


Accepted Manuscript

Title: Mastering system and power measures for servers in datacenter

Author: Georges Da Costa Jean-Marc Pierson Leandro Fontoura-Cupertino



PII: S2210-5379(17)30031-8
DOI: <http://dx.doi.org/doi:10.1016/j.suscom.2017.05.003>
Reference: SUSCOM 168

To appear in:

Received date: 14-2-2017
Accepted date: 1-5-2017

Please cite this article as: Georges Da Costa, Jean-Marc Pierson, Leandro Fontoura-Cupertino, Mastering system and power measures for servers in datacenter, <![CDATA[Sustainable Computing: Informatics and Systems]]> (2017), <http://dx.doi.org/10.1016/j.suscom.2017.05.003>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Mastering system and power measures for servers in datacenter

Georges Da Costa, Jean-Marc Pierson, Leandro Fontoura-Cupertino
IRIT, University of Toulouse, France

ABSTRACT

Using power meters and performance counters to get insight on system's behavior in terms of power consumption is common nowadays. The values coming from these external or internal meters are usually used directly by the research community, for instance to derive higher-level power models with learning techniques or to use them in decision tools such as schedulers in HPC and Cloud Computing. While it is reasonable when one wants only to have a broad view on the power consumption, they can not be used directly in most cases: We prove in this article that the problems of distributed measure and hardware limits are way more complex and create bias, and we give the keys to understand and chose the proper methodology to handle these bias to obtain relevant values for enhanced usage. A generic methodology is analyzed and its main lessons extracted for a direct usage by the research community to master system and power measures for servers in datacenter.

1. INTRODUCTION

While it is well known that measurements setups, jitter in acquisition monitoring, lost and repeated observations, inaccuracies of different meters, system counters availabilities, processor performance counters monitoring are all bias when dealing with power and system measures, there does not exist a consolidated analysis of these problems. The quality of performance and power measurements relies not only on the available physical infrastructure's accuracy but also on how the experiments are conducted. Interestingly, while everyone would agree on the importance of the process of measurements and the knowledge of the expected accuracy, many do not discuss enough of its importance. In particular, depending on the use cases some inaccuracies can be acceptable or not. For instance, power and performance counters (PMC) measurements' high accuracy is of great importance when used for regression models for creating power models, since the precision of the model will depend on the accuracy of the learning data [15]. Alternatively, when the question is only to estimate the maximum power consumption of an

infrastructure of several hundreds of servers, an accuracy of 100 Watts or 10s will not even make a difference. In the same vein, a cloud management system having to place, consolidate, migrate services among servers will probably only need accurate enough values every hour and not at high frequency.

For each of the analyzed bias, we will demonstrate through real experiments its impact in terms of accuracy of the power estimates, and the extra power needed when taking it into account. These results serve the community in order to take wise decisions on questions like: which biases should I take into account in a particular case? how to improve the behavior of my platform? We will exhibit lessons learnt for challenges that are commonly faced.

Two aspects will draw particular attention: First, we define a novel model of Power Supply Unit (PSU) power conversion losses. Second, we exhibit the overhead of tracking performance counters on Intel i7 processors and the need to limit the number of concurrent monitoring.

The remainder of this article is organized as follows : Section 2 describes the capabilities of hardware devices used to measure power consumption. In Section 3, the data acquisition infrastructure is detailed. It is followed by Section 4 concerning the power measures bias. Section 5 outlines the problematic of measuring system values. Finally we conclude this research work in Section 6.

2. POWER MONITORING DEVICES

Hardware power meters are the most accurate source of system's power measurements. The granularity of measurement is crucial for both, power measuring and modeling, and depends on the type of power meter used: external or internal. External meters are placed between the electric outlet and system's power supply unit, while internal meters are located inside the system [24]. Several studies rely on the precision of the power and system monitoring infrastructure to correlate both but without details on the exact methodology. In [33] or in [20] for example, authors model power consumption of VMs without explaining in details the bias of the measuring infrastructure and its impact on the resulting models.

2.1 Intra-node devices

Fine-grained measurements can be achieved by embedding power sensors into the system, enabling device specific mea-

Download English Version:

<https://daneshyari.com/en/article/4962761>

Download Persian Version:

<https://daneshyari.com/article/4962761>

[Daneshyari.com](https://daneshyari.com)