FaaSET: A Jupyter notebook to streamline every *facet* of serverless development

Robert Cordingly, Wes Lloyd
rcording@uw.edu, wlloyd@uw.edu

April 9th 2022

School of Engineering and Technology
University Of Washington, Tacoma
5th Workshop on Hot Topics in Cloud Computing Performance (HotCloudPerf 2022)

Outline

Serverless Computing

Serverless Function-as-a-Service platforms offer many appealing features:

- No infrastructure management
- Automatic scaling
- Fine grained usage-based billing models

But packaging, deploying, testing and running experiments across multiple FaaS platforms leads to unique challenges:

- Vendor lock-in requires specific tools, services, application design

+ FaaSET

Outline

- Introduction
- Supporting Tools
- FaaSET Workflow
  - Develop, Deploy, Test
  - Execute Experiments
  - Data Analysis
- Evaluation
- Conclusions

Supporting Tools

SAAF

SAAF supports profiling Function-as-a-Service (FaaS) workload performance, resource utilization, and infrastructure enabling accurate performance and cost characterizations.

SAAF supports profiling deployments to AWS Lambda, Google Cloud Functions, IBM Cloud Functions, OpenFaaS, and Azure Functions written in Java, Python, Javascript, and BASH.

FaaS Runner

FaaS Runner is a client-side Python application used in conjunction with SAAF and the FaaSET notebook.

FaaS Runner can invoke large batches of functions synchronously, or asynchronously and orchestrate complex pipelines of functions.

Experiments are defined using functions and experiment files that explain how functions show be executed and how the results from SAAF should be processed.

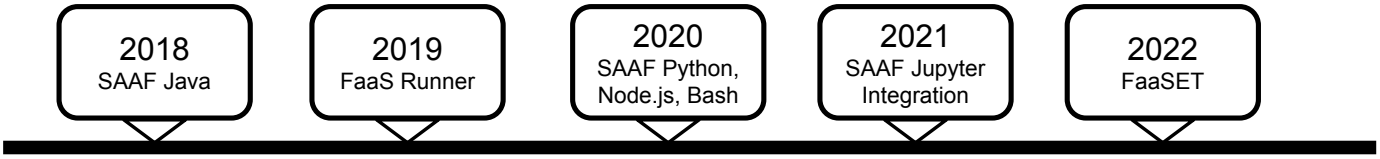

Outline

- Introduction
- Supporting Tools
- FaaSET Workflow
  - Develop, Deploy, Test
  - Execute Experiments
  - Data Analysis
- Evaluation
- Conclusions

Function-as-a-Service Experiment Toolkit (FaaSET)

FaaSET provides aggregated tools to write, deploy, test, and run experiments on FaaS platforms all in a unified Jupyter Notebook workspace.

FaaSET supports many commercial and open source FaaS platforms:

- AWS Lambda (x86 and ARM64 with or w/o Docker Containers)
- Google Cloud Functions (Gen 1 and 2)
- IBM Cloud Functions/OpenWhisk (with or w/o Docker Containers)
- Azure Cloud Functions
- OpenFaaS

