BenchPilot

Automated Benchmarking for Edge Micro-DCs

Joanna Georgiou - jgeorg02@ucy.ac.cy
Moysis Symeonides - msymeo03@ucy.ac.cy
Michalis Kasioulis - mkasio01@ucy.ac.cy
Demetris Trihinas - trihinas@ucy.ac.cy
George Pallis - gpallis@ucy.ac.cy
Marios D. Dikaiakos - mdd@ucy.ac.cy

MOTIVATION

The number of Internet of Things devices is growing rapidly along with their generated data that needs to be processed and analysed.

Edge Micro-DCs have recently started to be deployed to reduce the network delays occurring by the transfer of IoT data to Cloud, hence ensuring:

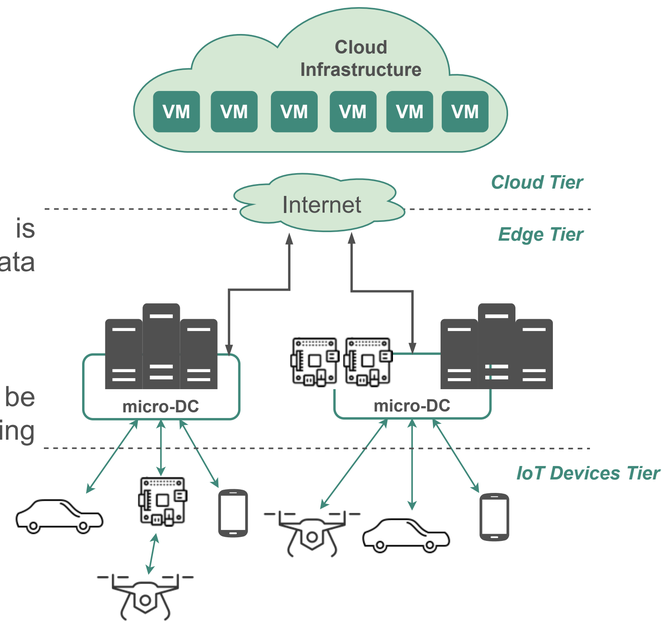

Smaller Response Time | **Less Network Pressure** | **Efficient Processing**

Even if the deployment of Streaming Distributed Processing Engines (SDPEs) on edge seems to be a good idea….
How might we investigate the performance of SDPEs on an edge Micro-DC?

BENCHPILOT

Open-source, modular & **highly customizable** benchmarking framework for edge micro-DCs.

Automates the benchmarking process on SDPEs, enabling users to **focus on performance analysis** instead of dealing with the complex and time-consuming setup.

Instantiates the underlying cluster, performs **repeatable experimentation**, and provides a **unified monitoring stack**.

ARCHITECTURE

BenchPilot's high-level declarative model for describing experiment testbeds

Choosing workload, setting repetition, and duration parameters, along with specific workload preferences

```
experiments:
 - workload:
    name: "marketing-campaign"
    record_name: "storm"
    repetition: 1
    duration: "10m"
    parameters:
      campaigns: 100
      tuples_per_second: 1000
      capacity_per_window: 10
      time_divisor: 10000
    cluster:
      manager: "xx.xx.xx.xx"
      nodes: [ "rp1", "rp2", "nc1" ]
    engine:
      name: "storm"
    parameters:
      partitions: 5
      ackers: 2
 - workload:
    name: "marketing-campaign"
    record_name: "flink"
    ...
```

