Collecting, cataloguing and searching performance information of Cloud resources.

Olaf Elzinga

Why?

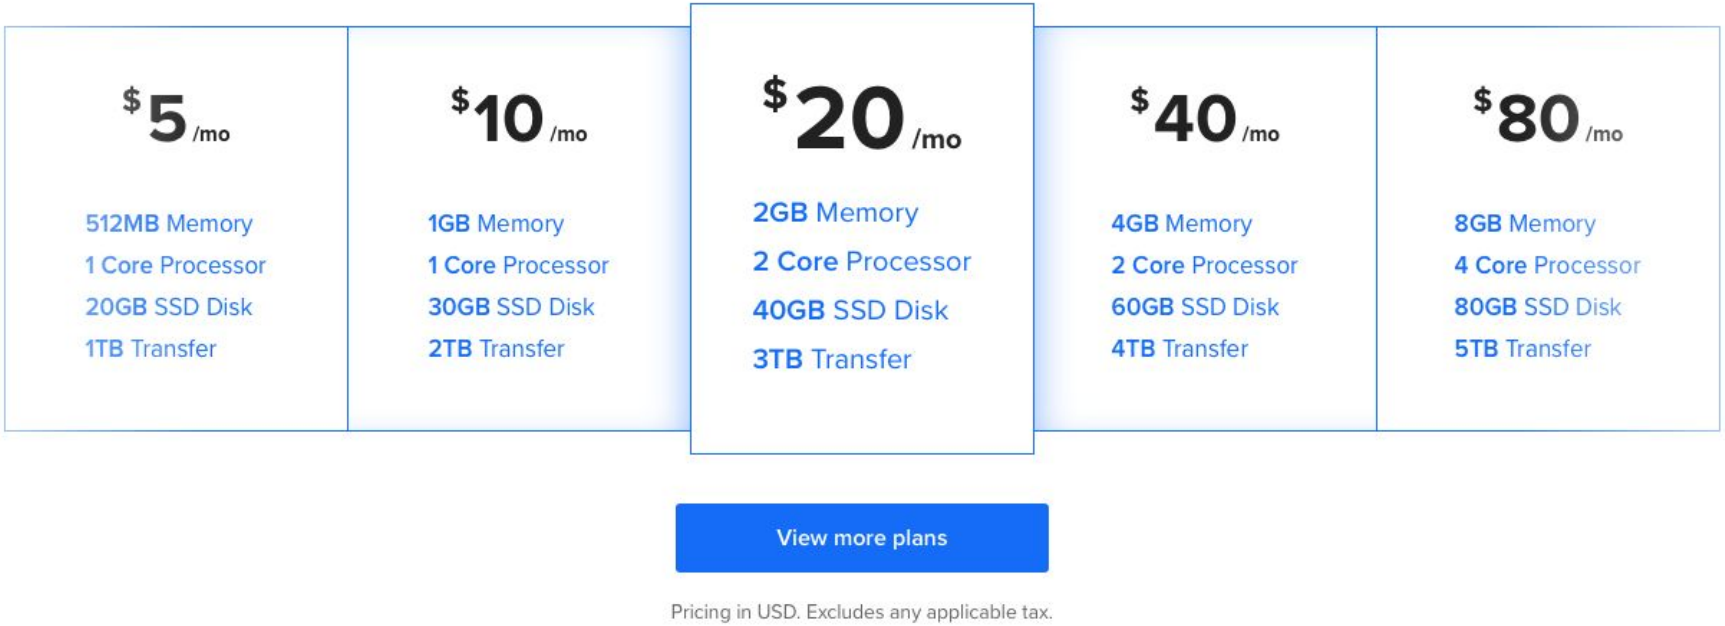

Source: https://www.digitalocean.com/pricing/

Research question

How can an automated cloud benchmark tool test any given application component to maintain a cloud performance catalogue?

State of the art review

Requirements for the automatic cloud benchmark tools:

- Publicly available
- Open-source
- Maintained
- Support for private and public IaaS providers

Related work

	Custom benchmarks	Schedule	Provider support	Catalogue result
Cloud WorkBench [1]	Yes	Yes	Only public	No
CloudBench [2]	No	No	Public and private	No

[1] Joel Scheuner, Philipp Leitner, Jürgen Cito, and Harald Gall. Cloud work bench– infrastructure-as-code based cloud benchmarking. In Cloud Computing Technology and Science (CloudCom), 2014 IEEE 6th International Conference on, pages 246–253. IEEE, 2014.