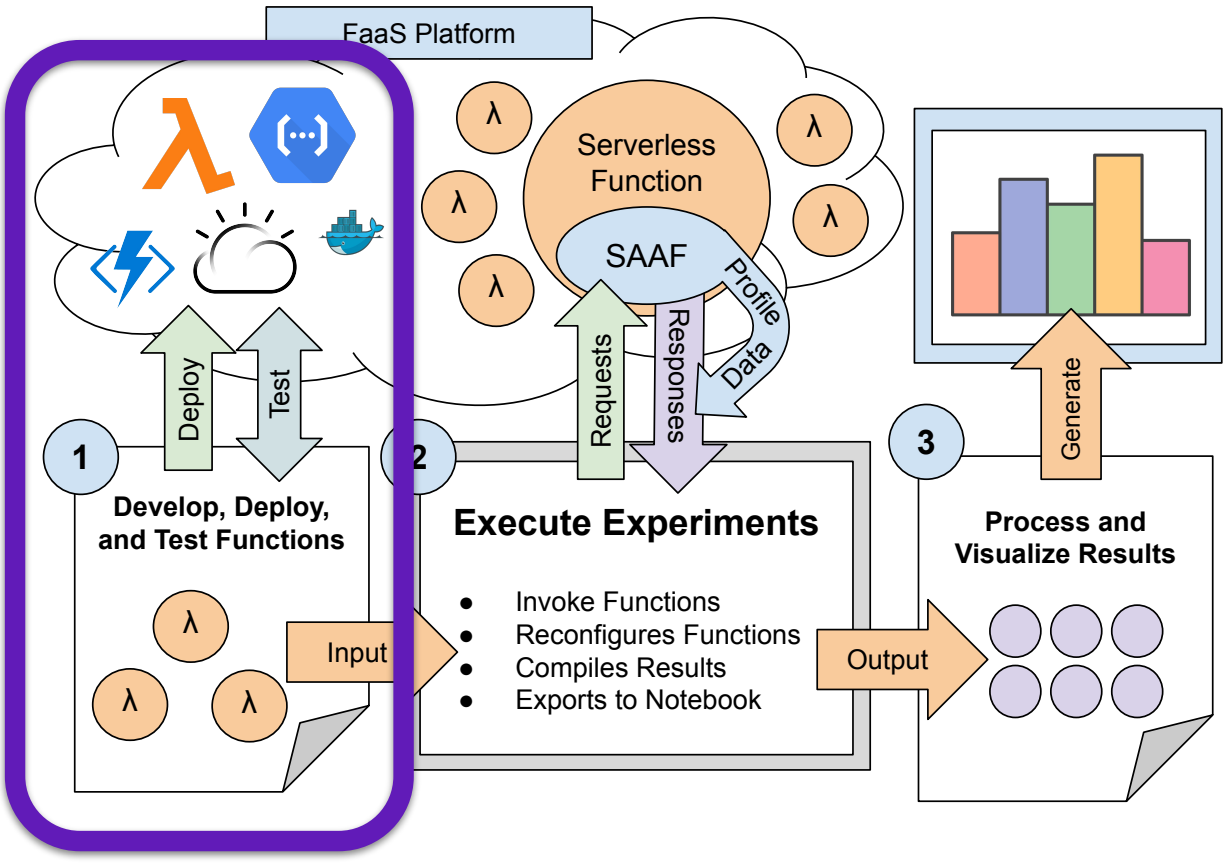

Function Development and Deployment with FaaSET

```python
def hello_world(request, context):
    return {"message": "Hello " + str(request["name"]) + "!"}
```


Function Development and Deployment with FaaSET

```python
def hello_world(request, context):
    return {"message": "Hello " + str(request["name"]) + "!"}
```

```python
import FaaSET
@FaaSET.cloud_function(platform="AWS", config={"memory":256})
def hello_world(request, context):
    return {"message": "Hello " + str(request["name"]) + "!"}
```

Function Development and Deployment with FaaSET

```
def hello_world(request, context):
   return {"message": "Hello " + str(request["name"]) + "!"}
```

```
import FaaSET
@FaaSET.cloud_function(platform="AWS", config={"memory":256})
def hello_world(request, context):
   return {"message": "Hello " + str(request["name"]) + "!"}
```


Function Development and Deployment with FaaSET

```
def hello_world(request, context):
   return {"message": "Hello " + str(request["name"]) + "!"}
```

```
import FaaSET
@FaaSET.cloud_function(platform="AWS", config={"memory":256})
def hello_world(request, context):
   return {"message": "Hello " + str(request["name"]) + "!"}
```

```
hello_world({'name': 'Bob'}, None)
```

```
>> Deploying to AWS Lambda…
>> {"message": "Hello Bob!"}
```

Function Development and Deployment with FaaSET

```python
def hello_world(request, context):
    return {"message": "Hello " + str(request["name"]) + "!"}
```

```python
import FaaSET
@FaaSET.cloud_function(platform="AWS", config={"memory":256})
def hello_world(request, context):
    return {"message": "Hello " + str(request["name"]) + "!"}
```

```python
hello_world({'name': 'Bob'}, None)
```

```
>> Deploying to AWS Lambda…
>> {"message": "Hello Bob!"}
```


Write Once Deploy Across Multiple Platforms

```python
@cloud_function(platform="AWS")
def hello_world(request, context):
    return {"message": "Hello Lambda!"}
```

```python
@cloud_function(platform="GCF")
def hello_world(request, context):
    return {"message": "Hello Google!"}
```

```python
@cloud_function(platform="IBM")
def hello_world(request, context):
    return {"message": "Hello IBM!"}
```

```python
@cloud_function(platform="Azure")
def hello_world(request, context):
    return {"message": "Hello Azure!"}
```

```python
@cloud_function(platform="AWS ARM")
def hello_world(request, context):
    return {"message": "Hello Lambda!"}
```

```python
@cloud_function(platform="OpenFaaS")
def hello_world(request, context):
    return {"message": "Hello OpenFaaS!"}
```

```python
@cloud_function(platform="GCF Gen2")
def hello_world(request, context):
    return {"message": "Hello Google!"}
```

```python
@cloud_function(platform="IBM Docker")
def hello_world(request, context):
    return {"message": "Hello IBM!"}
```

FaaSET Function Management Features:

- FaaSET tracks changes and only deploys when functions are modified
- Supports a simplified function development workflow using a Notebook while also allowing full control over function source code, packaging, and deployment
- Invoke existing functions already deployed in your notebooks
- Reconfigure functions on the fly without rebuilding package/containers
- Functions are automatically defined on startup allowing immediate access

FaaSET Function Management Features:

- FaaSET tracks changes and only deploys when functions are modified
- Supports a simplified function development workflow using a Notebook while also allowing full control over function source code, packaging, and deployment
- Invoke existing functions already deployed in your notebooks
- Reconfigure functions on the fly without rebuilding package/containers
- Functions are automatically defined on startup allowing immediate access

```
@FaaSET.cloud_function(platform="AWS")
def hello_world(request, context):
    return {"message": "Hello Bob!"}
```

```
hello_world({}, None)
hello_world({}, None)
```

```
>> Deploying to AWS Lambda…
>> {"message": "Hello Bob!"}
>> {"message": "Hello Bob!"}
```

