



# Google Summer of Code



## Final Project Report

By Abhishek Kumar

Mentors: Marcin Copik, Oana Rosca

Organization:

Scalable Parallel Computing  
Laboratory

Project:

[New serverless benchmarks](#)

This summer, I contributed to the SPCL SEBS project where I contributed primarily to the addition of a new serverless platform Knative to the SEBS project, and adding some benchmarks to the project as well.

## My Work during the GSOC period at a glance with PR links:

### 1. Knative Integration PR:

<https://github.com/spcl/serverless-benchmarks/pull/206> (COMPLETED ✓)

```
(python-venv) abhishek@abhishek-HP-Pavilion-Gaming-Laptop-15-ec2xxx:/spcl/serverless-benchmarks$ ./sebs.py benchmark invoke 503.graph-bfs test --config config/example.json --deployment knative --language python --language-version 3.9 --verbose
[11:28:58. 940524] Knative.Resources-bdf8 Using user-provided Docker registry for Knative.
[11:28:58. 940524] Knative.Resources-bdf8 Using user-provided configuration of storage for Knative.
[11:28:58. 940705] Knative.Resources-bdf8 User-provided configuration is different from cached storage, we will update existing Knative function.
[11:28:58. 951734] Knative-8elf We found 2 existing deployments! If you want to use any of them, please abort, and provide the resource id in your input config.
[11:28:58. 951792] Knative-8elf Deployment resource IDs in the cloud: ['5aa5816c', '91317bf2']
[11:28:58. 951918] Knative-8elf Generating unique resource name 7e647b9e
[11:28:58. 951918] Knative-8elf Creating new Docker image for 7e647b9e
[11:28:58. 961281] Knative-8elf Created Docker image for 7e647b9e
[11:28:58. 961281] minio,Minio-f783 Created bucket abhishekmarks-7e647b9e
[11:28:58. 962417] Benchmark-3bf3 Building benchmark 503.graph-bfs. Reason: no cached code package.
[11:28:58. 964331] Benchmark-3bf3 There is no Docker build image for knative run in python, thus skipping the Docker-based installation of dependencies.
[11:28:58. 965985] Knative-8elf Image abhidocker12:function.knative.503.graph-bfs.python-3.9 doesn't exist in the registry, building Docker package for 503.graph-bfs.
[11:28:58. 966041] Knative-8elf Running build command: func build --registry abhidocker12 --path /home/abhishek/spcl/serverless-benchmarks/503.graph-bfs/code/python/3.9
[11:29:22. 339984] Knative-8elf Build output:
[11:29:22. 339984] Knative-8elf Successful build and pushed Docker image abhidocker12:function.knative.503.graph-bfs.python-3.9 to registry: Docker Hub.
[11:29:22. 339984] Knative-8elf Creating new function! Reason: function 503.graph-bfs-python-3.9 not found in cache.
[11:29:22. 339604] Knative-8elf Creating new function! Reason: function 503.graph-bfs-python-3.9 not found in cache.
[11:29:22. 339613] Knative-8elf Building Knative function
[11:29:22. 396481] Knative-8elf Deploying new Knative function graph-bfs
[11:30:34. 505231] Knative-8elf Deployment succeeded:
[11:30:34. 505239] Knative-8elf
[11:30:34. 574041] Knative-8elf Function deployment URL fetched successfully.
[11:30:34. 574041] Knative-8elf Function deployment URL: http://graph-bfs.default.127.0.0.1.sslip.io
[11:30:34. 641494] Knative-HTTPTrigger-935d Invoke function http://graph-bfs.default.127.0.0.1.sslip.io
[11:30:38. 717401] Knative-HTTPTrigger-935d Invoke of function was successful
[11:30:38. 717592] Seb5-a441 Beginning repetition 2/5
[11:30:38. 717773] Knative-HTTPTrigger-935d Invoke function http://graph-bfs.default.127.0.0.1.sslip.io
[11:30:38. 889910] Knative-HTTPTrigger-935d Invoke of function was successful
[11:30:38. 889904] Seb5-a441 Beginning repetition 3/5
[11:30:38. 890281] Knative-HTTPTrigger-935d Invoke function http://graph-bfs.default.127.0.0.1.sslip.io
[11:30:38. 890281] Seb5-a441 Beginning repetition 4/5
[11:30:38. 897364] Seb5-a441 Beginning repetition 4/5
[11:30:38. 897547] Knative-HTTPTrigger-935d Invoke function http://graph-bfs.default.127.0.0.1.sslip.io
[11:30:38. 904761] Knative-HTTPTrigger-935d Invoke of function was successful
[11:30:38. 905036] Seb5-a441 Beginning repetition 5/5
[11:30:38. 905273] Knative-HTTPTrigger-935d Invoke function http://graph-bfs.default.127.0.0.1.sslip.io
[11:30:38. 911355] Knative-HTTPTrigger-935d Invoke of function was successful
[11:30:38. 912576] Seb5-a441 Save results to /home/abhishek/spcl/serverless-benchmarks/experiments.json
[11:30:38. 912576] Seb5-a441
```

### 2. PDF Generator Benchmark PR:

<https://github.com/spcl/serverless-benchmarks/pull/217> (COMPLETED ✓)

### 3. Image Captioning Benchmark PR:

<https://github.com/spcl/serverless-benchmarks/pull/218> (COMPLETED ✓)

### 4. PDF Generator Benchmark Input data PR:

<https://github.com/spcl/serverless-benchmarks-data/pull/4> (COMPLETED ✓)

5. Image Captioning Generator Benchmark data PR:  
<https://github.com/spcl/serverless-benchmarks-data/pull/4> (COMPLETED ✓)
6. EXTRA'S ( Update Support for Benchmark Language Version )  
<https://github.com/spcl/serverless-benchmarks/pull/198> (COMPLETED ✓)

## Here is the detailed workflow of my work:

### Project workflow:

1. New Benchmarks:
  - PDF Generation, e.g., turning a website into PDF.
  - Image Caption Generator.
2. Adding support for Knative as a new serverless platform.
3. Testing to check if everything works fine or not.

