Computer Science Review 28 (2018) 118–130

Contents lists available at ScienceDirect

Computer Science Review

journal homepage: www.elsevier.com/locate/cosrev

Distributed virtual machine consolidation: A systematic mapping study

Adnan Ashraf *, Benjamin Byholm, Ivan Porres

Faculty of Natural Sciences and Technology, Åbo Akademi University, Finland

ARTICLE INFO

Article history:
Received 9 May 2017
Received in revised form 19 February 2018
Accepted 27 February 2018

Keywords:
Cloud computing
Data center
Virtual machine
Consolidation
Placement
Energy-efficiency

ABSTRACT

Background: Virtual Machine (VM) consolidation is an effective technique to improve resource utilization and reduce energy footprint in cloud data centers. It can be implemented in a centralized or a distributed fashion. Distributed VM consolidation approaches are currently gaining popularity because they are often more scalable than their centralized counterparts and they avoid a single point of failure.

Objective: To present a comprehensive, unbiased overview of the state-of-the-art on distributed VM consolidation approaches.

Method: A Systematic Mapping Study (SMS) of the existing distributed VM consolidation approaches.

Results: 19 papers on distributed VM consolidation categorized in a variety of ways. The results show that the existing distributed VM consolidation approaches use four types of algorithms, optimize a number of different objectives, and are often evaluated with experiments involving simulations.

Conclusion: There is currently an increasing amount of interest on developing and evaluating novel distributed VM consolidation approaches. A number of research gaps exist where the focus of future research may be directed.

Contents

* Corresponding author.
E-mail address: adnan.ashraf@abo.fi (A. Ashraf).

https://doi.org/10.1016/j.cosrev.2018.02.003

1. Introduction

Energy footprint of cloud data centers is a matter of great concern for cloud providers [1]. According to a recent report [2], data centers in the United States consumed an estimated 70 billion kilowatt-hours of electricity in 2014, which corresponds to 1.8% of total United States electricity consumption. High energy consumption not only translates into a high operating cost, but also leads to huge carbon emissions. The ever increasing demand for computing resources to provide highly scalable and reliable services has caused an energy crisis [3]. The high energy consumption of data centers can partly be attributed to the large-scale installations of computing and cooling infrastructures, but more importantly it is due to the inefficient use of the computing resources [4]. Production servers seldom operate near their full capacity [5]. However, even at the completely idle state, they consume a substantial proportion of their peak power [6]. Therefore, under-utilized servers are highly inefficient.

Hardware virtualization technologies allow to share a Physical Machine (PM) among multiple, performance-isolated platforms called Virtual Machines (VMs) to improve resource utilization. Further improvement in resource utilization and reduction in energy consumption can be achieved by consolidating VMs on under-utilized PMs. The basic idea of VM consolidation is to migrate and place the VMs on as few PMs as possible and then release the remaining, unused PMs for termination or for switching to a low-power mode to conserve energy. A VM consolidation approach uses live VM migration to consolidate VMs on a reduced set of PMs. VM consolidation has emerged as one of the most effective and promising techniques to reduce energy footprint of cloud data centers [4,7].

Fig. 1 presents a simple hypothetical scenario to illustrate the VM consolidation process. The first half of Fig. 1 shows three PMs where each PM hosts multiple VMs and every VM uses a certain amount of the PM resources. It is assumed that due to some significant load variations, PM 2 and PM 3 have become under-utilized. The under-utilized PMs in such a scenario may continue to remain under-utilized for hours, days, or even weeks unless the existing VMs require more resources or some new VMs are placed on the under-utilized PMs. Therefore, it is difficult to provide a resource and energy efficient allocation of VMs without consolidation of VMs on the under-utilized PMs. The second half of Fig. 1 shows that after migrating all VMs from PM 2 to PM 3, PM 2 can be turned-off or switched to a low-power mode.

There is currently an increasing amount of interest on developing and evaluating efficient VM consolidation approaches for cloud data centers. Over the past few years, researchers have used a multitude of ways to develop novel VM consolidation approaches. Some of these approaches have been recently reported in the form of nonsystematic literature reviews such as [8] and [9]. However, the drawback of these existing nonsystematic studies is that they provide a partial and possibly biased overview of the state-of-the-art on VM consolidation. For a comprehensive and unbiased coverage of the existing literature on VM consolidation, there is a need to study the existing VM consolidation approaches in a systematic way.

VM consolidation can be implemented in a centralized or a distributed fashion. Traditional VM consolidation approaches, such as [4,10–15], tend to be centralized. A centralized VM consolidation approach uses a centralized algorithm on a centralized architecture and does not provide support for multiple, geographically distributed data centers. The main drawbacks of centralized VM consolidation approaches include limited scalability and lack of robustness due to a single point of failure. On the other hand, a distributed or decentralized VM consolidation approach uses a distributed algorithm or a distributed architecture for PMs [7,16] or provides support for multiple, geographically distributed data centers [17,18]. Distributed VM consolidation is a recurring theme in recent VM consolidation approaches such as [7,19,16]. Distributed approaches are gaining popularity because they have benefits over centralized approaches. They are often more scalable than their centralized counterparts and they avoid a single point of failure [19,20]. Feller et al. [16] showed that their proposed VM consolidation algorithm does not compute a solution (in a reasonable amount of time) on a centralized architecture, but finds a good solution on a distributed architecture. Lucanin and Brandic [17] reported that their VM consolidation algorithm for geographically distributed data centers finds a good solution for a large-scale problem comprising ten thousand VMs. Sedaghat et al. [21] showed that their proposed distributed VM consolidation algorithm scales to tens of thousands of PMs and VMs without compromising on the quality of the solution. Sedaghat et al. [22] reported that their proposed distributed VM consolidation algorithm finds a near-optimal solution for 100,000 PMs in a reasonable amount of time. Marzolla et al. [19] showed that their proposed distributed VM consolidation algorithm is resilient to major failures and outages involving a thousand PMs. Therefore, distributed VM consolidation approaches are more suitable for large-scale data centers involving thousands of VMs and PMs.

We present a systematic study of the existing distributed VM consolidation approaches. The objective is to present a comprehensive, unbiased overview of the state-of-the-art on distributed VM consolidation approaches. Considering the broad nature of the research objective, it was not appropriate to launch a Systematic Literature Review (SLR). Therefore, we launched a Systematic Mapping Study (SMS) [23–25]. A SMS follows the same principled process as a SLR, but: (1) it has a broader scope, (2) it uses different criteria for inclusions/exclusions and quality assessments, and (3) its data collection and synthesis tend to be more qualitative than for a SLR [26]. It is "intended to 'map out' the research that has been undertaken rather than to answer a detailed research question" [24].

We proceed as follows. Section 2 presents the design and schedule of our study. The results of the SMS are presented in Section 3. In Section 4, we discuss major threats to the validity of the results presented in this paper. Finally, we present our conclusions in Section 5.

2. Study design

One of the most important differences between a nonsystematic literature review and a SMS is that a SMS follows an unbiased and repeatable process. Moreover, the process is documented as a review protocol. Therefore, we defined a review protocol for our SMS on distributed VM consolidation approaches. In this section, we present the design of our study and the review process.