www.cost804.org

Action IC0804

Energy and QoS-Aware Workload Management in Clouds

Focus group proposal

Duration: 12 months

Authors

Eugen Feller (INRIA Rennes)
Dr. Ivona Bradnic (TU Vienna)

Institutes

INRIA Rennes - Bretagne Atlantique, Campus universitaire de Beaulieu, 35042 Rennes Cedex, France

Vienna University of Technology, Argentinierstrasse 8/184-1, 1040 Vienna, Austria

30. March 2011

Group leaders

Contents

1 Motivation

The growing demands for resources by compute intensive service and scientific applications have led the data center providers to build a large amount of warehouse-sized data centers. These data centers require a significant amount of power to be operated and hence consume a lot of energy. In particular, studies have shown that data centers alone have consumed 61 billion kWh of U.S. energy in 2006. This is enough energy to power 5,8 million average U.S households and results in approximately $4.5 billion/year of energy costs [Age07]. The amount of consumed energy is most likely to increase up to 120 billion kWh by 2011 in case no further energy conservation steps are taken [Age07]. Moreover, the way energy is generated influences our environment either directly by the carbon footprint or indirectly by the nuclear waste. According to the Environment Protection Agency (EPA) decreasing the energy consumption could reduce the carbon dioxide emissions by 15 to 47 million metric tons in 2011 [Age07].

Due to the increasing user demands for computing power and the emergence of the *cloud computing* paradigm, long term energy consumption will most likely increase in case users and cloud providers are not made energy-aware. Some cloud providers (e.g. Yahoo!) have already started to build their new data centers in areas where *green energy* is available [Yah09]. Moreover, old hardware is replaced with new, more energy-efficient one. However, most of the data centers are still powered by coal-fired or nuclear plants and require dizzying amounts of energy [Int10]. Furthermore, despite the enormous potential for energy savings due to low average data center utilization of approximately 15-20% [Vog08], current cloud providers (e.g. Amazon EC2) do not detail how users could be involved in energy conservation and how energy consumption is optimized by their cloud management software. Moreover, although most of the current cloud providers offer some means of specifying service level agreements (SLAs), they are solely limited to resource availability and do not consider many other important metrics such as response time and throughout. Consequently, they cannot use these metrics to minimize the energy consumption while trying to guarantee a certain level of QoS.

Given the low average utilization in current data centers, big fraction of the resources can be used to take energy conservation decisions such as consolidating the workload on a subset of machines and suspending or turning off the resulting idle servers. In combination with the ability of today's ubiquitous virtualization (e.g., Xen [BDF+03], KVM [Hab08], OpenVZ [ope], etc.) and Single-System-Image (SSI) [MLV+04] solutions to migrate workload among physical machines, cluster sizes can be dynamically adapted depending on the resource demands. Thereby, several issues still need to be resolved: Which workload needs to be moved? Which machines are the most suitable ones with respect to energy and performance to host this workload? Answers to these questions ideally require accurate models for power, performance, migration costs and algorithms which use these models to take decisions.

In the following chapters we propose the creation of a focus group to deal with the energy and QoS-aware workload management issues in current and future cloud providers at all levels of the cloud computing stack (i.e., SaaS, PaaS and IaaS). Therefore, Chapter 2 details why we have decided to initiate this focus group, describes the problem to be solved and what should be the outcome of this collaboration. Afterwards, Chapter 3 outlines a workplan for the next 12 months. Finally, Chapter 4 presents the currently involved partners.

2 Description

During the last two meetings of COST Action IC0804 in Passau and Portugal we have seen many talks related to energy conservation in wireless and wired networks. On the other hand, we have also seen several efforts dealing with energy-efficient resource management of virtualized infrastructures such as clouds. These works can be divided into four categories: analytical and empirical models (e.g., machine learning) [BGN$^+$10] for *estimating the energy consumption*, *workload (i.e., VM) placement algorithms* [BDCPS09], *SLA and workload management algorithms* [MBS10] and *system architectures* [FRM$^+$10]. The following figure depicts the position of the focus group partners in each category.

