Title	Towards the Automatic Detection of Efficient Computing Assets in a Heterogeneous Cloud Environment
Authors(s)	Iglesias, Jesus Omana, Stokes, Nicola, Ventresque, Anthony, Murphy, Liam, B.E., Thorburn, James
Publication date	2013-06-03
Publication information	Iglesias, Jesus Omana, Nicola Stokes, Anthony Ventresque, Liam Murphy B.E., and James Thorburn. "Towards the Automatic Detection of Efficient Computing Assets in a Heterogeneous Cloud Environment." Institute of Electrical and Electronic Engineers (IEEE), 2013.
Conference details	IEEE 6th International Conference on Cloud Computing, Santa Clara Marriott, CA, USA, 28 June - 3 July, 2013
Publisher	Institute of Electrical and Electronic Engineers (IEEE)
Item record/more information	http://hdl.handle.net/10197/7145
Publisher's statement	© © 2013 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Publisher's version (DOI)	10.1109/CLOUD.2013.136

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Towards the Automatic Detection of Efficient Computing Assets in a Heterogeneous Cloud Environment

Jesus Omana Iglesias, Nicola Stokes, Anthony Ventresque, Liam Murphy
School of Computer Science and Informatics, University College Dublin
Dublin, Ireland

10289844@ucdconnect.ie, {nicola.stokes, anthony.ventresque, liam.murphy}@ucd.ie

James Thorburn
IBM software Group
Toronto, Canada
jthorbur@ca.ibm.com

Abstract—In a heterogeneous cloud environment, the manual grading of computing assets is the first step in the process of configuring IT infrastructures to ensure optimal utilization of resources. Grading the efficiency of computing assets is however, a difficult, subjective and time consuming manual task. Thus, an automatic efficiency grading algorithm is highly desirable. In this paper, we compare the effectiveness of the different criteria used in the manual grading task for automatically determining the efficiency grading of a computing asset. We report results on a dataset of 1,200 assets from two different data centers in IBM Toronto. Our preliminary results show that electrical costs (associated with power and cooling) appear to be even more informative than hardware and age based criteria as a means of determining the efficiency grade of an asset. Our analysis also indicates that the effectiveness of the various efficiency criteria is dependent on the asset demographic of the data centre under consideration.

Keywords—Asset Utilization cost, Asset Efficiency Grading;

I. INTRODUCTION

In recent years, large investments have been made to build data processing centers; purpose-built facilities composed of thousands of servers, providing storage and computing services within and across organizational boundaries. These large data centers seek to achieve economies of scale by consolidating massive computing capacity and providing it to end users via virtualization. However, it has been shown that typically servers seldom operate near their maximum utilization, instead, operating most of the time at between 10 and 50 percent of their maximum utilization levels [1], [2]. Optimizing this type of IT infrastructure is a continuous process, in which organizations (i.e., private/hybrid cloud providers, or organizations with a large IT infrastructure) spend a fixed budget every year purchasing new computing assets (i.e., servers, storage, racks and workstations). In order to lower costs, they must also migrate systems (i.e., an entire asset composed of one OS instance, or a virtual machine) from less efficient assets to more efficient ones, and eventually retiring the oldest, most inefficient assets. Efficiency grading of computing assets is however, a difficult, subjective and time-consuming task performed manually in organisations by Capital Planners. In this paper, we evaluate the effectiveness of efficiency assessment criteria that are

commonly used in the manual assessment of asset efficiency grading. This represents, to the best of our knowledge, the first preliminary study of how to automate asset efficiency grading in a heterogeneous cloud environment.

II. EFFICIENCY GRADING CRITERIA

In manual efficiency grading a variety of *rules of thumb* are used by the assessors. For instance, a Capital Planner will tend to consider an asset inefficient if it has one or more of the following characteristics: 1) RAM is less than 8GB; 2) CPU processor with less than or equal to two cores; 3) Assets older than four years. One of the problems with manual grading is that these rules are applied arbitrarily, hence assessment is very subjective. In addition, sometimes graders may be partially influenced by seemingly less important attributes such as hardware vendor or CPU name. There are many potential advantages to an automatic grading approach: it may improve the quality, consistency and reliability of the grading; it can provide an actual efficiency ranking score for each asset instead of just assigning it to say one of three efficiency categories. In this section, we introduce a number of efficiency criteria which we evaluate in this paper with respect to their ability to automatically label the efficiency of assets.