[2] Marcio Silva, Michael R Hines, Diego Gallo, Qi Liu, Kyung Dong Ryu, and Dilma Da Silva. Cloudbench: experiment automation for cloud environments. In Cloud Engineering (IC2E), 2013 IEEE International Conference on, pages 302–311. IEEE, 2013.

Technical gaps

- Catalogue the collected results
- Ability to add providers
- Possibility to add custom benchmarks

Requirements

Requirements for the users:

- Easy to use
- Fully automatic and possible to scheduling benchmarks
- Custom benchmarks to test different performance attributes
- Catalogue results

Requirements for developers:

- Modular in design

Cloud Performance Collector

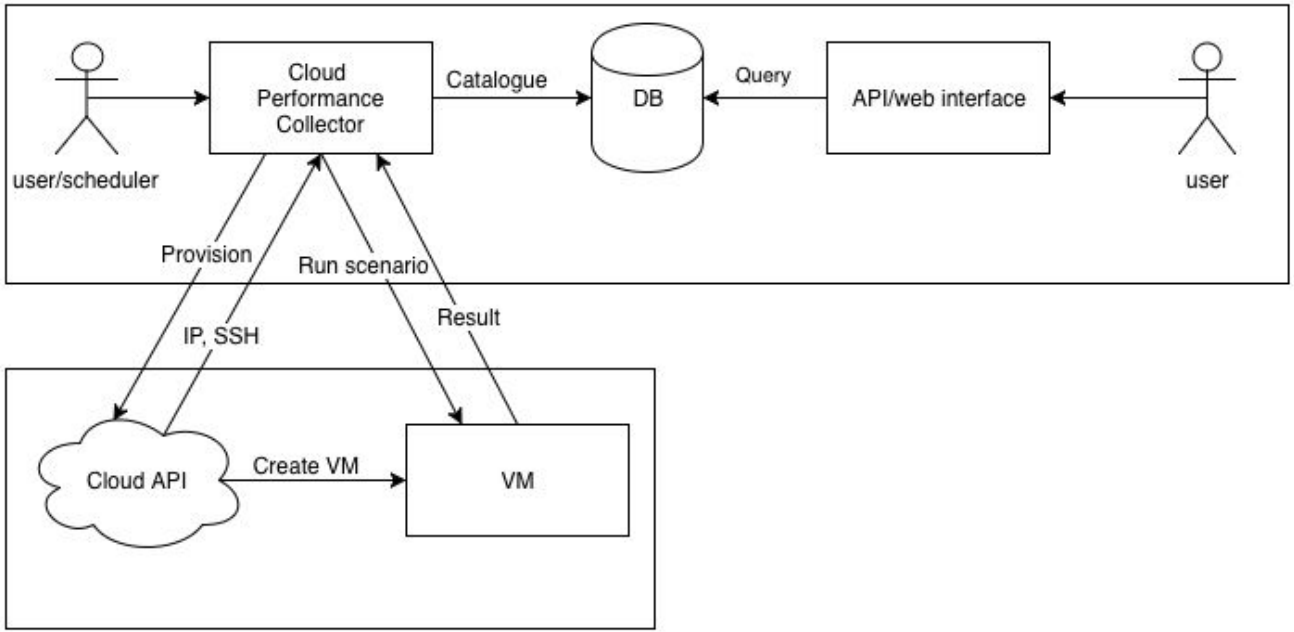

Cloud Performance Collector: modules

- Provider module
  - Provision VM
  - Release VM
- Deploy and run module
  - Installing, configuring and running the benchmarks
- Result module
  - Parse the useful parts of the benchmark output