Calling twice only deploys on the first call

```
@FaaSET.cloud_function(platform="AWS")
def hello_world(request, context):
    return {"message": "Hola Bob!"}
```

```
hello_world({}, None)
>> Deploying to AWS Lambda…
>> {"message": "Hola Bob!"}
```

Redeploys after function is changed

FaaSET Function Management:

- FaaSET tracks changes and only deploys when functions are modified
- Supports a simplified function development workflow using a Notebook while also allowing full control over function source code, packaging, and deployment
- Invoke existing functions already deployed in your notebooks
- Reconfigure functions on the fly without rebuilding package/containers
- Functions are automatically defined on startup allowing immediate access

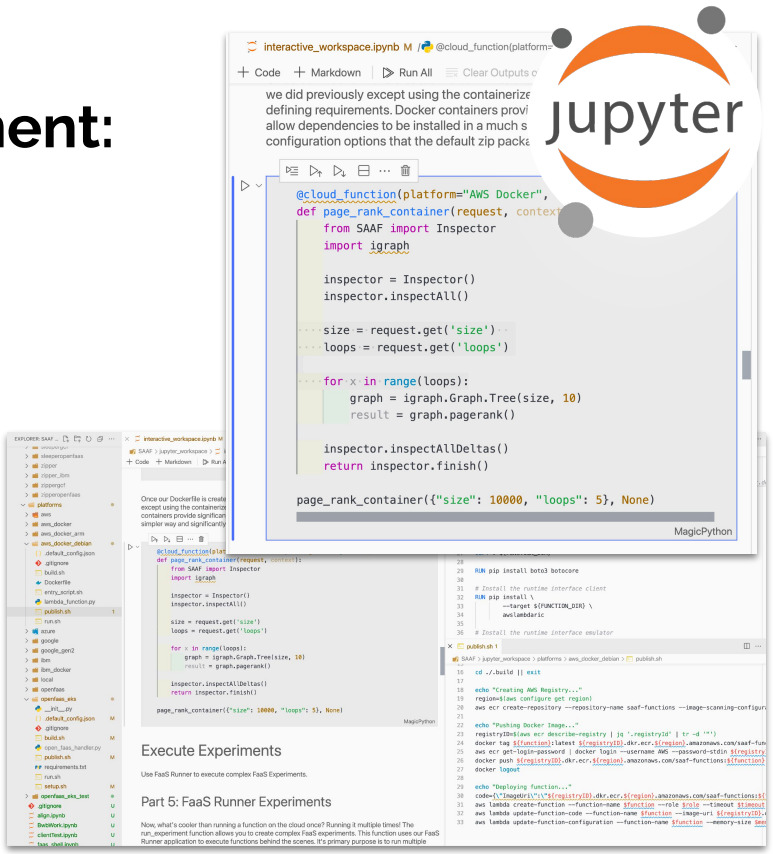

FaaSET Function Management Features:

- FaaSET tracks changes and only deploys when functions are modified
- Supports a simplified function development workflow using a Notebook while also allowing full control over function source code, packaging, and deployment
- Invoke existing functions already deployed in your notebooks
- Reconfigure functions on the fly without rebuilding package/containers
- Functions are automatically defined on startup allowing immediate access

```
@cloud_function(platform="IBM",deploy=False)
def java_function(request, context):
    pass
```

```
@cloud_function(platform="Azure",deploy=False)
def nlp_pipeline(request, context):
    pass
```

```
@cloud_function(platform="GCF",deploy=False)
def imageprocessor(request, context):
    pass
```

```
@cloud_function(platform="OpenFaaS",deploy=F…
def node_info(request, context):
    pass
```

FaaSET Function Management Features:

- FaaSET tracks changes and only deploys when functions are modified
- Supports a simplified function development workflow using a Notebook while also allowing full control over function source code, packaging, and deployment
- Invoke existing functions already deployed in your notebooks
- Reconfigure functions on the fly without rebuilding package/containers
- Functions are automatically defined on startup allowing immediate access

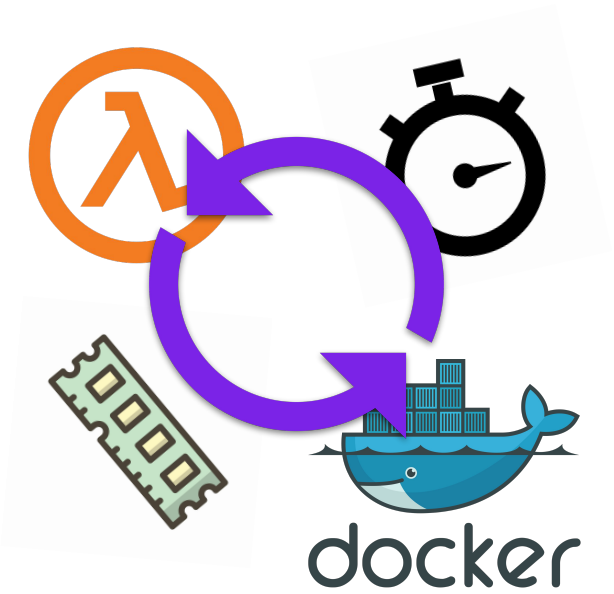

FaaSET Function Management Features:

