

Eco-Efficient Cloud Resource Monitoring and Analysis

An approach for saving energy and reducing the carbon footprint of cloud infrastructures

Axel Tenschert, Pavel Skvortsov, Michael Gienger
HLRS – High Performance Computing Center
University of Stuttgart
Stuttgart, Germany
{tenschert, skvortsov, gienger}@hls.de

Abstract—This paper presents an approach for monitoring cloud computing resources by using customized monitoring metrics. In addition, based on the monitoring results, this approach includes a data mining analysis for making assumptions regarding consumed power of infrastructure and virtual machines (VMs) aiming at a reduced carbon footprint.

Index Terms—Cloud computing, Monitoring, Data Mining, Eco-efficiency, Carbon footprint

I. INTRODUCTION

This work presents an enhanced monitoring approach considering the infrastructure of providers' sites and the virtual machines (VMs) hosted on physical nodes. It considers the GAMES [1] and the OPTIMIS [2] projects, that were both developing methods for energy-aware resource usage. As the monitoring infrastructure (figure 1) is based on standardized software, customization is the goal for ensuring the relevant values to reduce the power consumption and in addition, the eco-efficiency of the infrastructure providers and the overall carbon footprint.

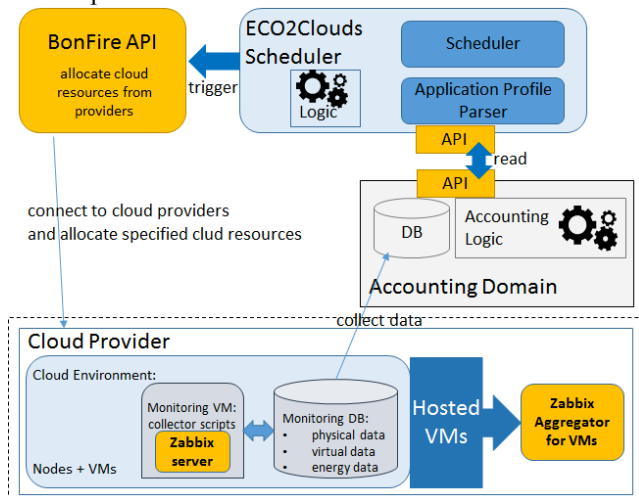


Figure 1: The monitoring infrastructure

The presented monitoring approach demonstrates a monitoring system making use of a Zabbix [3] monitoring

server being connected to Zabbix client agents installed on each physical node. Additionally, a data mining analysis is performed on the monitored data set.

The developed and implemented concept for monitoring cloud resources makes use of the European funded BonFIRE [4] project that was extended by the ECO₂Clouds [5] project. ECO₂Clouds adapts and extends the already provided monitoring capabilities by the providers EPCC, INRIA and HLRS with the aim to reduce the power consumption and improve the carbon footprint through monitoring and analyzing required data. One major step of the overall approach is to elaborate the power consumption of VMs.

This paper is structured as follows: (1) Introduction, (2) Monitoring and Metrics, (3) Data Mining Analysis and (4) Conclusions.

II. MONITORING AND METRICS

The ECO₂Clouds monitoring system is a layered approach distinguishing mainly between the physical infrastructure and the virtualization layer. Additionally, applications can be monitored as well but as this information is highly dependent on the application, this paper won't target the specific application monitoring. For those three layers, a Zabbix client agent is gathering monitoring information and sends it to the Zabbix server in order to provide a monitoring information repository to assist in optimization decision making. To enable power measurements for all three layers, power distribution units are attached to the physical servers. The monitored parameters are derived as defined by customized monitoring metrics based on standard programming languages like Python, Ruby or even Bash.

A. The infrastructure and virtualization layer

The infrastructure layer is represented by the physical infrastructure of a cloud provider, including physical nodes, storages or network devices. Thus, the developed monitoring approach needs to measure parameters of those physical components in order to enable a data analysis based on the current state of the infrastructure for making assumptions regarding the overall system behavior. Besides the

infrastructure layer, especially the virtualization layer is a crucial part as more or less the virtualization defines a completely separated environment for applications. Thus, it is essentially important to monitor those two layers in a fine grained manner. Furthermore, for providing accurate information for an optimized power consumption, both layers have to be linked together by the creation of calculated metric as presented in figure 2.

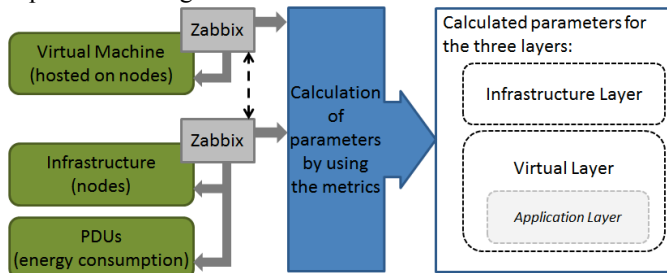


Fig. 2: The layered monitoring approach

It presents the structure of the layered approach by presenting the virtual machines (VMs), the infrastructure and the power distribution units (PDUs) for measuring the power consumption of physical nodes and showing the connection of those three via Zabbix to the three layers. To perform all necessary calculations being defined in the set of monitoring metrics the Zabbix server (infrastructure) and the aggregator (Zabbix at VM level) need to be able to communicate to each other.