Cloud Performance Collector: workflow

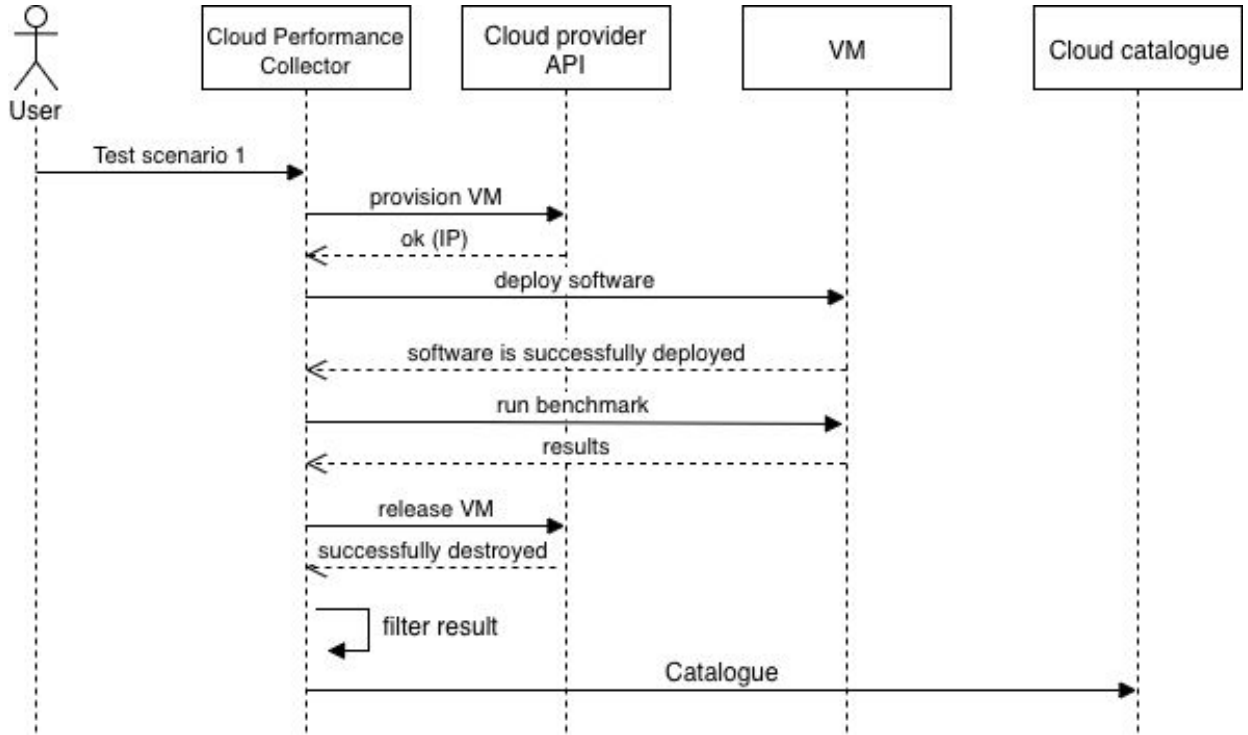

Cloud Performance Collector: prototype

- CLI
- Provider modules written in bash
- Installing, configuring and running the benchmarks via Ansible [1]
- Benchmarks as Dockerfile
- Scheduling via crontab

Execution example:

- bash modules/providers/geni/geni nictaXL

[1] https://www.ansible.com

Experimental setup

ExoGeni:

- University of Amsterdam (UvA)
- National ICT Australia (NICTA)
- Raytheon BBN Technologies (BBN)

Experimental setup: ExoGeni resources

Type	CPU	Memory	SSD
M	1 vCPU	3 GB	25 GB
L	2 vCPU	6 GB	50 GB
XL	3 vCPU	12 GB	100 GB

Experimental setup: questions

- Will VM instances with the same specifications and image from the same provider give similar performance?
- Will the same VM instance with the same workload provide a constant level of performance over time?
- Will a given application component perform the same compared to the synthetic benchmarks?

Experiment 1:

- Measure the difference in performance between different VMs with the same image
- Using a different VM instance every 2 hours
- Measured 24 times (every hour)

Benchmark	Component		Metrics
Sysbench	CPU	Calculate the primeness of 100,000 numbers	Duration (sec)
Stream	Memory	Triad A[i] = B[i] + scalar * C[i]	Throughput MB/s
iozone	Disk	Read and write 64Kb using a file of 2GB	Throughput MB/s

Experiment 1: results CPU (sysbench)