- FaaSET tracks changes and only deploys when functions are modified
- Supports a simplified function development workflow using a Notebook while also allowing full control over function source code, packaging, and deployment
- Invoke existing functions already deployed in your notebooks
- Reconfigure functions on the fly without rebuilding package/containers
- Functions are automatically defined on startup allowing immediate access

```
import FaaSET
>> Loading platforms…
>> Platforms: AWS, GCF, IBM, Azure, AWS Docker…
>> Loading functions…
>> Functions: hello_world, nlp_pipeline, node_info…
```

```
FaaSET.hello_world({}, None)
>> {"message": "Hola Bob!"}
```

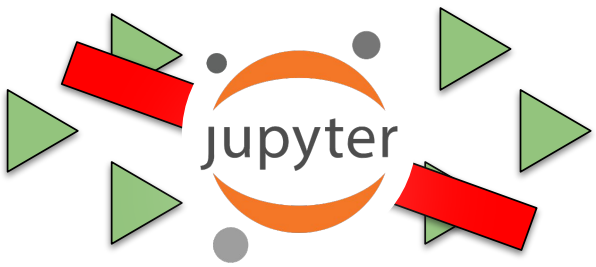

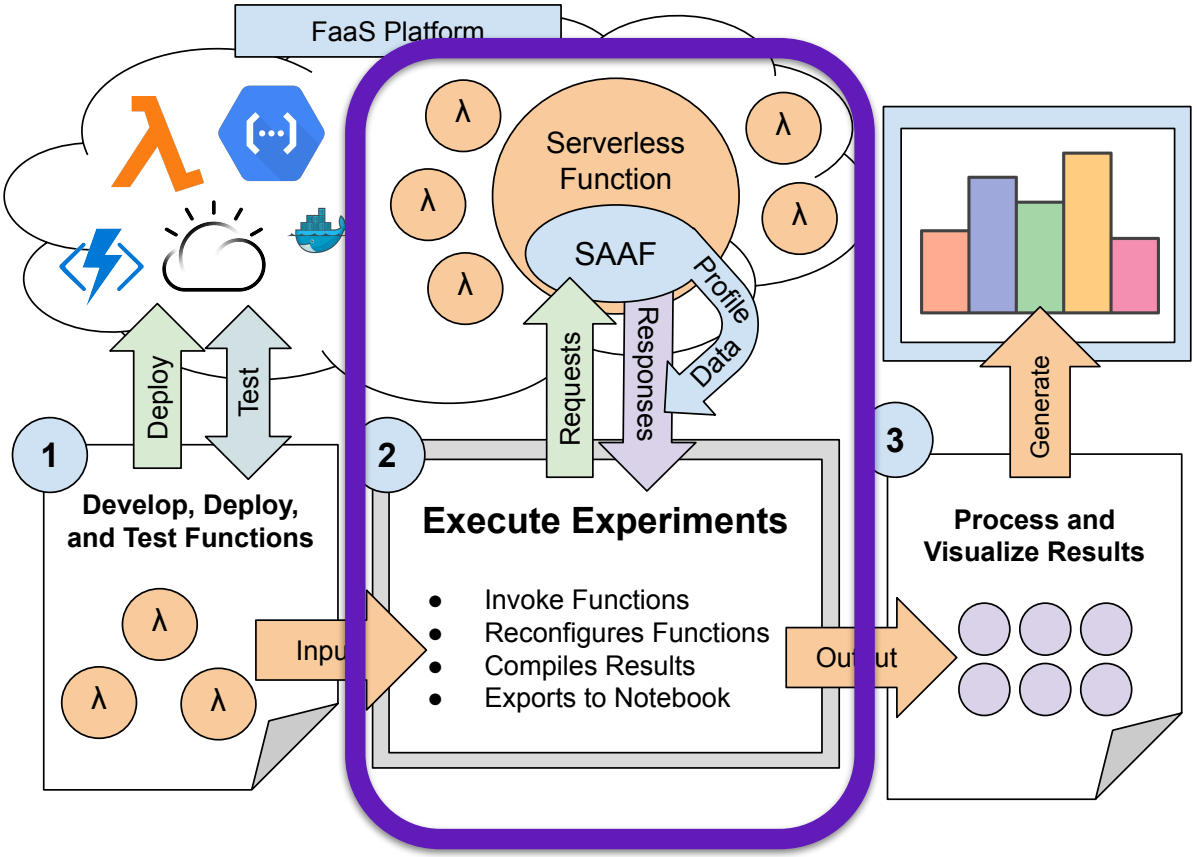

Experiments with FaaSET + FaaS Runner

- FaaSET includes the FaaS Runner tool
- Define experiment parameters such as number of runs, threads, payloads, and call order
- Utilize FaaSET's reconfiguration tool to automate complex experiments
- Results are imported into the Notebook as a Pandas dataframes