### 1st - 7th June 2024:

This week:

1. Worked on testing out benchmarks on the Azure platform:
  - a. Fixed dependencies issues of the NodeJs, and updated the images that support new Azure function versions. Successfully tested out latest images against all benchmarks on Azure functions, current latest images are NodeJs 16, 18, and 20. The latest artifacts are fetched from here  
<https://mcr.microsoft.com/en-us/product/azure-functions/node/tags> (Only supported tags)
  - b. Updated the Azure function extension bundle to [4.0.0 , 5.0.0) from [1.\* , 2.0.0). The previous version was outdated and was giving warnings on the Azure platform.

Related Issue: <https://github.com/spcl/serverless-benchmarks/issues/196>

### Plan for next week:

1. Knative support research and plan updation on matter most channel for further review and discussion with a proper research review (Although I'll experiment with the benchmarks on the platform).
2. Reviewing Prajin's container support as discussed in today's meeting, if got stuck somewhere will try to reach Prajin for a quick meet to follow up.

## 14th June

- Knative research done! Moving forward with the integration part with SEBS, looked at the implementation of open whisk as suggested.
- This week was more of research and understanding the code workflow of Open Whisk and the prajin's PR.
- I'll raise a PR with a raw implementation of Knative by the end of Tuesday next week for a review.

### Next week's plan:

- Knative platform integration PR for review.
- Code changes after review.
- Benchmark addition ( To be discussed with mentors )

## 20th June

- Raised a skeleton PR for Knative integration by referring the knative serving but as discussed with Marcin Copik, I have to shift towards the Kn Func which is much easier.
- Working on Knative function creation, deployment, and invocation using the kn func or func ( will go with the kn func ).

### Next week's plan:

- To come up with a working main implementation for knative and updating the PR.
- Code changes after review.
- Benchmark additions.

## 27th June

- Tested benchmarks explicitly using the func CLI for knative services and invoked them ( Have done this in order to understand the underlying logic happening inside knative).
- Created the create, deploy and invok functions for the knative deployment.
- Working on the integration them in the main interface with SEBS.

Next week's plan:

- Completing the Knative integration.
- Adding a script for knative environment deployment from the user end, to ease user workflow.
- Adding documentation and PR review.

## 4th July

- Support for knative platform specific wrappers is done for both NodeJS and Python.
- Main integration part is done.
- Script for spinning up a knative enabled kubernetes cluster is also done(just working on few nit picks Macin suggested me).

### Next week's Plan

- Fully functional knative serverless platform (Testing part is to be done here)
- Documentation for the new serverless system.

## 11th July

- Fully functional Knative integration is done, ran benchmarks and worked as expected.
- Updated microbenchmarks importing issues for knative environment.
- Created a fully functional kind cluster script that will spinup the cluster for the knative environment which later can be used for benchmarks deployments.

### Next week's Plan

- A proper documented Knative PR will be raised by the end of Monday, This PR will be raised separately from the current integration one as there are plenty of changes in the current one ( Just to ease reviewing).
- Start working on the 2nd part of my GSOC which is adding new serverless benchmarks for SEBS.

## July 19

- Raised knative integration final PR with regression and docs for review.
- Ran some main benchmarks.
- Got some errors with deployments of some benchmarks trying to see what's going wrong.

### Next week plan:

- New serverless benchmark addition.
- Solving the current minor issues of knative work.

## July 26th

- Knative Re-deployment issue fixed.
- Tested out benchmarks as well, and are working.
- Regression support is added for knative.
- Researched about the Image captioning benchmark.

### Next week plan:

- Adding new benchmarks and testing them.
- If any changes needed on Knative PR, will work on it.

## Aug 2nd

- Shifted Knative function image built from pack to s2i.
- Added different base image during function development.
- Added support for dynamically fetching storage credentials for knative.
- Tested out the whole knative system and support is DONE.

- Researched about the text-to-speech conversion and image captioning generator benchmark.

**Next week plan:**

- Adding benchmark and testing them.

### **Aug 9th**

- Updated Knative Update function and update function configuration functions to add memory config of benchmarks in their respective functions yaml file.
- Updated Knative documentation.
- Researched how to create the Image caption generator and PDF Generation benchmarks.
- Wrote raw code for the benchmarks.

**Next week plan:**

- Image Captioning & PDF generator Benchmark addition to the project.

### **Aug 16th**

- Researched about the different open source libraries/project we can use for new benchmarks.
- This week was more of a researching and reading blogs on different implementations of PDF gen, Image caption generator benchmarks.

**Next week's plan:**

- Raising PR for benchmarks and working on reviews.

### **Aug 25th**

- Raised PDF generation benchmark PR and Image Captioning generator benchmark PR along with their corresponding data.