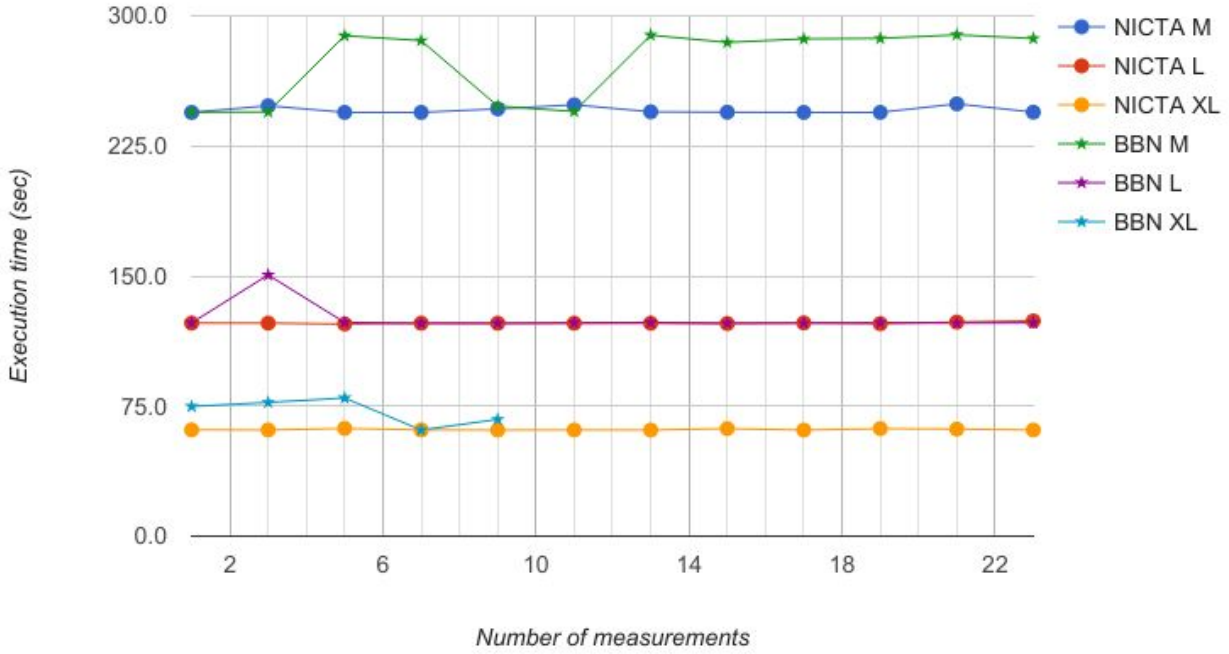

Experiment 1: results memory (STREAM)