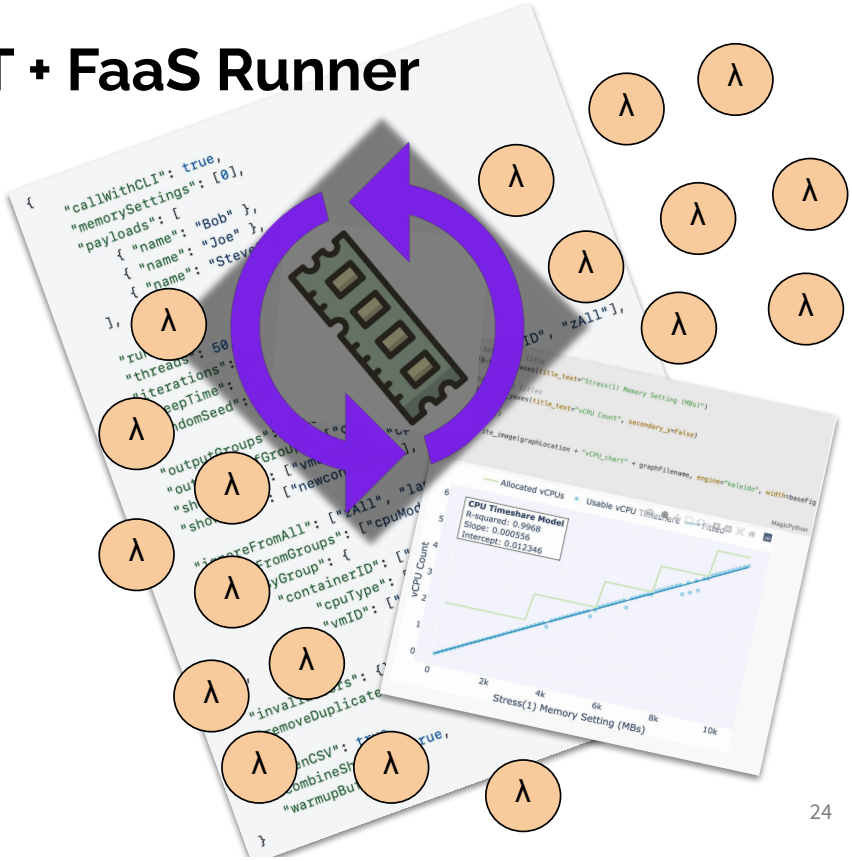

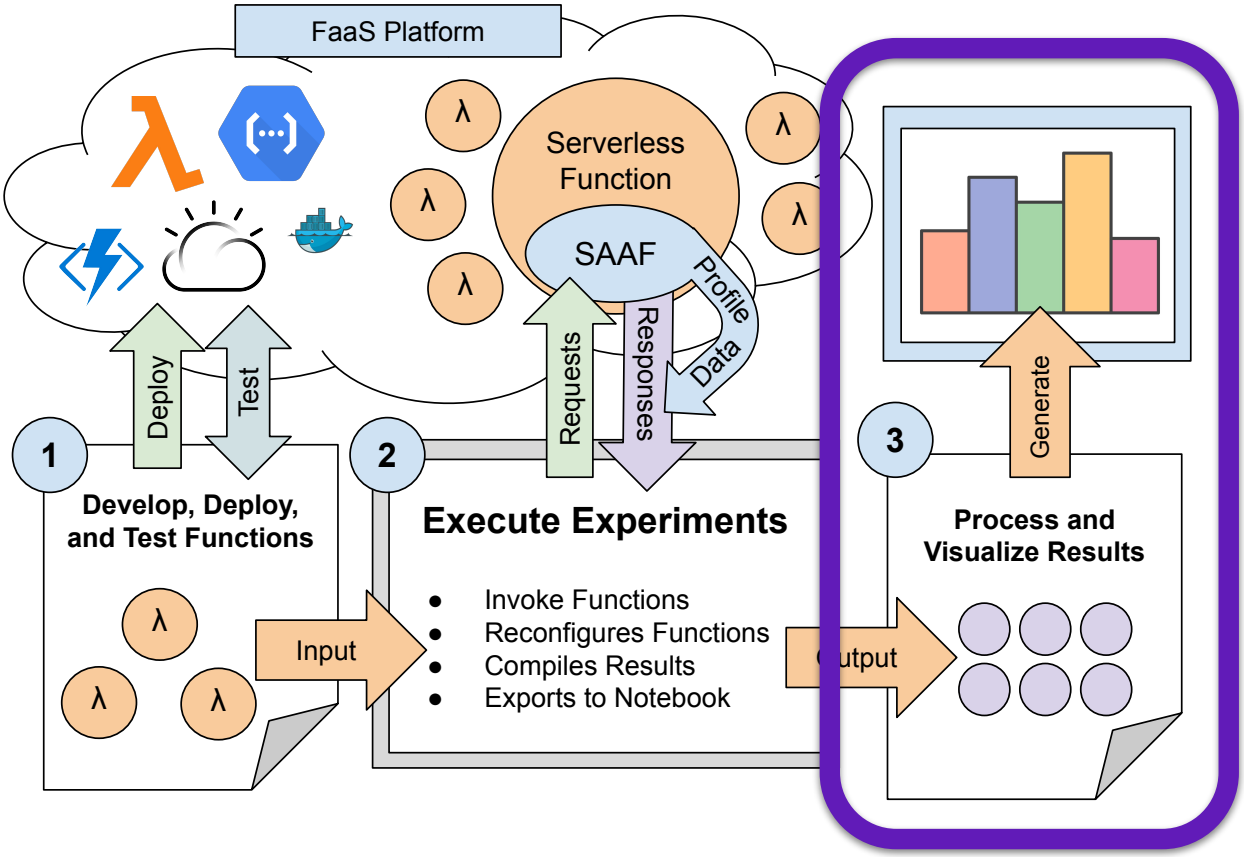

Data Analysis with FaaSET + SAAF

- FaaSET integrates the Serverless Application Analytics Framework
  - SAAF is placed in the deployment package of the function
  - Collects information about the host infrastructure, resource utilization metrics, and FaaS platform
- Combining FaaSET and SAAF improves accessibility of observations into FaaS platforms from a Jupyter Notebook:
  - Tenancy, warm/cold infrastructure, latency, round-trip time, and more

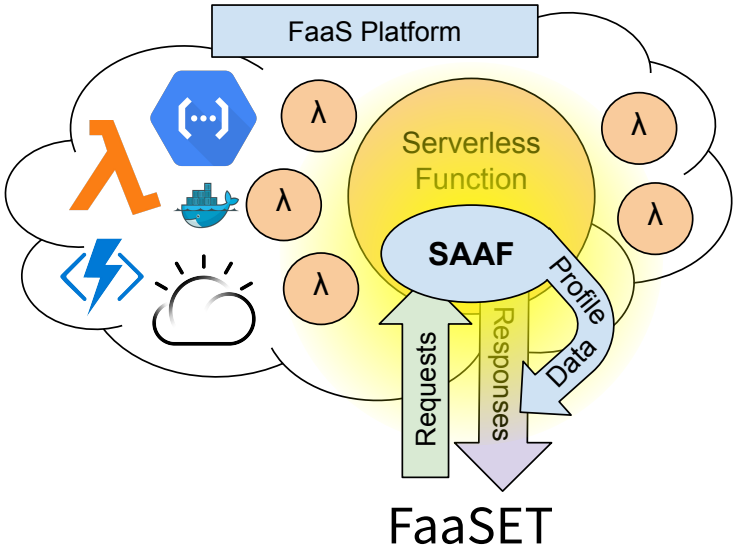

Data Analysis in FaaSET

- Since FaaSET is designed to be used inside Jupyter Notebooks, existing libraries can be used for data analysis:
