Pay as You Go in the Cloud: One Watt at a Time

Kayo Teramoto *
Yale University

H. Howie Huang (Advisor)
George Washington University

1 Introduction

The introduction of virtual machines and cloud computing in the market has raised the need for a fair pricing that appropriately reflects the nature and costs of the service: providing computing resources to consumers through a data center. Currently, cloud computer providers such as Amazon EC2 charge consumers by the hour with a rate that varies by the allocation size of the VM (e.g., small and large), the data center location (e.g., US east and Asia), the type of the purchase (e.g., on-demand and reserved), and the software (e.g., Linux and Windows) [1]. This pricing scheme accounts for the variety in VM specifications, but overlooks the disparities in resource usage and the related infrastructural costs when running different workloads on a VM.

A significant percentage of data center costs comes from power and power distribution and cooling [2]. Power fluctuates with changes in CPU, disk, and memory usage. Hence, a VM running a resource-intensive application expends more power and is more costly than a similar VM using less hardware and consuming less energy. However, since power is not factored in the per-hour rate, the latter VM must share the cost of providing for the former.

We thus propose a pay-as-you-go per-watt-hour pricing model. The main advantages are:

- A fairer division of the cost based on VM power consumption is achieved; pricing becomes a more accurate cost estimate of the received service.
- Cloud computing providers' profit can become more stable. With a per-hour-rate, Amazon's profit may decrease when resource usage is maximized (due to less flexibility in VM consolidation). With our model, Amazon can ensure a guaranteed profit from each user and more easily adjust for changes in resource cost. Depending on the cost of VM provision, a heavy user may be charged a premium for above-average power consumption.
- Consumers pay only for their accountable power consumption; no more paying for unused minutes (e.g., Amazon EC2's full-hour billing).

*Kayo Teramoto worked as a Summer 2012 undergraduate research assistant at GWU.

2 Power Based Pricing Model

Since a complicated and esoteric model risks losing consumers, we keep our pricing simple:

$$Price_{VM,i} = r \cdot (AccPower_{VM,i}) \tag{1}$$

i is the VM ID number, r is the dollars-per-watt-hour rate, and $AccPower_{VM,i}$ is the VM's accountable power consumption.

Note: Although power is not the only cost of a data center, it is reasonable to say that more of the other costs are due to or correlated with the demands of a highly active VM than a less active VM.

2.1 Power Model

A physical power meter cannot be used on a VM, but previous works [3] has shown that a reasonably accurate aggregate power formula can be constructed using readily available data. Here we propose a new model that excludes Dom0 (the host domain that manages the system) and partitions all of the power among the guest VMs in proportion to their CPU usage (*CPU*) and number of read and write requests (*Disk*).

$$\begin{aligned} TotalPower = {} & a\sum_{i=1}^{n}(CPU_{VM,i}) + c_{cpu} \\ & + b\sum_{i=1}^{n}(Disk_{VM,i}) + c_{disk} + Idle \end{aligned} \tag{2}$$

a and b are CPU and disk coefficients, c_{cpu} and c_{disk} are constants, and $Idle$ is the idle power of the physical machine (while logging the data).

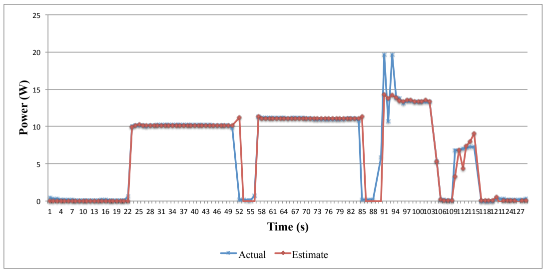

Figure 1: VM Power Consumption: Actual vs. Estimate

The accountable power of VM $i = 1 \ldots n$ is then calculated in proportion to the other VMs running on the machine:

$$AccPower_{VM,i} = a \cdot CPU_{VM,i} + b \cdot Disk_{VM,i} + \frac{Idle}{T}$$
$$+ \left(\frac{CPU_{VM,i}}{\sum_{i=1}^{n} (CPU_{VM,i})} \right) c_{cpu} + \left(\frac{Disk_{VM,i}}{\sum_{i=1}^{n} Disk_{VM,i}} \right) c_{disk} \quad (3)$$

T is the number of VMs of a specific instance/size that can be run on one physical machine.

Since many of the unobserved resource states are correlated to CPU and disk I/O usage, the model recalculates the coefficients and constants when errors in estimated total power increase (generally from workload changes).

3 Evaluation

Accountable Power Model: Three configurations were tested: a Dell machine with 2.93GHz Intel Core2Duo processor, and a 3.5" HDD or an SSD, and a low-power machine with 1.6GHz Intel ATOM processor, and a 2.5" HDD. We installed Xen 4.1.0 hypervisor on all the machines. Per-VM CPU and disk I/O data were collected at every second using a perl script. A Wattsup? power meter measured the total physical power consumption. We ran the CPU, memory, and fileio benchmarks from Sysbench to simulate workloads on the VMs.

Our aggregate power consumption estimates on all three configurations yielded average errors of 0.5-1%. Although a ground truth is lacking for our accountable power model, the estimated VM accountable power held when compared with the difference in power between one VM running and the idle power of the machine (see Figure 1). We thus reason that our model is fair.

Comparison of Pricing Models: To evaluate the impact of the new proposed pricing model, we compare four consumers running identical VMs in the same data center for 200 hours with the following usages:

Consumer	CPU %	IO Req. No.(%)
User1	5	417 (5)
User2	30	2505 (30)
User3	55	4594 (55)
User4	80	6682 (80)

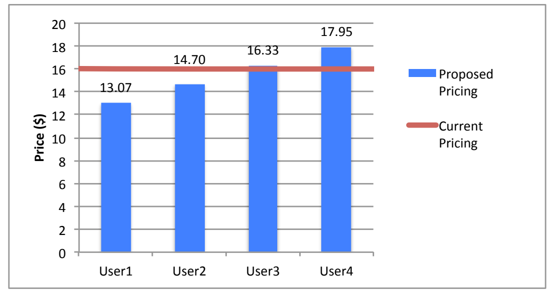

Figure 2: Current Pricing vs. Proposed Pricing

As shown in Figure 2, with the per-hour model, each of the users would pay the same amount at the end of the month ($16 with the rate of $0.080/hour for a standard on-demand small instance). In contrast, using the SSD models, User4 is consuming 37.3% more power than User1 and shall consequently be charged 37.3% more ($4.87 difference when using a $r = \$0.00142/watt \cdot hour$ which was calculated on the assumption that providers expect a $64 total revenue from 4 VMs running an average of 50% CPU and 50% disk I/O).

4 Discussion

In summary, we proposed a per-watt-hour pricing model that makes the division of costs fairer among similar VMs and frees less active VMs from having to pay for others. Future work will be done to further evaluate our accountable power consumption model and experiment using cloud services. Additional variables, such as network, will also be considered in our future VM power consumption models.

5 Acknowledgments

This work is in part supported by the National Science Foundation grant OCI-0937875.

References

[1] AMAZON EC2. http://aws.amazon.com/ec2/.

[2] BELADY, C. In the data center, power and cooling costs more than the it equipment it supports. *Electronics cooling 13*, 1 (2007), 24.

[3] KANSAL, A., ZHAO, F., LIU, J., KOTHARI, N., AND BHATTACHARYA, A. A. Virtual machine power metering and provisioning. In *Proceedings of the 1st ACM symposium on Cloud computing* (New York, NY, USA, 2010), SoCC '10, ACM, pp. 39–50.