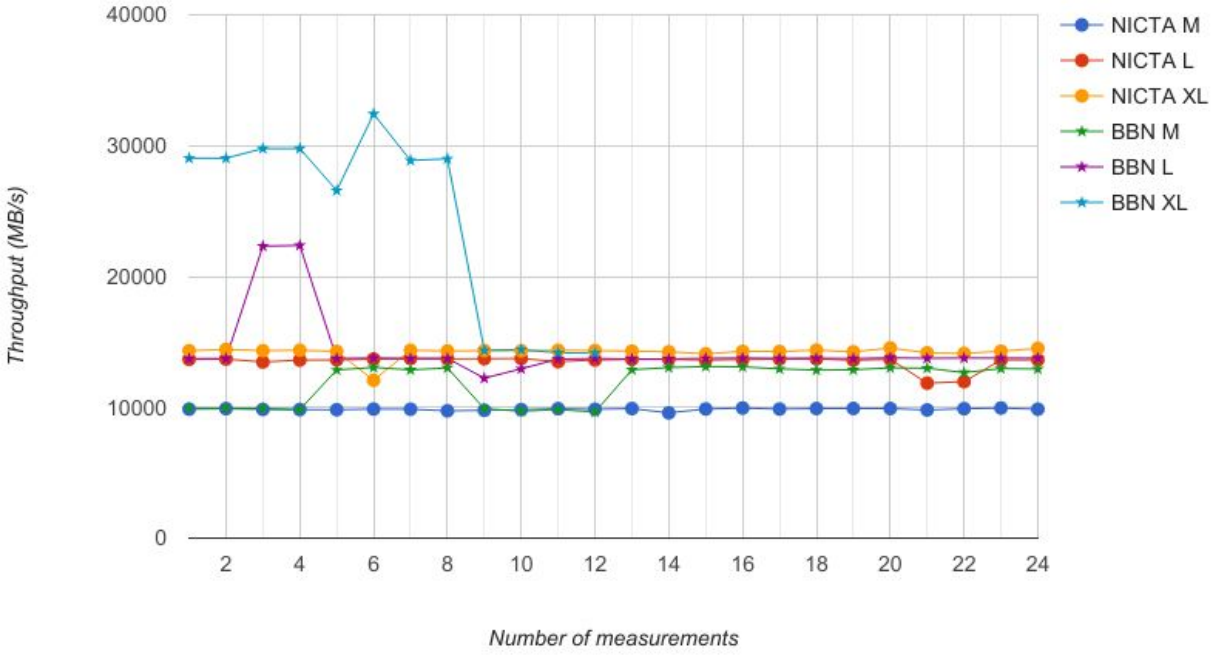

Experiment 1: results disk I/O read and write (iozone)