  - Numpy
  - Pandas
  - Matplotlib
  - Plotly
  - Scikit-learn
  - Scipy
- FaaS Runner experiments directly output results in Pandas dataframes

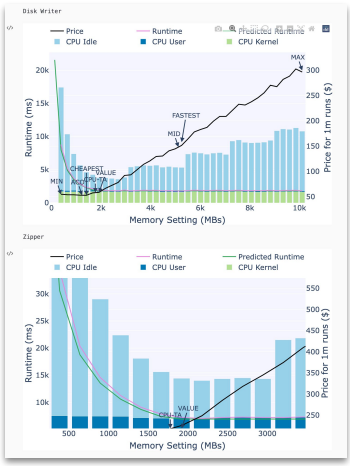

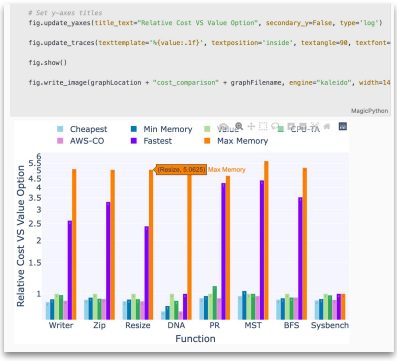

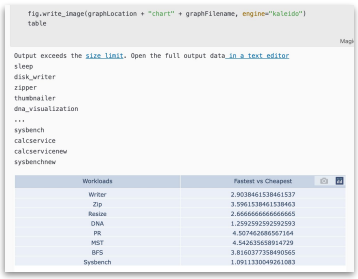

Outline

- Introduction
- Supporting Tools
- FaaSET Workflow
  - Develop, Deploy, Test
  - Execute Experiments
  - Data Analysis
- Evaluation
- Conclusions

Hosting FaaS Experiments

- To run experiments on FaaS platforms a host is required to invoke the functions. This can be a local PC, powerful cloud virtual machine, or Jupyter Notebook-as-a-Service platforms (e.g. Google Colaboratory).

