Chapter 6
Energy-Efficient Monitoring and Controlling of Computer Systems

Micha vor dem Berge
christmann informationstechnik & medien GmbH & Co. KG, Germany

Wolfgang Christmann
christmann informationstechnik & medien GmbH & Co. KG, Germany

ABSTRACT

This chapter aims to describe the current state and the ongoing efforts to integrate a monitoring and controlling architecture in modern computing environments. Benefits to be gained from this architecture include: More knowledge about the available computing resources and the capability to establish a management platform that optimizes energy efficiency as well as the availability of the computer environment in a time of a rapidly growing need for computer systems. In two use cases, namely server virtualisation and High Performance Computing (HPC), implementations of a management system for energy-efficiency were evaluated, showing that it is possible to increase the energy-efficiency by more than 10%, depending on the use case and the workload.

INTRODUCTION

Many of today's computer systems are set up with a strong focus on availability and reliability on the one hand and low acquisition costs on the other hand. When customers have to choose which computer systems to buy, the Total Costs of Ownership (TCO) are often thrust aside as if they were of lower importance than the purchasing price, although they are the proper measure for overall costs. The TCO usually consist of the acquisition or depreciation costs, maintenance and service costs as well as disposal costs or resale value. In case of a computer system that has to provide

DOI: 10.4018/978-1-4666-2190-9.ch006

high availability, costs resulting from downtime can be considered as well. Most server systems are running 24 hours a day, 7 days a week and 365 days a year. Because of this, the energy costs for operation and cooling are among the biggest expenses during the lifetime of a server. These costs can be reduced by managing all servers in a sensible and energy-efficient way.

Another aspect of the reduction of energy consumption is the accompanying reduction of CO2 emissions. The worldwide CO2 emissions of data centres are already approximately equal to the total CO2 emissions of airlines and account for almost 2% of global production and has a rapidly growing trend. The worldwide data centre electricity consumption grew up from 70.8 Billion kWh in the year 2000 to 152.2 Billion kWh in 2005 (Koomey, 2008).

Thus, the efficiency of IT systems and data centres becomes a priority for research and industry. High Performance Computing is considered one of the major issues to be addressed in this context. Companies like Google or Facebook know the problems of running thousands of servers (which regularly show up in the form of an immense electricity bill). They have begun to build their own energy-efficient servers.

In this chapter we are going to explain how complex HPC environments can be monitored and managed in an energy-efficient way while keeping the additional payload for the entire system as small as possible. This can be done via standardized software and hardware access layers or by new monitoring and controlling architectures that enable the user to reach a maximum of energy efficiency. One new type of monitoring and controlling architecture for cluster servers will be introduced, which was developed by the German company Christmann, see Christmann (2012).

BACKGROUND

As described in the introduction, the increasing worldwide demand on energy for data centres is becoming a serious problem. If the energy required running modern data centres would increase at a constant rate over the next years, we would theoretically need all power plants of the world just to run the data centres in a few years. For this reason, the European Council has confirmed the objective to save 20% of the EU's energy consumption compared to projections for 2020 (CEC, 2009), while the Climate Group (2008) has estimated that ICT-enabled improvements could save about 15% of total carbon emissions by 2020, even in other sectors than the ICT sector. Every IT company must deal with this emerging problem that more and more companies are already aware of. By now there are several rather conventional factors that can increase the energy efficiency of data centres, which will be briefly presented in the following passage.

The first possibility to increase energy efficiency is to scrutinise existing data and applications. Often it is possible to shut down some servers just because the running application or existing data is no longer used, but in many scenarios the necessity of applications is not monitored, which makes it impossible to evaluate the necessity of servers. Also, the amount of resources used by different programs might vary a lot and should be compared critically.

The second factor is to use server virtualisation. By doing so it is possible to migrate old servers to virtual machines (VM) and to consolidate several VMs on one physical server (host). In most cases, a prerequisite is a newer hardware with CPU support for Virtualisation Technology (called Intel VT or AMD-V, see details at AMD

Related Content

Methods for Statistical and Visual Comparison of Imputation Methods for Missing Data in Software Cost Estimation
Lefteris Angelis, Panagiotis Sentas, Nikolaos Mittasand Panagiota Chatzipetrou (2011). *Modern Software Engineering Concepts and Practices: Advanced Approaches (pp. 221-241).*
www.irma-international.org/chapter/methods-statistical-visual-comparison-imputation/51974

EfficientNet-B0 Model for Face Mask Detection Based on Social Information Retrieval
Moolchand Sharma, Harsh Gunwant, Pranay Saggar, Luv Guptaand Deepak Gupta (2022). *International Journal of Information System Modeling and Design (pp. 1-15).*
www.irma-international.org/article/efficientnet-b0-model-for-face-mask-detection-based-on-social-information-retrieval/313444

Object-Role Modeling: Principles and Benefits
Terry Halpin (2010). *International Journal of Information System Modeling and Design (pp. 33-57).*
www.irma-international.org/article/object-role-modeling/40952

MAE-CoReNet: Masking Autoencoder-Convolutional Reformer Net
Di Wang, Li Wangand Yuyang Zhou (2025). *International Journal of Information System Modeling and Design (pp. 1-23).*
www.irma-international.org/article/mae-corenet/371399

A Two-Level Multi-Modal Analysis for Depression Detection From Online Social Media
Dhrubasish Sarkar, Piyush Kumar, Poulomi Samanta, Suchandra Duttaand Moumita Chatterjee (2022). *International Journal of Software Innovation (pp. 1-22).*
www.irma-international.org/article/a-two-level-multi-modal-analysis-for-depression-detection-from-online-social-media/309114