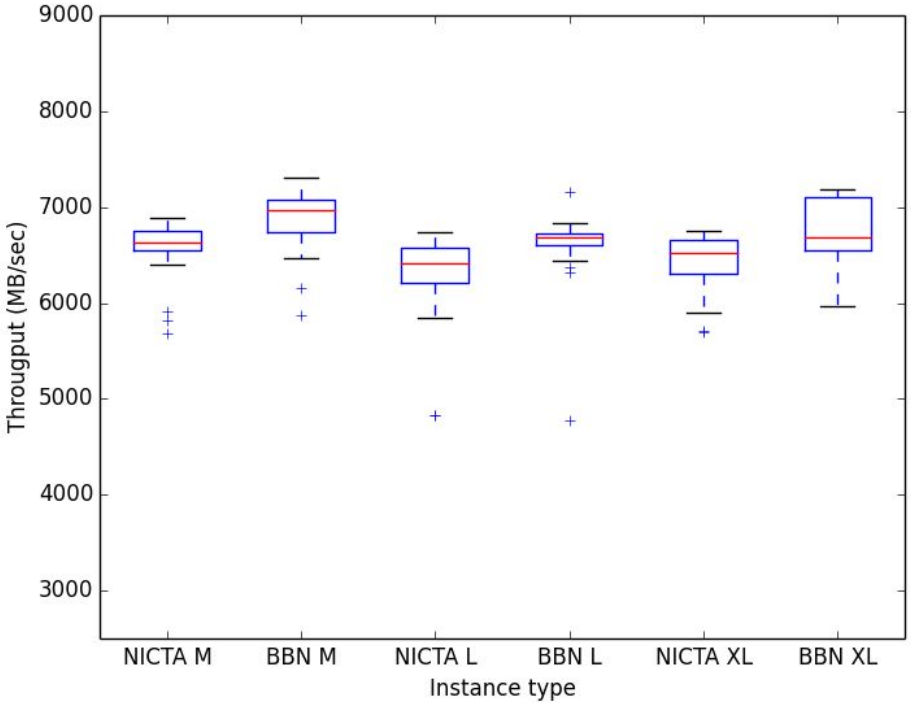

Read

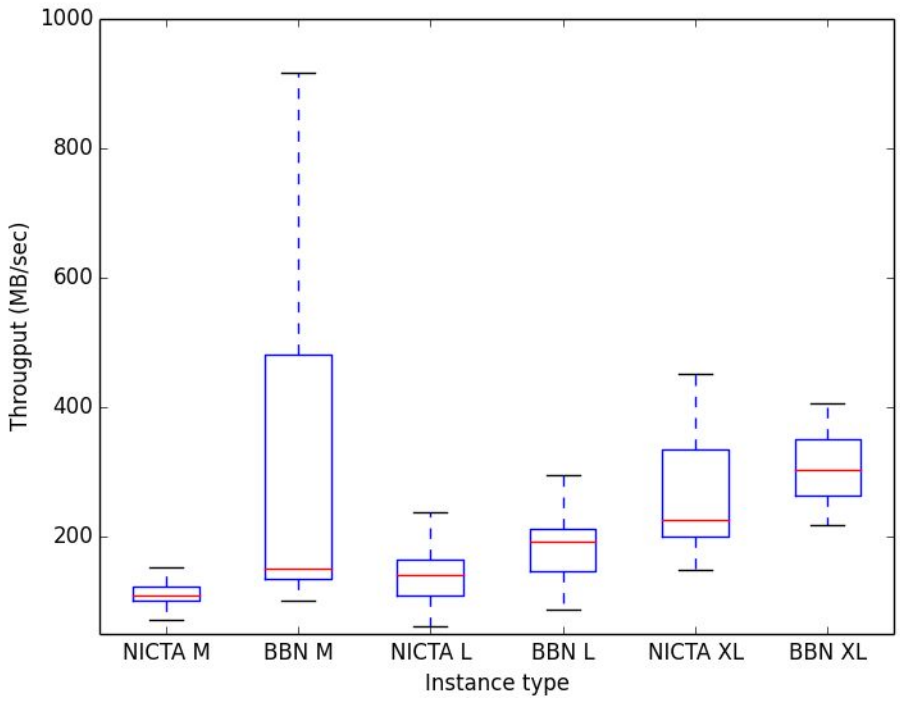

Write

Experiment 2:

- Using the same VM instance for every benchmark
- Use the same benchmark tools as experiment 1
- Measured 24 times (every hour)

Experiment 2: results CPU (sysbench) & memory (STREAM)