Coming back to the *infrastructure layer*, it refers to the characteristics of the sites and the resources available in each of them. Thus a set of metrics was defined for measuring relevant infrastructure parameters and calculate them for enabling data analysis.

B. The virtualization layer

The virtualization layer contains the underlying hypervisor including the virtual machines (VMs) hosting the applications (or part of applications) that are used to manage the cloud infrastructure (e.g., the monitoring system) or that are offered to the users (e.g., user specific business applications). In order to create separated environments, the virtualization layer can be freely configured: the amount of CPU, memory and disk space can be aligned to the used application hosted inside.

Metrics at the virtualization layer aim to characterize the VMs on which the applications are running. They are analyzed to evaluate if the current deployment can be further improved and thus optimized. The analysis of the VM energy consumption aims to understand how the energy consumed by the host is distributed among the deployed VMs. Moreover, new metrics were defined inspired by the data center metrics (infrastructure) proposed in the last years, especially by The Green Grid Consortium (GreenGrid). The idea is to redefine the classical infrastructural metrics, like PUE and Data Centre Energy Productivity (DCeP), at virtualization level to measure the impact of the application tasks in terms of energy consumption and carbon emissions.

As already stated, the applications inside a VM can also be monitored in particular. But due to the application character, the metrics are highly specialized and cannot be regarded in a generic fashion.

C. The monitoring metrics

This subsection will describe some of the most important metrics to improve the energy consumption as well as the carbon footprint for the infrastructure providers as well as for the users.

For improving the carbon footprint of an infrastructure provider, metrics like “Carbon Usage Effectiveness of the Provider”, “Green Efficiency Coefficient of the Electrical Supplier” or the “Site Infrastructure Efficiency” are considered. They are dealing with the amount of energy consumed and in particular, on which basis the requested energy was generated. As presented in the Data Mining section, there are big differences in the amount of CO₂ created to power the machines at different days or even times. For infrastructure providers, the mandatory metrics are global: monitoring a single host is not sufficient. In order to understand the infrastructure load, metrics like “Service Availability” as well as “Storage and CPU as well as Memory Utilization” are also of interest and have to be considered.

The power consumption of VMs depends on the size of a VM defined through the used memory, the data I/O identified through the send and receive activities, the disk activity identified through the read and write operations and the consumed CPU seconds. Thus, the infrastructure and the virtualization layer are involved to derive the information for this important metric (figure 3).

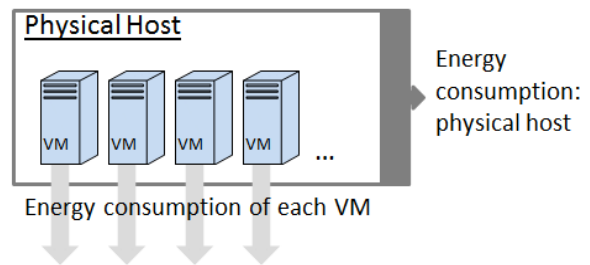


Fig. 3: Power consumption per VM

Finally, the power information of a VM can be correlated with the “Green Efficiency Coefficient of the Electrical Supplier” resulting in additional metrics. But the results are impressive: the calculation of the virtual machine power consumption is possible and even more, the produced carbon can be observed as well. Especially for distributed computing and data centers, there are differences for the location and timeslot hosting the virtual instances.

III. DATA MINING ANALYSIS

The monitoring infrastructure described before is producing large amounts of data in a short period of time. This data is being collected into Accounting DB by the Accounting Service, which runs on a separate host. During our experiments, we have observed that the size of the collected

metrics data varies between 10 and 20 Mb per day. As the result, the fast-growing SQL database of the Accounting Service host may shortly reach the storage capacity limits of the Accounting node. This would trigger the automatic deletion of the old data by Zabbix. In order to avoid the loss of historical raw data and to allow for its statistical analysis, we have developed the ECO₂Clouds Data Mining service (DM Service). Next, we present architecture, concepts and analysis principles of the DM Service in more detail.

A. Architecture

DM Service consists of two major components (figure 4). The first component is running on the Accounting Service to perform transfer of non-reduced metrics data to a remote data storage (DM Storage), and to generate a reduced data set.

The second component is DM Storage, which gathers the non-reduced metrics data and performs statistical analysis over them, e.g., correlation analysis over a large enough portion of data. The resulting stable parameters of this analysis are inserted into the Accounting DB as a separate table.