Analytic Performance Insights

Accessing detailed reports and charts through Jupyter

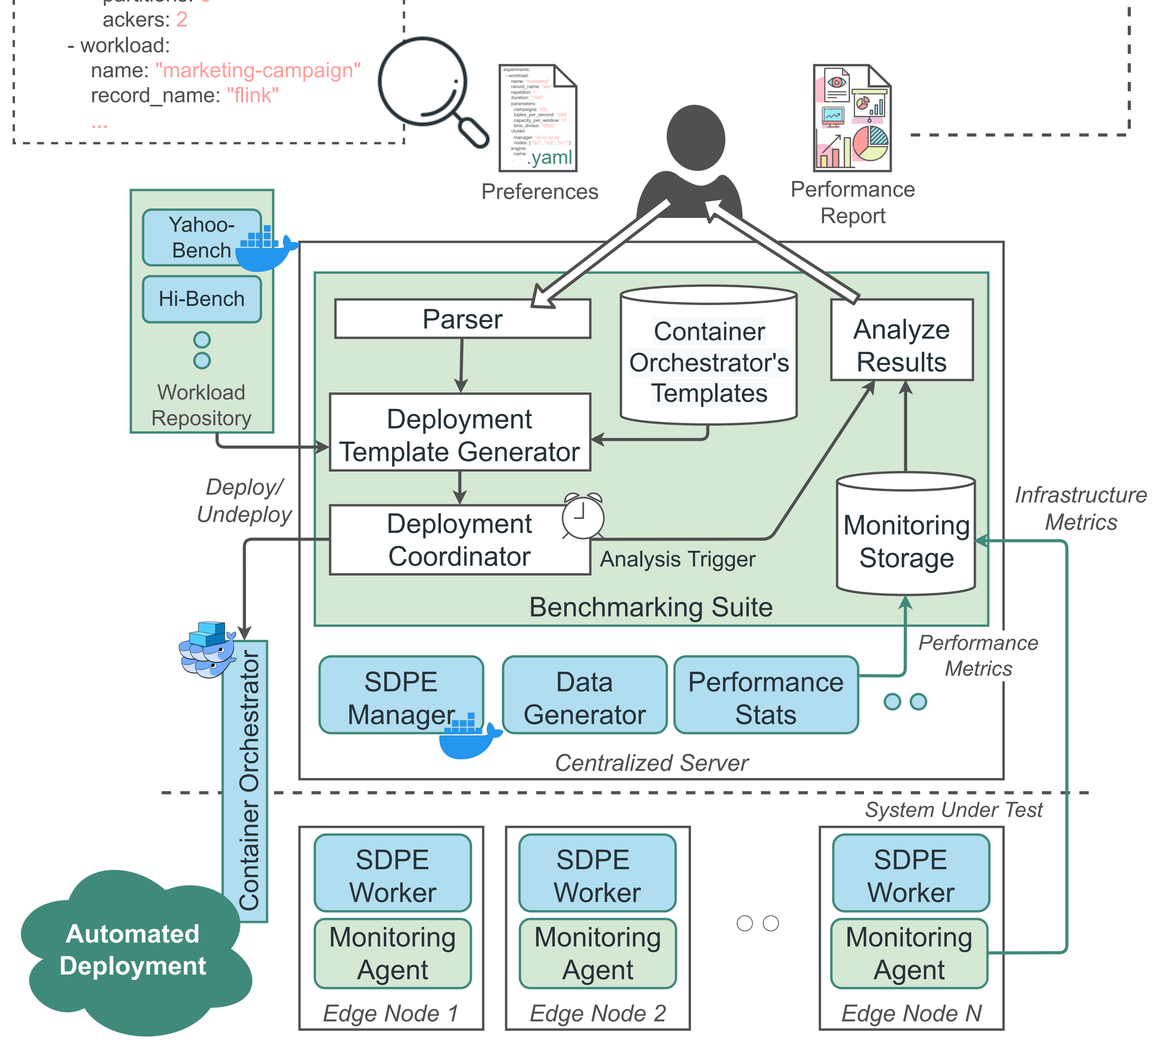

EXPERIMENTS

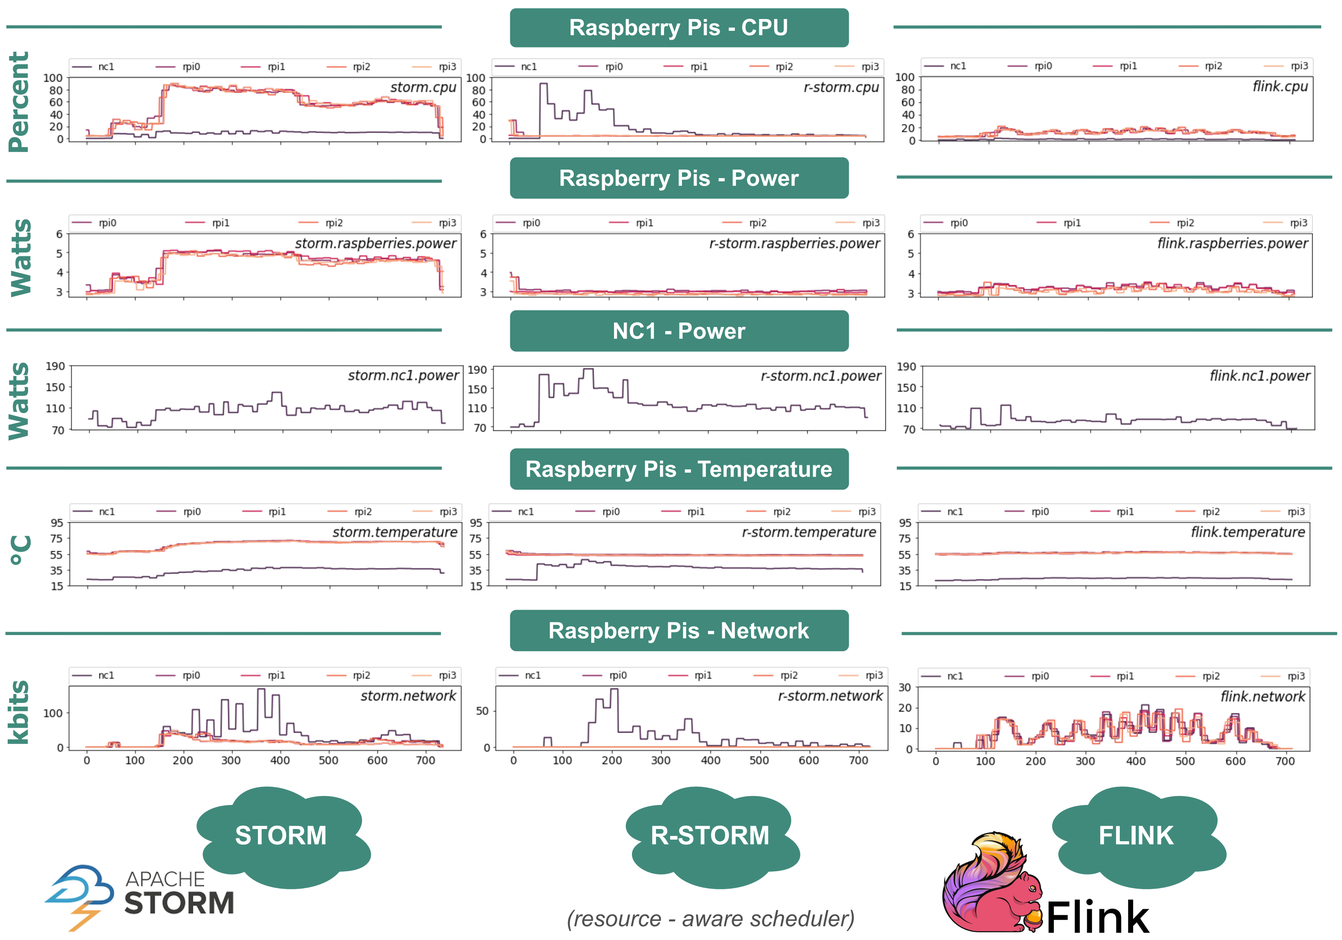

STORM (APACHE STORM)

R-STORM *(resource - aware scheduler)*

FLINK (Flink)

** Data were collected through BenchPilot's monitoring stack*

Experiment Run: 10min
Load Injection: 150sec
Cooldown: 50sec

Lightweight
50k tuples per second data generation rate and 50k campaigns

Heavyweight
500k tuples per second and 50k campaigns

Workload Latency

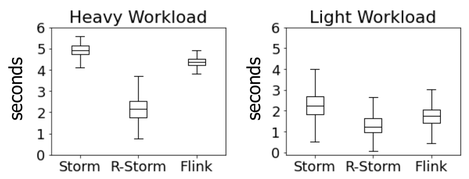

The **testbed** consists of the following SDPE worker nodes:

NC1: Dell PowerEdge R610 Server, with 12cores@2.4GHz & 12GB RAM

RP0-RP3: Raspberry Pi 4 Model B, with quad core ARM Cortex-A72@1.5GHz & 4GB RAM

University of Cyprus
Laboratory for Internet Computing
http://linc.ucy.ac.cy
RAIS

This work is partially supported by the EU Commission through **RAINBOW** 871403 (ICT-15-2019-2020) project, and **RAIS** (Real-time analytics for the Internet of Sports), under grant agreement No 813162

Workload used for experiments:
Yahoo Streaming Benchmarks
(https://github.com/yahoo/streaming-benchmarks)

BenchPilot

This work is based on: "BenchPilot: Repeatable & Reproducible Benchmarking for Edge Micro-DCs", Joanna Georgiou, Moysis Symeonides, Demetris Trihinas, George Pallis, Marios D. Dikaiakos, **"Proceedings of the 27th IEEE Symposium on Computers and Communications"** (**ISCC**), **2022**