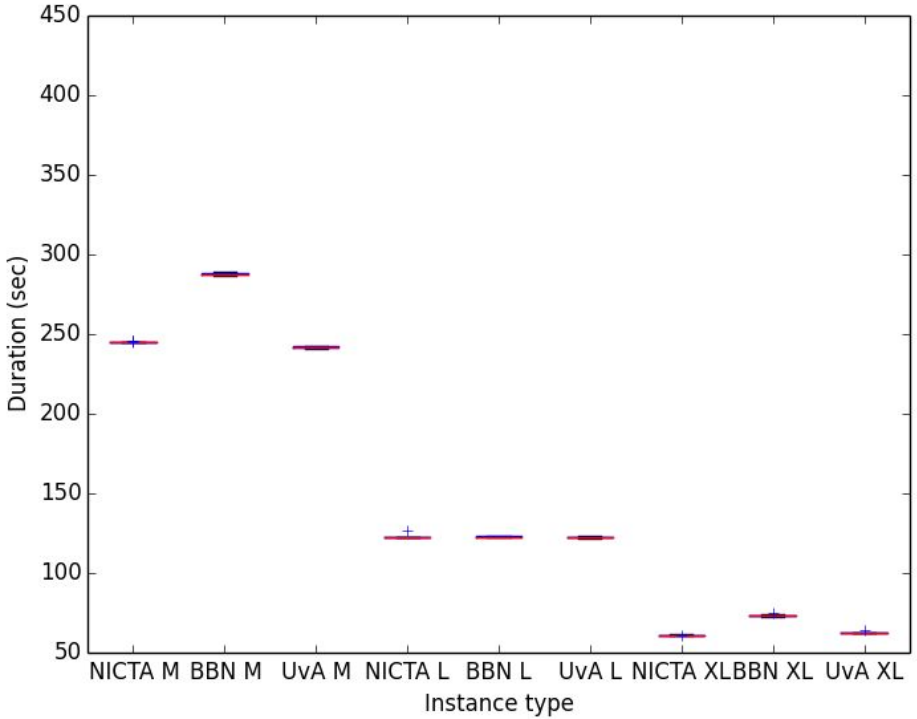

CPU

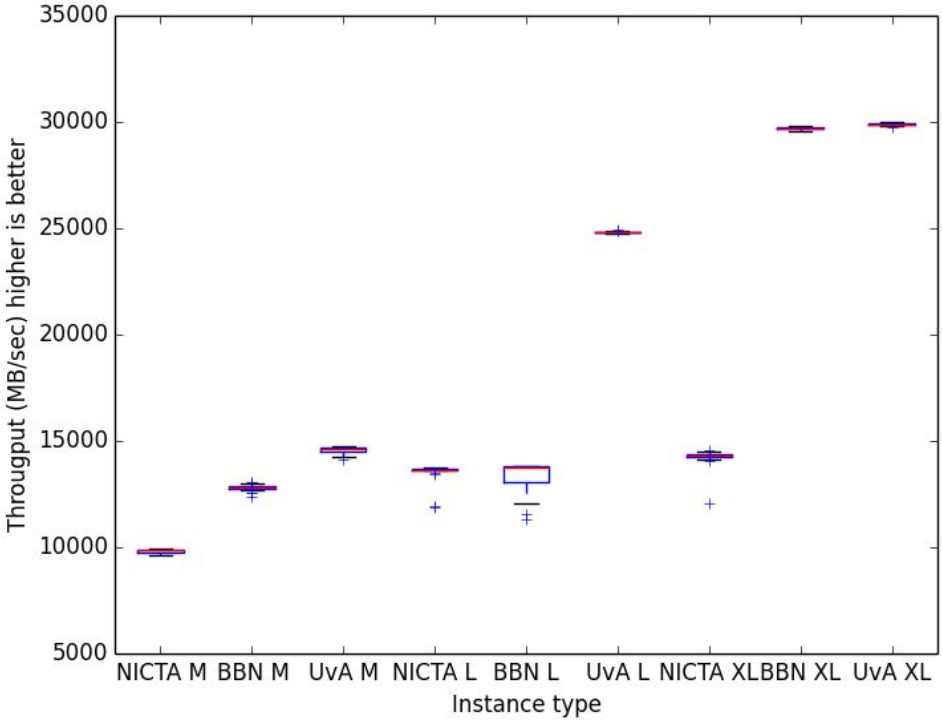

Memory

Experiment 2: results disk I/O read and write (iozone)