- We compared the performance of running 1,000 concurrent function invocations on AWS Lambda using a local PC (i9-9900k CPU, 1Gbps Network), powerful EC2 instance (c5.metal), and the free tier of Google Colab.

Host	CPU	vCPUs	Memory	Network
PC	i9-9900k @ 3.6 GHz	10	32 GBs	~1 Gbps
c5.metal	Xeon 8275L @ 3 GHz	96	192 GBs	25 Gbps
Colab	Xeon @ 2.2 Ghz	2	12 GBs	Unknown

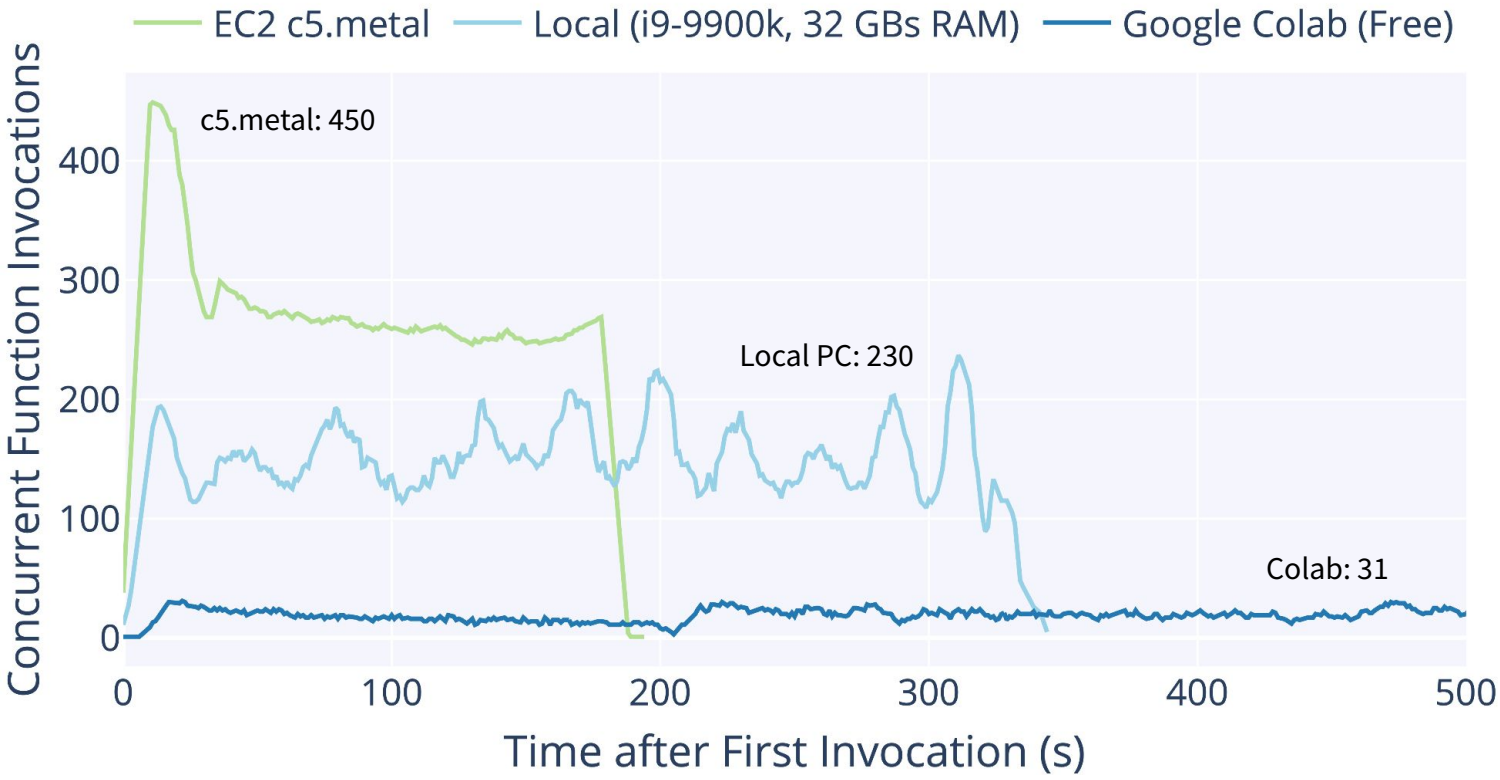

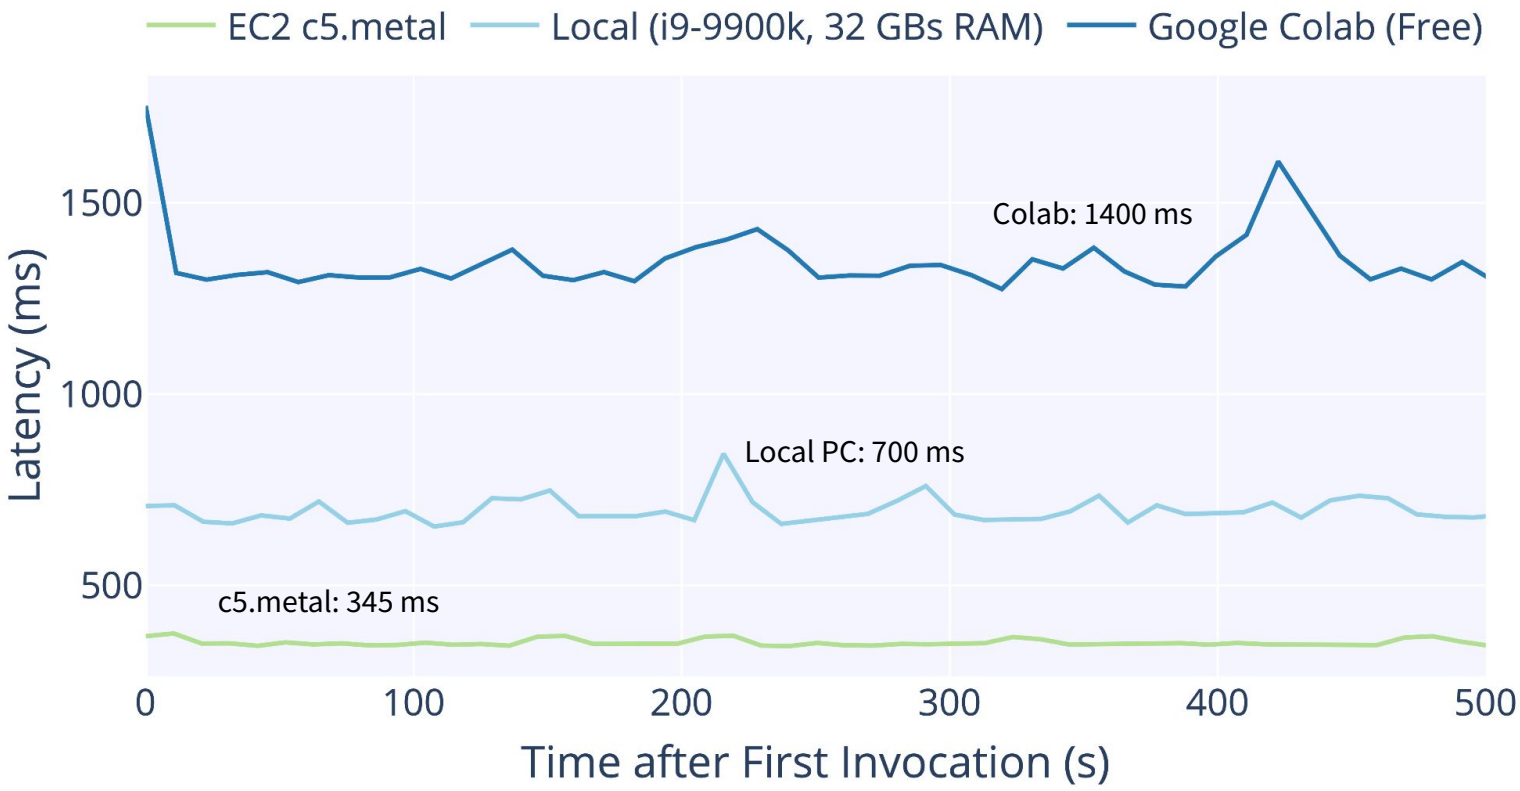

Outline

- Introduction
- Supporting Tools
- FaaSET Workflow
  - Develop, Deploy, Test
  - Execute Experiments
  - Data Analysis
- Evaluation
- Conclusions

Conclusions

- FaaSET provides many features for deploying, testing, and running experiments on FaaS platforms.
  - FaaSET's goal is to provide a streamlined development environment for developers or researchers running experiments on FaaS platforms.

- While Google Colaboratory has the worse performance, it is the easiest to set up, is free, and has useful collaboration features making it great for small experiments:
  Try FaaSET by visiting: https://bit.ly/3DNVeOE

  Note: FaaSET on Google Colaboratory only supports AWS Lambda

Thank You!

Get FaaSET on GitHub: https://github.com/wlloyduw/SAAF
Try FaaSET in Google Colab: https://bit.ly/3DNVeOE

Any Questions?

Get FaaSET on GitHub: https://github.com/wlloyduw/SAAF
Try FaaSET in Google Colab: https://bit.ly/3DNVeOE

FaaSET: A Jupyter notebook to streamline every *facet* of serverless development

Robert Cordingly, Wes Lloyd
rcording@uw.edu, wlloyd@uw.edu

April 9th 2022

School of Engineering and Technology
University Of Washington, Tacoma
5th Workshop on Hot Topics in Cloud Computing Performance (HotCloudPerf 2022)

Results Summary

- The host evaluation experiment found:
  - Maximum concurrent function calls:
    - c5.metal: 450
    - i9 PC: 230
    - Google Colab: 31
  - Average total Latency:
    - c5.metal: 345 ms
    - i9 PC: 700 ms
    - Google Colab: 1400 ms