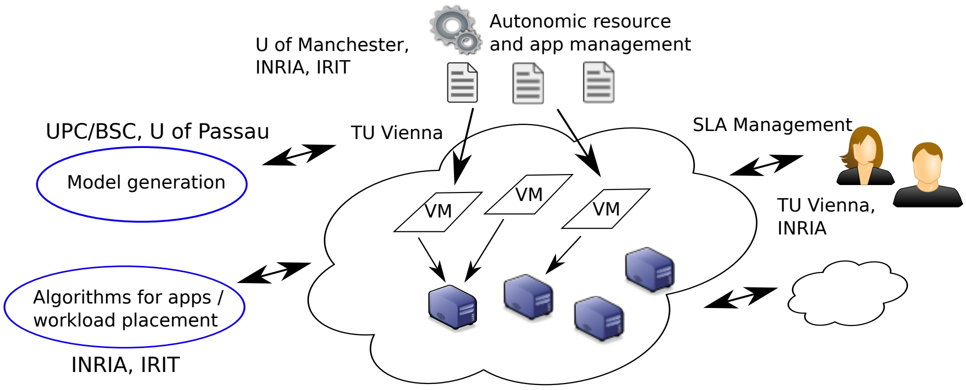

Figure 1: Partners position in the focus group

On the modeling side we have seen a clear lack of models for estimating the impact of workload migration on performance and energy. Thereby, current algorithms are not able to take accurate placement decisions. Furthermore, existing models for estimating the workload energy consumption do not take into account the variety of available resources (e.g., CPU, RAM, disk, etc.).

From the algorithmic point of view, workload mapping to physical machines is usually modeled as an instance of a multidimensional bin-packing problem in which the bins represent the physical machines and the items the workload to be packed. The objective of the algorithm is then to find a mapping of the items to the bins such that the number of bins is minimized, and all the the predefined constraints: capacity limits of bins, items resource (e.g., CPU, RAM) and QoS requirements (i.e., response time, throughput) are satisfied. Instances of this problem have been extensively studied in the past and shown to be NP-hard [LMV02]. Therefore, a variety of approaches for obtaining either approximate (e.g., heuristics) or exact solutions (e.g., linear/constraint programming) have been proposed in the area of operations research.

However, workload placement problem still differs from the traditional bin-packing problem as many workloads are of non-additive nature. Hence, co-allocation of workload with same characteristics (e.g., CPU bound) can lead to significant performance degradations. Furthermore, state-of-the-art heuristics (e.g., First-Fit) to solve these kind of problems can not be directly applied in production environments as they tend to waste a lot of resources. For example, typically First-Fit-Decreasing (FFD) heuristic assumes that all the items are available a priori, sorts them in decreasing order according to the load a certain resource (e.g., CPU) and starts assigning each item to a bin starting from the first one which has enough capacity to host it. Thereby, much of the remaining bin capacity (e.g., RAM, I/O, etc.) is wasted. Furthermore, these kinds of simple heuristics do not fully utilize the migration functionality provided by ubiquitous virtualization and SSI solutions to optimize the placement. In particular, already existing workload within a server is not allowed to be migrated upon new workload arrival in order to optimize the resource usage. For example, it could be advantageous to migrate already existing workload in favor of arriving workload in order to increase the resource utilization. This behavior results in resource wastings even for simple cases where only one resource (e.g., CPU) is considered. On the other hand, methods which compute exact solutions can only be used for problems of smaller size as they require exponential time to compute a placement. Thus, their practicability to be used in large-scale environments with many thousands of nodes, such as clouds, is limited.

In addition, SLAs expressed in terms of QoS (i.e., response time and throughput) are often ignored while performing workload consolidation. This is crucial, especially in the context of cloud computing where users are required to pay for the utilized resources and thus want to be provided with a certain level of QoS. Consequently, energy and QoS-aware workload placement algorithms are required which take into account both, the user and data center objectives: *Maximize performance* and *minimize the energy consumption*.

Finally, most of the cloud management systems available nowadays are highly centralized, do not provide any means of fault-tolerance and are solely limited to virtualization technology, although alternative solutions (e.g., SSI) are available and could be complementary used in existing and future cloud computing environments. First attempts to provide scalability, fault-tolerance and software heterogeneity to clouds at the Infrastructure as a Service (IaaS) level have been made and need to be further investigated [FRM$^+$10].

The objective of this focus group is to federate the currently loosely coupled researchers in the COST action working on models, algorithms and system architectures for energy-efficient cloud computing environments. This involves everyone currently working or planing to work on: *analytical/empirical models for workload energy consumption* and *migration costs (i.e., performance and energy)*, *approximate or exact algorithms for energy and QoS-aware workload placement* and finally *decentralized energy-aware cloud architectures*.

The idea is to investigate how different works could be improved and combined together in order to provide a holistic problem solution, which has a solid theoretical background and practical applicability. Results from this focus group could lead to joint publications and further collaborations between the partners.

3 Workplan

The following workplan is preliminary and needs to be discussed during the meeting in Budapest.