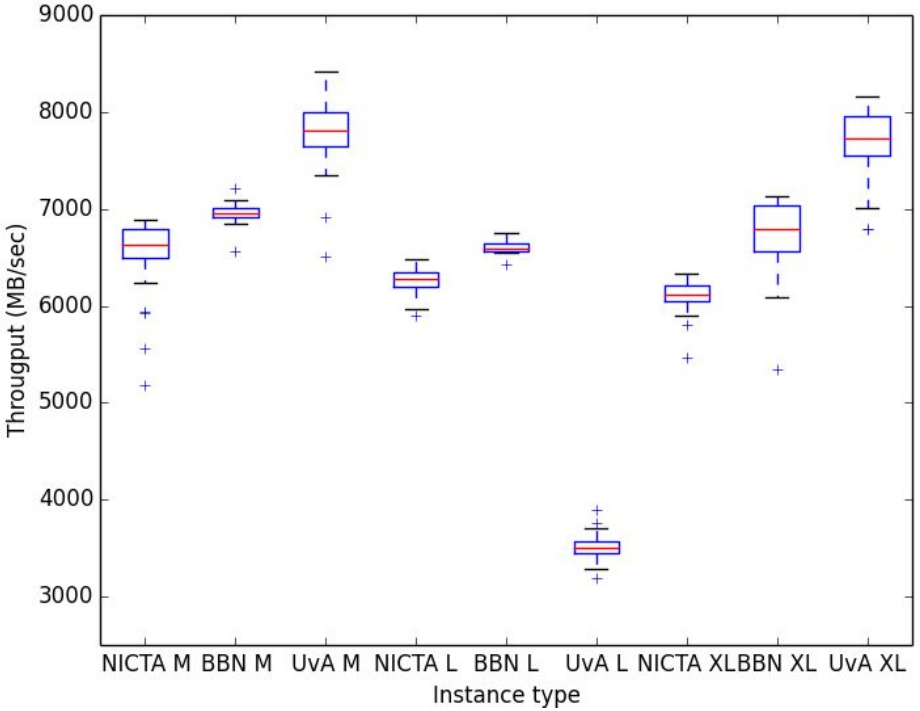

Read

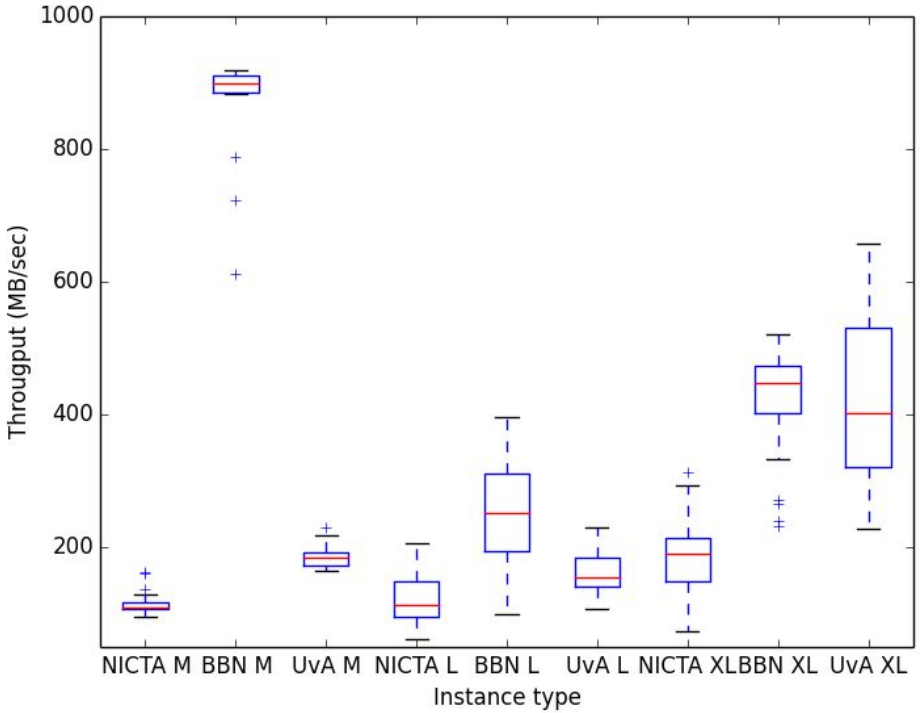

Write

Experiment 3:

- Using docker container with the application Montage
  - An astronomical image mosaic engine
- Measuring how long it takes to create the astronomical image
- Measured 24 times (every hour)

Experiment 3: results Montage

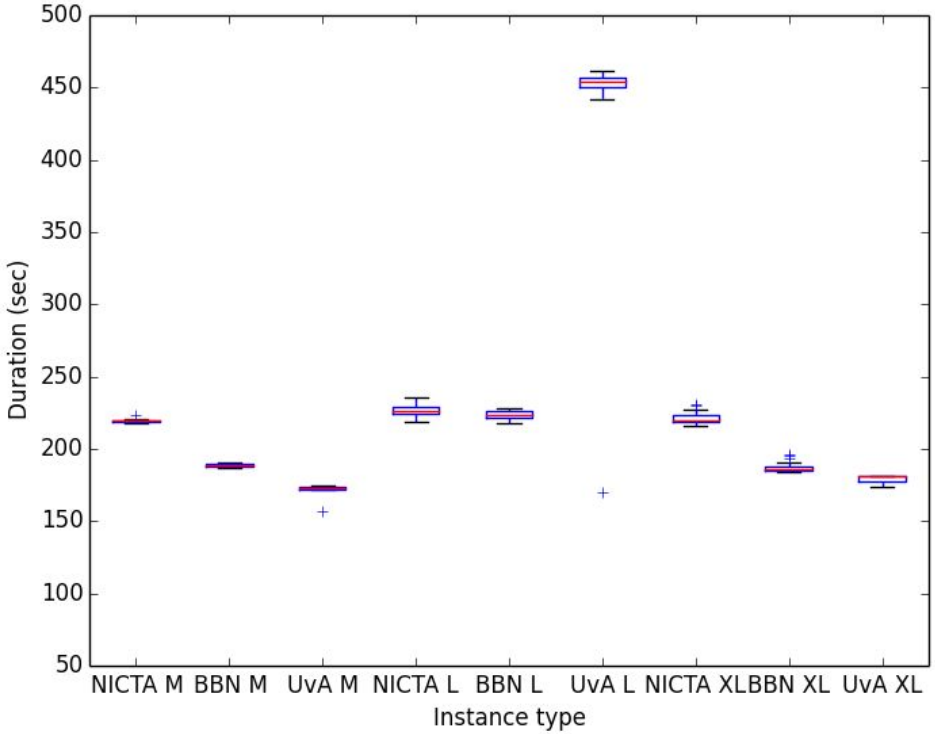

Conclusion

- Performance can vary between different VMs within an ExoGeni rack
- The same VM instance perform similar over time
- Largest instance is not always the right choice
- Problems provisioning VMs and suddenly were unreachable (UvA rack)

Future work

- Test it with a larger amount of applications
- Test the network performance of resources
- Design the cloud catalogue

Questions?