Chapter 18
Towards Energy-Efficient, Scalable, and Resilient IaaS Clouds

Eugen Feller
INRIA Centre Rennes - Bretagne Atlantique, France

Louis Rilling
Kerlabs, France

Christine Morin
INRIA Centre Rennes - Bretagne Atlantique, France

ABSTRACT

With increasing numbers of energy hungry data centers, energy conservation has now become a major design constraint for current and future Infrastructure-as-a-Service (IaaS) cloud providers. In order to efficiently manage such large-scale environments, three important properties have to be fulfilled by the management frameworks: (1) scalability, (2) fault-tolerance, and (3) energy-awareness. However, the scalability and fault tolerance capabilities of existing open-source IaaS cloud management frameworks are limited. Moreover, they are far from being energy-aware. This chapter first surveys existing efforts on building IaaS platforms. This includes both, system architectures and energy-aware virtual machine (VM) placement algorithms. Afterwards, it describes the architecture and implementation of a novel scalable, fault-tolerant, and energy-aware VM manager called Snooze. Finally, a nature-inspired energy-aware VM placement approach based on the Ant Colony Optimization is introduced.

INTRODUCTION

Cloud computing has recently evolved as a new computing paradigm which promises virtually unlimited resources. Customers can rent resources based on the pay-as-you-go model. Resources are transparently provisioned by the cloud provider according to the customers' requirements. However, customers' growing demands for computing power are now facilitating the cloud service providers to deploy increasing amounts of energy hungry data centers (Greenpeace, 2010). Consequently, energy

DOI: 10.4018/978-1-4666-1631-8.ch018

costs for operating and cooling the equipment of such data centers have increased significantly up to a point where they are able to surpass the hardware acquisition costs.

Besides the possibility to replace the hardware with more energy-efficient one, reducing the energy wasted because of hardware over-provisioning is crucial. Today's data centers infrastructure is typically over-provisioned in order to sustain the service availability during periods of peak resource demand. However, resource demand in current data centers is usually of a bursty nature and thus results in a low average utilization of approximately 15-20% (Vogels, 2008). Therefore, a big fraction of the resources can be used to take energy conservation decisions such as suspending or turning off unnecessary servers, while still preserving the customer's performance requirements.

Several open-source cloud projects have been recently started to provide alternative solutions to public Infrastructure as a Service (IaaS) cloud providers. Examples of such cloud management frameworks include Eucalyptus (Nurmi et al., 2009), OpenNebula (*OpenNebula*, 2011), and Nimbus (*Nimbus*, 2011).

Given that ubiquitous virtualization solutions are able to live migrate the VMs and servers can be turned on and off at any time, clusters can be turned into dynamic pools of resources by these frameworks. However, two main drawbacks exist which prevent existing cloud management frameworks to efficiently manage current and future large-scale infrastructures: (1) high degree of centralization, (2) lack of advanced energy and QoS-aware VM placement algorithms. While the first one leads to single point of failure (SPOF) and limited scalability, the latter results in high energy costs and heat dissipation. This in turn decreases the hardware reliability.

In order to solve the first drawback more decentralized IaaS management frameworks are required. Similarly, energy-conservation algorithms are needed which are able to operate efficiently in such dynamic large-scale environments.

One possible approach to conserve energy is to perform VM consolidation. VMs are packed on the least number of physical nodes and over provisioned servers are transitioned into a lower power state (e.g., suspend). This problem can be modeled as an instance of the well-known multi-dimensional bin-packing (MDBP) problem and has been mostly studied by means of simulations in several works (Stillwell et al., 2010; Li et al., 2009). Because of the NP-hard nature of the problem and the need to compute the solutions in a reasonable amount of time, approximation approaches (i.e., heuristics) have shown to provide good results. However, many of the existing approaches nowadays still: (1) ignore the multi-dimensional character of the problem (Buyya et al., 2010; Stillwell et al., 2010), (2) adapt simple greedy algorithms (e.g., FFD), which tend to waste a lot of resources (Setzer & Stage, 2010) and are highly centralized.

This chapter is organized as follows. In the first part recent achievements in designing and implementing IaaS cloud computing frameworks are reviewed. The VM placement problem is defined as an instance of the multi-dimensional bin-packing (MDBP) problem and some of the existing bin-packing algorithms including their application to the energy and QoS-aware VM placement problem are discussed.

The second part of the chapter presents the ongoing work towards enabling energy-efficiency, scalability, resilience and self-management in IaaS clouds. A novel VM manager called Snooze is introduced whose design was driven by these properties.

In addition, in order to overcome the drawbacks of current VM placement algorithms, a novel nature-inspired energy-aware VM placement algorithm based on the Ant Colony Optimization is introduced.

Finally, the chapter closes with conclusions and future research directions.