Schedule	Description
20.05.2011	1st meeting (Budapest)
08.08.2011	First deliverable (State of the Art Green Clouds in Europe)
14.11.2011	2nd meeting (TBD)
30.01.2012	Joint publication (Integrated QoS-Aware Architecture for Green Clouds)

4 Partners

Institution	Responsible
Vienna University of Technology	Ivona Brandic
INRIA Rennes (MYRIADS research team)	Christine Morin
EMN/INRIA Nantes (ASCOLA research team)	Jean-Marc Menaud
IRIT	Jean-Marc Pierson
University of Passau	Hermann de Meer
UPC/BSC	Jordi Torres
Vienna University	Helmut Hlavacs
University of Manchester	Rizos Sakellariou

References

[Age07] AGENCY, U.S. Environmental P.: *EPA REPORT TO CONGRESS ON SERVER AND DATA CENTER ENERGY EFFICIENCY APPENDICES*. 2007

[BDCPS09] BORGETTO, D.; DA COSTA, G.; PIERSON, J.-M.; SAYAH, A.: Energy-aware resource allocation. In: *Grid Computing, 2009 10th IEEE/ACM International Conference on*, 2009, S. 183 –188.

[BDF+03] BARHAM, Paul; DRAGOVIC, Boris; FRASER, Keir; HAND, Steven; HARRIS, Tim; HO, Alex; NEUGEBAUER, Rolf; PRATT, Ian; WARFIELD, Andrew: Xen and the art of virtualization. In: *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles*. New York, NY, USA : ACM, 2003. – ISBN 1–58113–757–5, S. 164–177.

[BGN+10] BERRAL, Josep L.; GOIRI, Inigo; NOU, Ramon; JULIA, Ferran; GUITART, Jordi; GAVALDA, Ricard; TORRES, Jordi: Towards energy-aware scheduling in data centers using machine learning. In: *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. New York, NY, USA : ACM, 2010 (e-Energy '10). – ISBN 978–1–4503–0042–1, 215–224.

[FRM+10] FELLER, Eugen; RILLING, Louis; MORIN, Christine; LOTTIAUX, Renaud; LEPRINCE, Daniel: Snooze: A Scalable, Fault-Tolerant and Distributed Consolidation Manager for Large-Scale Clusters. In: *2010 IEEE/ACM International Conference on Green Computing and Communications (GreenCom-2010)*. Hangzhou, China, December 2010,

[Hab08] HABIB, Irfan: Virtualization with KVM. In: *Linux J.* 2008 (2008), Nr. 166, S. 8. – ISSN 1075–3583.

[Int10] INTERNATIONAL, Greenpeace: Make IT Green: Cloud computing and its contribution to climate change. Version: 2010. `http://www.greenpeace.org/raw/content/international/press/reports/make-it-green-cloud-computing.pdf`. 2010. – Forschungsbericht.

[LMV02] LODI, Andrea; MARTELLO, Silvano; VIGO, Daniele: Recent advances on two-dimensional bin packing problems. In: *Discrete Appl. Math.* 123 (2002), Nr. 1-3, S. 379–396. `http://dx.doi.org/http://dx.doi.org/10.1016/S0166-218X(01)00347-X`. – DOI http://dx.doi.org/10.1016/S0166–218X(01)00347–X. – ISSN 0166–218X.

[MBS10] MAURER, Michael; BRANDIC, Ivona; SAKELLARIOU, Rizos: Simulating Autonomic SLA Enactment in Clouds using Case Based Reasoning. In: *ServiceWave 2010*. Ghent, Belgium, December 2010,

[MLV+04] MORIN, C.; LOTTIAUX, R.; VALLEE, G.; GALLARD, P.; MARGERY, D.; BERTHOU, J.-Y.; SCHERSON, I.D.: Kerrighed and data parallelism: cluster computing on single system image operating systems. In: *Cluster Computing, IEEE International Conference on* 0 (2004), S. 277–286. `http://dx.doi.org/http://doi.ieeecomputersociety.org/10.1109/CLUSTR.2004.1392625`. – DOI http://doi.ieeecomputersociety.org/10.1109/CLUSTR.2004.1392625. ISBN 0–7803–8694–9.

[ope] *OpenVZ*. `http://openvz.org/`

[Vog08] VOGELS, Werner: Beyond Server Consolidation. In: *Queue* 6 (2008), Nr. 1, S. 20–26. http://dx.doi.org/http://doi.acm.org/10.1145/1348583.1348590. – DOI http://doi.acm.org/10.1145/1348583.1348590. – ISSN 1542–7730.

[Yah09] YAHOO!: *Yahoo! data center will be powered by Niagara Falls*. http://green.yahoo.com/blog/ecogeek/1125/yahoo-data-center-will-be-powered-by-niagara-falls.html. Version: 2009