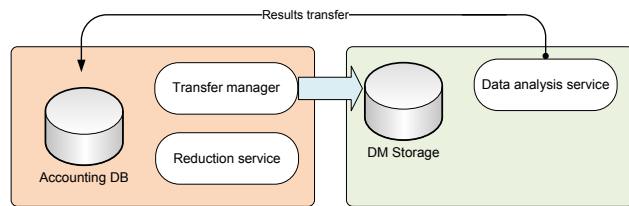


Fig. 4: Data Mining Architecture

B. Concept

The Data Mining workflow consists of the following steps. (1) Triggered daily, the Accounting Service inserts the new non-reduced data into the database of DM Storage. (2) Triggered daily, the Accounting Service performs the reduction of the metrics data; the results of data reduction are inserted into a separate database on the Accounting VM. (3) The initial non-reduced data are deleted – thus, each time the Accounting Service has to perform operations over a metrics data set which was produced only during the last 24 hours

The resulting loss of data at the Accounting DB does not affect the Scheduler functionality, because it does not use the historical data for physical hosts, while for experiments we only reduce the data of finished (non-active) experiments only.

On the DM Storage side, the statistical analysis of the data is performed. The results are inserted into Accounting DB.

C. Implementation

On the Accounting VM, we install a bash script which performs the following operations: (1) Creates a temporary SQL-dump of the current non-reduced “e2c_collector” database. (2) Imports the generated SQL-dump into DM Storage by inserting the rows incrementally to the previously stored there data. (3) Removes the temporary SQL-dump from the Accounting VM. (4) Calls the reduction service implemented as a Java program, which performs the data reduction and inserts its results into a separate database on the same VM.

D. Data Analysis

On the DM Storage side, the metrics data analysis is being performed. The aim of statistical data analysis is to find correlation patterns, which may help to define the deployment strategies depending, for instance, on the time of the day.

One example a pattern of the power consumption and emissions depending on the time of the day is shown in Fig. 5.

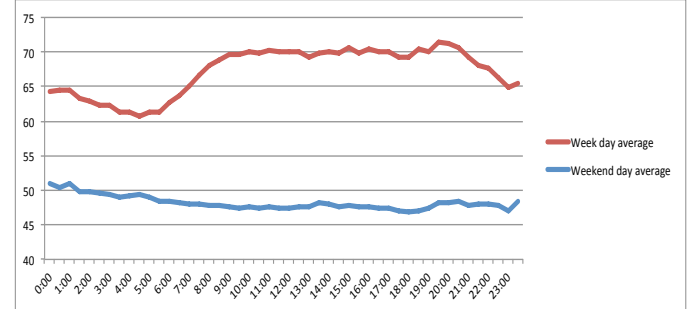


Fig. 5: Emissions during the week and weekend days [6] (see ECO₂Clouds project publications of WP3 for detailed analysis methods)

IV. CONCLUSIONS

The presented monitoring architecture uses customized monitoring metrics and enables a data analysis providing result data regarding a cloud system behavior. The focus of the approach is the monitoring and analysis of data related to the used energy mix and consumed power on infrastructure and virtualization level.

Thus, it becomes possible to align the VM scheduling to the workload of the physical infrastructure by considering the power consumption and the carbon footprint through knowing about the energy mix. Herewith, costs for consumed power and CO₂ emissions are reduced for a cloud infrastructure.

Future work will be using the presented monitoring infrastructure and the implemented metrics as basis for running experiments and supporting the deployment of use cases.

ACKNOWLEDGMENT

This work has been supported by the ECO₂Clouds project (<http://eco2clouds.eu/>) and has been partly funded by the European Commission’s IST activity of the 7th Framework Programme under contract number 318048. This paper expresses the opinions of the authors and not necessarily those of the European Commission. The European Commission is not liable for any use that may be made of the information contained in this paper.

REFERENCES

- [1] M. Bertoincini, B. Pernici, I. Salomie, S. Wesner: GAMES: Green Active Management of Energy in IT Service Centres. CAiSE Forum 2010: 238-252.
- [2] G. Katsaros, J. Subirats, J. O. Fitó, J. Guitart, P. Gilet, D. Espling, A service framework for energy-aware monitoring and VM management in Clouds, Future Generation Computer Systems, Vol. 29, 8, October 2013.
- [3] The Zabbix website documentation, <https://www.zabbix.com>, Last visited: 27.06.2014.

- [4] The BonFIRE Project Website, <http://www.bonfire-project.eu/>, Last visited: 27.06.2014.
- [5] The ECO₂Clouds Project Website, <http://eco2clouds.eu/>, Last visited: 27.06.2014.
- [6] C. Cappiello, P. Melià, B. Pernici, P. Plebani, M. Vitali. Sustainable choices for cloud applications: a focus on CO₂ emissions. Proceedings of the ICT4S, 2014.