Guest Lecture: Cloud Computing – Class Activities Building Modern Web Applications - Winter 2021-1

Julien Gascon-Samson, Assistant Professor

ÉTS Montréal / University of Quebec Département of Software and IT Engineering Montreal, Canada



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ → ヨ → の Q @

Tuesday, Nov 30, 2021





• Why is resource polling an important characteristic of cloud environments?

A- Cloud Characteristics



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ → ヨ → の Q @

- Why is resource polling an important characteristic of cloud environments?
- We have presented some advantages of the "measured service" characteristic in class, both from the perspective of cloud users and cloud proviers. Describe one inconvenience of that characteristic for the provider.

A- Cloud Characteristics



- Why is resource polling an important characteristic of cloud environments?
- We have presented some advantages of the "measured service" characteristic in class, both from the perspective of cloud users and cloud proviers. Describe one inconvenience of that characteristic for the provider.
- The Heroku platform simultaneously hosts many web apps from many customers on the same shared infrastructure. Which cloud characteristic corresponds to that statement?





To which service model each of the following correspond?

Editing a document online on Google Docs





・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・ ・ つ へ ()

- Editing a document online on Google Docs
- ② Testing a new Linux Kernel on an Amazon VM



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ → ヨ → の Q @

- Editing a document online on Google Docs
- ② Testing a new Linux Kernel on an Amazon VM
- Our Purchasing gifts on Amazon



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ ● ● ● ● ●

- Editing a document online on Google Docs
- ② Testing a new Linux Kernel on an Amazon VM
- Our Purchasing gifts on Amazon
- Accessing a MySQL database service



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ ● ● ● ● ●

- Editing a document online on Google Docs
- Presting a new Linux Kernel on an Amazon VM
- Our Purchasing gifts on Amazon
- Accessing a MySQL database service
- Opploying a Python application



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ ● ● ● ● ●

- Editing a document online on Google Docs
- Presting a new Linux Kernel on an Amazon VM
- Ourchasing gifts on Amazon
- Accessing a MySQL database service
- Opploying a Python application
- Provisioning a virtual machine





・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・ ・ つ へ ()

To which service model each of the following correspond?

Developing and deploying "lambda" functions that will be executed in reaction to specific events.



・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・ ・ つ へ ()

- Developing and deploying "lambda" functions that will be executed in reaction to specific events.
- Opploying a web app on Heroku.



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ ● ● ● ● ●

- Developing and deploying "lambda" functions that will be executed in reaction to specific events.
- Opploying a web app on Heroku.
- **③** Requesting a set of VMs with high performance GPUs.



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ ● ● ● ● ●

- Developing and deploying "lambda" functions that will be executed in reaction to specific events.
- Opploying a web app on Heroku.
- **③** Requesting a set of VMs with high performance GPUs.
- Ownloading a game on steam



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ ● ● ● ● ●

- Developing and deploying "lambda" functions that will be executed in reaction to specific events.
- Opploying a web app on Heroku.
- **③** Requesting a set of VMs with high performance GPUs.
- Ownloading a game on steam
- **o** Watching the CPEN400A lectures on Youtube.



▲ロ ▶ ▲周 ▶ ▲ ヨ ▶ ▲ ヨ ▶ ● ● ● ● ●

- Developing and deploying "lambda" functions that will be executed in reaction to specific events.
- Opploying a web app on Heroku.
- **③** Requesting a set of VMs with high performance GPUs.
- Ownloading a game on steam
- **o** Watching the CPEN400A lectures on Youtube.
- O Configuring a SDN for a group of VMs.



・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・ ・ つ へ ()

Choose the most accurate deployment model for each of the following (public cloud, private cloud, community cloud, hybrid cloud):

Oue to strict privacy laws, UBC offers a storage service operated by the university and hosted on the university premises.



・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・ ・ つ へ ()

Choose the most accurate deployment model for each of the following (public cloud, private cloud, community cloud, hybrid cloud):

- Oue to strict privacy laws, UBC offers a storage service operated by the university and hosted on the university premises.
- Intel-ligent operates a cloud service for its own internal needs whenever the demands goes above the available capacity, the extra load is sent to the MS Azure cloud.



・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・ ・ つ へ ()

Choose the most accurate deployment model for each of the following (public cloud, private cloud, community cloud, hybrid cloud):

- Oue to strict privacy laws, UBC offers a storage service operated by the university and hosted on the university premises.
- Intel-ligent operates a cloud service for its own internal needs whenever the demands goes above the available capacity, the extra load is sent to the MS Azure cloud.
- You Android phone automatically stores the photos that you have taken onto the Google Photos service.



Choose the most accurate deployment model for each of the following (public cloud, private cloud, community cloud, hybrid cloud):

- Oue to strict privacy laws, UBC offers a storage service operated by the university and hosted on the university premises.
- Intel-ligent operates a cloud service for its own internal needs whenever the demands goes above the available capacity, the extra load is sent to the MS Azure cloud.
- You Android phone automatically stores the photos that you have taken onto the Google Photos service.
- The healthcare network provides a cloud service that allows for storing patient information from multiple establishments (hospitals, clinics, pharmacies, etc.) scattered across the province. The information is stored on a layer of servers deployed/spread over the various establishments. The service respects various compliance policies pertaining to storing critical patient information.

D- Load Balancing and Optimization



<u>Context:</u> Uber Best Cloud (UBC) has 25 IoT devices that are deployed around the world. Each device has to <u>send a message to each other device</u> once per second (i.e., each device will need to send 24 messages per second, so a total of 600 messages per second will be sent). As a peer-to-peer communication is not possible due to firewall issues, each message must be relayed through a datacenter in the Amazon EC2 cloud.

The file *cloud-optimization.csv* contains arbitrary latency values (round-trip time / ping) between each of the 25 devices, and each of the available Amazon cloud regions¹. You are hired as a consultant, and you are tasked with determining which Amazon region(s) should be used to relay the messages exchanged by the IoT devices, in order to minimize the average latency.

The file *cloud-optimization.js* contains a code skeleton that you must complete.

We assume that the service will be deployed on <u>only one cloud region</u>. You must implement the findBestRegion function, which will determine the most optimal AWS region that should be used for all message exchanges ($Node_A \rightarrow Region_R \rightarrow Node_B$) - that is, which region R would minimize the average latency. The function should print the average latency for all regions, and return the name of the region that minimizes the latency.

We assume that the service is replicated on all cloud regions, and that every single message is routed through the region *i* that minimizes its latency (i.e., Node_A → Region_i → Node_B, which "best" region *i* would minimize the latency for that message). You are asked to implement a second function findLowestLatencyReplicated that will return the lowest average latency that can be obtained by fully replicating the service. You should also print the percentage of messages that will be processed at each region.

GitHub repository for the files mentioned above:

https://github.com/juliengs/gti780-exemples-A2021/tree/master/cours1-ds-cloud

¹All values are randomly generated.

うせん 御 ふぼうえばく (型・人口・