```python
import pymongo

mongo_url = 'mongodb+srv://user:password@example.org/
    ↪ database'
mongodb_client = pymongo.MongoClient(mongo_url)
mongodb = mongodb_client.my_database


def main(params):
product_id = params['product_id']
stock_change = int(params['stock_change'])

result = mongodb.inventory.find_one_and_update(
{'product_id': product_id},
{'$inc': {'count': stock_change}},
upsert=True,
return_document=pymongo.collection.ReturnDocument.AFTER
)

return {
'product_id': result['product_id'],
'count': result['count']
}
```

Listing 8: File __main__.py

Deploy the function in OpenWhisk

Deploy the function in OpenWhisk:

```
$ mkdir tmp-build
$ cp __main__.py tmp-build/
$ pip3 install dnspython pymongo -t tmp-build/
$ cd tmp-build
$ zip -r ../exec.zip ./*
$ cd ..
```

As a ZIP-File. . .

In order to run the function in OpenWhisk, the dependencies
`dnspython` and `pymongo` need to be installed with the Python
Package Manager `pip3` (Python Installs Packages)! Those
dependencies are stored with the application inside a ZIP-File.

KFURT
ERSITY
OF APPLIED SCIENCES

Creating Actions and Triggers

In the OpenWhisk platform events are characterized by **Trigger**.
An **Action** is used to invoke the function. A **Rule** binds an **Action**
to a **Trigger**.

Creating an action for updating the database:

```
$ wsk action create updateInventory exec.zip --kind
    python:3
```

Creating Triggers for increment and decrement operations:

```
$ wsk trigger create itemSold --param stock_change -1
$ wsk trigger create itemRestocked --param stock_change
    1
```

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Creating a Rule and Trigger an Event

In order to invoke an action, a trigger needs to be fired. The Rule
restockRule and saleRule are bound to the updateInventory
action.

Creating A Rule to combine Triggers and Actions:

```
$ wsk rule create restockRule itemRestocked
    updateInventory
$ wsk rule create saleRule itemSold updateInventory
```

Firing the Trigger:

```
$ wsk trigger fire itemRestocked ——param product_id 1234
    ——param stock_change 100
$ wsk trigger fire itemSold ——param product_id 5678
```

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

Trigger an Event

```
$ wsk trigger fire itemRestocked -p product_id 42 -p stock_change 10
ok: triggered /_/itemRestocked with id f3204cd9dc16412ba04cd9dc16212b02
$ wsk activation result --last
{
    "count": 39,
    "product_id": 42
}
$ wsk trigger fire itemSold -p product_id 42
ok: triggered /_/itemSold with id ed476347731f4f75876347731fdf751f
$ wsk trigger fire itemSold -p product_id 42
ok: triggered /_/itemSold with id c1a15500f7e149c1a15500f7e1a9c144
$ wsk activation result --last
{
    "count": 37,
    "product_id": 42
}
```

Figure: Action Invocation

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

OpenWhisk Benefits

OpenWhisk has the following benefits:

- Open Source
- Deployment of functions
- Autoscaling
- Robust and flexible (ideal for production)
- Migration to public offering IBM Cloud Functions possible

OpenWhisk and IBM Cloud Functions

OpenWhisk is the basis of the public offering IBM Cloud Functions. Therefore applications developed for OpenWhisk can be ported to IBM Cloud Functions and vice versa without additional refactoring. This fact gives Enterpises more flexibility in developing their service offering!

KFURT
ERSITY
OF APPLIED SCIENCES

Conclusion

Function as a Service characteristics:

- More fine grained buisiness model (payment per invocation)
- Functions have no side effects, stateless model
- Scaling of functions with Container Virtualization (Docker)
- Shorter developement and deployment cycles (DevOps)
- Suitable technology for microservices

Outlook

FaaS is a new technology in the field of Cloud Platform Services.
With the developement of IoT, Smart Homes and other
event-driven technologies the number of private FaaS Frameworks
and public FaaS offerings will grow in the near future!

FRANKFURT
UNIVERSITY
OF APPLIED SCIENCES

References I

[1] M. Eder, "Hypervisor- vs. Container-based Virtualization," https://www.net.in.tum.de/fileadmin/TUM/NET/NET-2016-07-1/ NET-2016-07-1_01.pdf, accessed April 20, 2019.

[2] "Docker Docs," https://docs.docker.com/, accessed April 20, 2019.

[3] "Docker Swarm," https://docs.docker.com/engine/swarm/, accessed April 20, 2019.

[4] "Golem – Mehr Zeit für den Code," https://www.golem.de/news/ serverless-computing-mehr-zeit-fuer-den-code-1811-137516.html, accessed April 20, 2019.

[5] "AWS Lambda," https://aws.amazon.com/lambda/, accessed April 20, 2019.

[6] "Google Cloud Functions BETA," https://cloud.google.com/functions/, accessed April 20, 2019.

References II

[7] "IBM Cloud Functions," https://www.ibm.com/cloud/functions, accessed April 20, 2019.

[8] "Apache OpenWhisk," https://openwhisk.apache.org/, accessed April 20, 2019.

[9] I. Baldini, P. C. Castro, K. S. Chang, P. Cheng, S. J. Fink, V. Ishakian, N. Mitchell, V. Muthusamy, R. M. Rabbah, A. Slominski, and P. Suter, "Serverless Computing: Current Trends and Open Problems," *CoRR*, vol. abs/1706.03178, 2017. [Online]. Available: http://arxiv.org/abs/1706.03178

[10] T. Lynn, P. Rosati, A. Lejeune, and V. Emeakaroha, "A Preliminary Review of Enterprise Serverless Cloud Computing (Function-as-a-Service) Platforms," in *2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, Dec 2017, pp. 162–169.

[11] "OpenFaaS - Serverless Functions Made Simple,"
https://docs.openfaas.com/, accessed April 20, 2019.

[12] "Your Serverless Raspberry Pi cluster with Docker,"
https://blog.alexellis.io/your-serverless-raspberry-pi-cluster/,
accessed April 20, 2019.