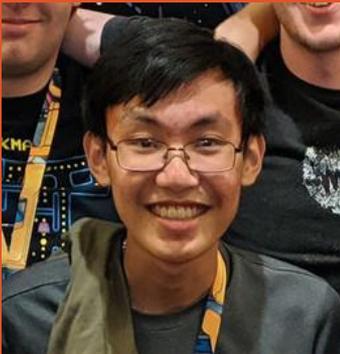


FAISAL KHAN



ANDREW QUACH



CLOUD INFRASTRUCTURE BILLING

Automating virtualization resource billing



PROBLEM

Parts of campus are allowed to self-provision these virtual machines without any middle man. Consumers can reserve any amount computing resources they deem necessary.

This pushes the burden of tracking usage to IT infrastructure. Currently, the billing calculation process is a mix of automated and manual procedures.

To complicate things further, not all virtual machines are treated equal. Different parts of campus are charged different rates. Some are not even charged at all.

DEVELOPERS

- Andrew Quach: quachand@oregonstate.edu
- Faisal Khan: khanfa@oregonstate.edu

REQUIREMENTS

Continuously scrape data from vSphere for VMs and scrape data from storage and store in a database (hourly).

Generate a spreadsheet based on scraped data in FUPLOAD format and send to accounting.

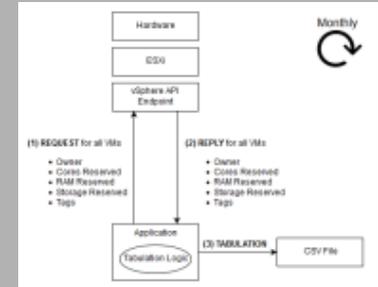
Automatically push tags to VMs to add in spreadsheet calculations based on user-changeable configuration file.

Created with pyVmomi and python3 technology stack as base.

CLIENTS

- Majorie McLagan: Majorie.mclagan@oregonstate.edu
- Stacy Brock: brocks@oregonstate.edu

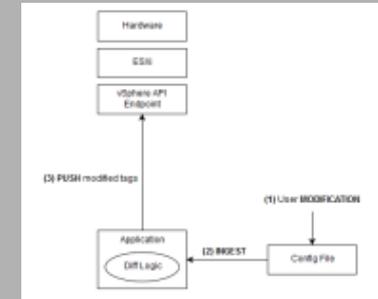
BILLING DESIGN



From the vSphere API, we will extract all necessary information to calculate billing amount. This API call mechanism (1) and (2) will happen every hour and stored in an sqlite3 database.

After pulling this information, the bill is tabulated for all VM owners and outputted to a CSV file. This (3) happens every month. The output file is automatically sent to accounting.

TAGGING DESIGN



The billing application requires some user inputted data. This input will be in a YAML configuration file. The tagging application continuously monitors for VMs without tags and updates the VMs with metadata about the index used and whether it is managed.