Performance Variation of the Public Cloud

Guilherme G. Haetinger, Eduardo Roloff and Philippe Navaux

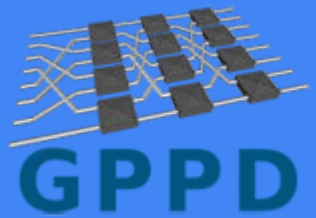

Introduction -

- Motivation?
- Research conduction?
- Tools used?
- Expected results?

Microsoft Azure -

- Azure’s Virtual Machinery:
  - A wide range of options
  - From 13 - 2000 dollars/month
  - 4 - 32 processing cores
  - Each machine category with its own processing purpose, from simple processing requests to High Performance
- Chosen machine: A8
  - HPC focused machine - Intel Xeon E5 processor
  - 8 cores
  - Perfect for cluster workloads

Microsoft Azure -

- Azure’s locations:
  - 26 datacenters all over the globe
  - USA has the biggest amount of data centers per country
- Chosen regions:
  - North USA
  - South USA
  - East USA
  - West USA
  - Purpose behind choosing 4 locations in the same Country?

Leveled Core Individual Efficiency Benchmark (LCIE) - Temporary name!

- Custom made benchmark in development to be Cloud fit
- Characteristics:
  - Takes in consideration the machine's number of cores
  - Flexible - can measure the efficiency on the same core 'n' times
  - Segmented on Load Levels:
    - User enumerated sections
    - Represented in percentages
    - Leveled by a user-given proportion
  - Each core is set to generate a number of random integers equal to the proportion times the Load Level it has been fit in.

Leveled Core Individual Efficiency Benchmark (LCIE) - Temporary name!

- The LCIE benchmark was applied to the four regions chosen with these following conditions:
  - 4 Load Levels and a proportion of 100000
  - Executed 10 times in each data center
  - The result in each region is set as a tuple of 4 values:
    - The mean of the results in the Load Level of 25%
    - The mean of the results in the Load Level of 50%
    - The mean of the results in the Load Level of 75%
    - The mean of the results in the Load Level of 100%
  - Considering these results, a graphic was plotted as the final result, as well as the variation between the results in each Load Level was calculated having the West USA data center efficiency as parameter.

Results & Interpretation -

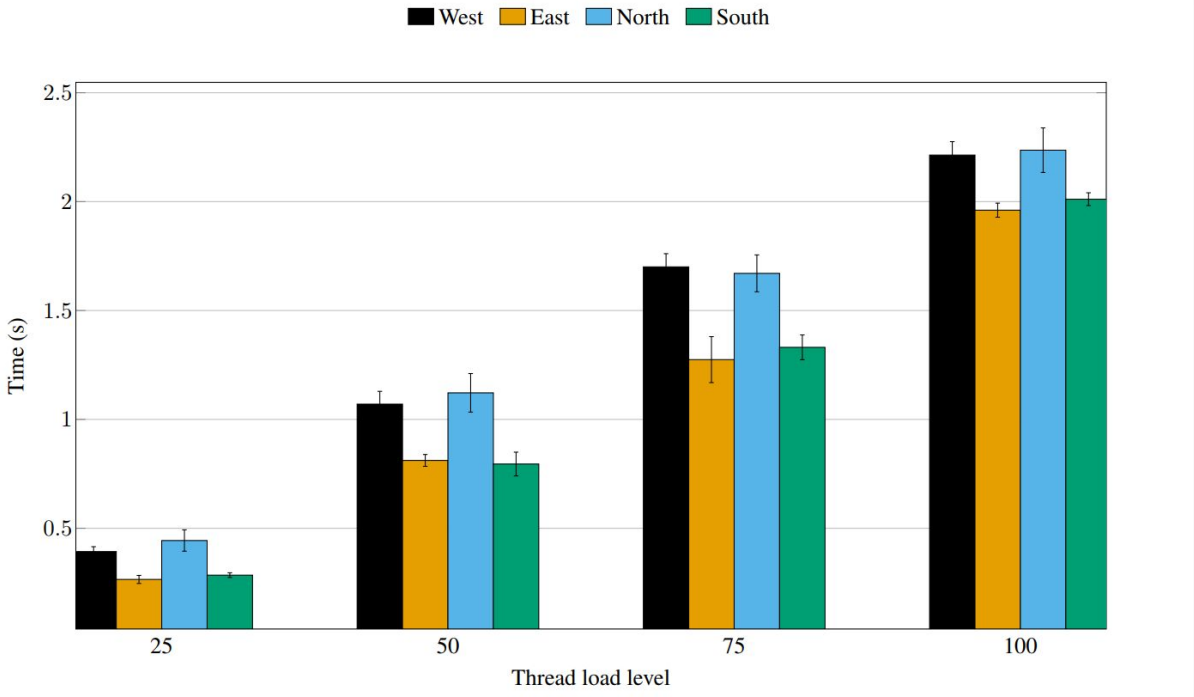

Load Level	East	North	South
25	48.25%	-11.43%	38.03%
50	31.85%	-4.60%	34.62%
75	33.43%	1.79%	27.76%
100	12.91%	-1.00%	10.07%

Results & Interpretation -

- As the results show, there was a great difference between the data centers results. However, the variation between each location is almost constant.
- As we can see, the efficiency ranking between the four data centers is:
    1. East USA
    2. South USA
    3. West USA
    4. North USA

Results & Interpretation -

- The variation seen has many possible explanations:
  - The results were a product of the particular Virtual Machines we allocated, since we didn't reallocate them.
  - The higher number of allocations in areas with tremendous money investments on technology(most seen in the West, Center-Northeast).
  - We need further investigation...

Conclusion -

- These results were unexpected, even though research related to computer processing and efficiency is never previsible.
- It is the company's duty to present the same effectiveness to all locations it disposes their services. The data gathered in this project shows a problem that many people are exposed to even if they don't notice it.
- Both the differences of performance among these four locations and the differences observed in performance gains and losses for the different load levels are promising research topics that need further investigation.

Performance Variation of the Public Cloud

Guilherme G. Haetinger, Eduardo Roloff and Philippe Navaux