- *CPU* values are not comparable in a heterogeneous computing environment. It is therefore necessary to normalise the CPU values; we use the SPEC CPU2006¹ benchmark values, namely an average between SPECint_rate_base2006 and SPECfp_rate_base2006, which measure the throughput of a machine running simultaneous tasks over a certain amount of time.
- *RAM* is measured in GBs.
- *Hardware Capacity* is the average between the CPU value and the RAM capacity of the asset.
- *Electrical Cost* is obtained by multiplying the maximum power consumption (Watts) of an asset by the energy price (\$/KWh) at a particular location, this result is then multiplied by the number of hours per month that an asset is powered on (e.g., 720 hours) plus the

¹<http://www.spec.org/cpu2006/>

cooling cost. For simplicity we consider that for each watt spent on powering an asset, another watt is needed for cooling it [3]

- *Asset Year* is the year in which the asset was purchased and deployed.

III. EXPERIMENTAL METHODOLOGY AND RESULTS

In this section, we present our experimental methodology and results. Our dataset for these experiments consists of 1,200 assets from two different data centers in IBM Toronto which hold a variety of hardware platforms (IBM System X, IBM System P, Sun, and HP) ranging in age from 2012 to 2008. Along with an asset's 'characteristics' the dataset also provides its corresponding *manual efficiency grade* assigned by a Capital Planner. There are three gradings: 704 assets are graded as 'More Efficient', 449 as 'Less Efficient' and 47 as 'Least Efficient'.

The experiments were carried out as follows: For each of the efficiency grading criteria described in Section II, a ranked array of assets based on their 'criteria score' is generated. Each asset identifier is then replaced with its manual efficiency grading. If the resultant array is a perfect ordering of 'More Efficient' followed by 'Less Efficient' followed by 'Least Efficient', then this criteria has managed to correctly grade all the assets. We use an evaluation metric called *accuracy* to quantify the similarity between the efficiency criteria array and the manual grading array. This metric is calculated by dividing the total number of agreements between the two arrays by the total number of assets in the dataset.

Table I presents the results of our experiments. *Random* represents a randomly assignment of efficiency labels to assets, where the number of efficiency grade instances in the dataset is maintained in the random labelling. *Majority label* assign assets to the most common efficiency label, i.e., the 'More Efficient' class. As expected our *Random* and *Majority Label* criteria achieve the lowest accuracy scores. These results represent a lower bound on performance for the task. We can also see that the best performing criteria is *Electrical Cost*, which indicates that the most costly assets in terms of cooling and power are also the More Efficient ones. This is an interesting result as manual efficiency grading, while somewhat arbitrary, still tends to be more heavily influenced by hardware criteria and asset year as described in the 'rules-of-thumb' presented in Section II. A closer inspection of our dataset reveals why the electrical cost criteria is so effective. Looking at the dataset, we see that if assets are roughly categorised into two asset types - racks and workstations. The data centre demographic has the following characteristics: 100% of the assets in the 'More Efficient' manual grading are racks, 94% of the assets in the 'Least Efficient' grading are workstations, and 93% are workstations and 7% are racks in the middle grading, 'Less Efficient'. Since racks are always more power hungry

Efficiency Criteria	Accuracy %
<i>Electrical Cost</i>	96.24%
<i>RAM</i>	88.15%
<i>Hardware Capacity</i>	72.81%
<i>CPU</i>	72.47%
<i>Asset Year</i>	72.10%
<i>Majority Label</i>	58.71%
<i>Random</i>	49.05%

Table I
GRADING ACCURACY OF ASSET EFFICIENCY CRITERIA WHEN
COMPARED TO MANUAL EFFICIENCY GRADINGS

than workstations, it is clear now why the electrical cost criteria is such a strong predictor of the asset's efficiency grade. This observation leads us to the conclusion that the performance of the various criteria we explore will change significantly with respect to the asset demographic of the data centre. So while the results in this paper are interesting, they do not generalise to other data centres. Therefore, the contribution of this paper is the presentation of an experimental methodology for exploring efficiency grading criteria. Our intention now is to extend this work by exploring the effectiveness of these criteria on different data collections with different asset demographic characteristics.

IV. CONCLUSION AND FUTURE WORK

This paper investigates the use of common grading criteria or rules-of-thumb used in data centres to manually assess the efficiency of a computing asset. To the best of our knowledge this is the first publication to explore the automatic efficiency grading of assets. In future experiments, we plan to show that the effectiveness of the various efficiency criteria is dependent on the asset demographic of the data centre under consideration. Another important direction for future work is to investigate a *supervised machine learning approach* as a means of automatically grading assets using a combination of evidence from all the criteria scores.

ACKNOWLEDGMENT

This work was supported by Science Foundation Ireland grant 10/CE/I1855 to Lero in cooperation with IBM.

REFERENCES

- [1] Luiz André Barroso and Urs Hölzle. The Case for Energy-Proportional Computing. *Computer*, 40(12):33–37, December 2007.
- [2] Albert Greenberg, James Hamilton, David A Maltz, and Parveen Patel. The Cost of a Cloud : Research Problems in Data Center Networks. 39(1):68–73, 2009.
- [3] Jonathan G Koomey. Worldwide electricity used in data centers. *Environmental Research Letters*, 3(3), 2008.