Computer Networks 108 (2016) 78–96

Contents lists available at ScienceDirect

Computer Networks

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

Review article

Building a cloud on earth: A study of cloud computing data center simulators☆

Mohamed Abu Sharkh[a,*], Ali Kanso[b], Abdallah Shami[a], Peter Öhlén[b]

[a] Department of Electrical and Computer Engineering, Western University, London, Canada
[b] Ericsson Research, Montreal, Canada

ARTICLE INFO

Article history:
Received 4 August 2015
Revised 30 May 2016
Accepted 27 June 2016
Available online 22 July 2016

Keywords:
Cloud computing
Cloud simulators
Scalability
Data centers
Virtualization
Network and systems monitoring and measurements

ABSTRACT

As cloud computing technologies finalize their transformation into the standard technologies for businesses of all sizes, they face more scrutiny than ever. Clients are expecting the benefits of turning infrastructure, platform and network into services payable per use without tolerating any service hiccups caused by performance bottlenecks or overprovision. This puts cloud providers under pressure to deliver data center management solutions and deployment plans in minimal time and with failure allowance close to none. Any comprehensive solution evaluation could gain much from the use of cloud simulators. Cloud simulators have the advantage of practicality over both mathematical proofs and real testbeds. They support any amount of heterogeneous use cases demanded by the cloud provider. Despite being a relatively new concept, multiple cloud simulators were developed. However, they are still in the phase of adapting to the scenarios, objectives and characteristics of the cloud. This paper examines a selected set of the current cloud simulators in terms of vision, features, and architecture. Strong points and limitations are discussed. Moreover, this paper presents a framework for cloud simulator design that can serve as an elaborate design checklist. A discussion of the open research challenges concludes the paper.

1. Introduction

As the cloud computing client base grows, the cloud service providers face the challenge of adapting to the needs of this growth in both the technical and business dimensions. Cloud providers should maintain this gradual enhancement without losing focus on delivering the level of service their clients demand. The dynamic and unpredictable nature of cloud computing adds extra layers of complexity to the providers' tasks. This challenge is magnified by the rise of Big Data concepts that are pushing a new wave of solutions and use case scenarios. Newly developed cloud solutions should be able to handle the scale client data have reached. It is expected of the available infrastructure and software stack to serve the 2.5 Quintillion Bytes (2.3 trillion Gigabytes) that are created every day [1]. IBM estimates that there will be 18.9 billion network connections by 2016. This covers all types of connections especially the ones originated at smartphones carried by any one of the 6 billion expected carriers all over the world.

In the article on Big data facts, myths and possibilities [2], O'Neal presents an example of this change. "With machine to machine (M2M), though, this paradigm changes. Systems generate vast amounts of data independent of human business processes – collecting, analyzing and then deploying this information for use represents a net-new activity for most organizations, in both IT and business operations."

When you add to this mix the number of technical/performance related parameters involved in forming a cloud solution, the task looks increasingly challenging. Providers need efficient tools to support solution design decisions related to deployment models, resource allocation, scheduling, and performance adjustments. Evaluating solutions directly and from the beginning of the solution development process using the real infrastructure is not always practical due to cost factors. The idea of benchmarking using a subset of the infrastructure (like a set of servers with the same configuration and the same topology of the data center for example) would not guarantee a wholesome vision of scalability issues. Scalability is a key player in this scenario and it (along with the cost) constitutes a challenge when using real testbeds. Solution evaluation using analytical methods is rendered infeasible due to the increasing complexity as the scale of the problem grows. Simulation arises here as a tool that – while is not enough alone to handle the whole cloud solution evaluation process – can rather

☆ This work is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC-STPGP 447230) and Ericsson Research.

* Corresponding author.
E-mail addresses: mabusha@uwo.ca (M. Abu Sharkh), ali.kanso@ericsson.com (A. Kanso), ashami2@uwo.ca (A. Shami), peter.ohlen@ericsson.com (P. Öhlén).

http://dx.doi.org/10.1016/j.comnet.2016.06.037
1389-1286/

play major roles before a solution is deployed on real hardware. These roles can be considerably less expensive and with less risk. Cloud simulation involves modeling a real or a proposed cloud system using computer software. It is notably useful when changes to the actual system are either difficult to implement, involve high costs, or are impractical.

Several attempts have been made to develop a competitive cloud simulator. Each cloud simulator differs in vision, the focal points and the resulting features. In this work, we examine the major cloud simulators available to researchers and industry engineers and compare them in terms of the main simulated components, application model, network model, and architecture. Previous attempts to survey cloud simulators can be seen in the literature [3–5] and [86]. In this work, we strive to offer an updated and more comprehensive view of this topic. We present the limitations found in each simulator using an approach that depicts what it is and what it is not. We also aim at illustrating the ground on which the current simulators stand. This helps us to construct a framework for the cloud simulator design process which would ideally cover the industry and research community needs.

Moreover, the paper offers a deep analysis of the open research challenges related to this topic. Challenges covered include realistic user application patterns, cloud deployment and pricing, reliability and high availability, challenges originated outside the data center and Big data considerations.

The coming sections are divided as follows:

Section 2 details the specific roles expected of a cloud simulator. Section 3 is a discussion of the design decisions included in the process of developing a cloud simulator in terms of visions for hardware, applications and network sides. Next, Section 4 traverses a chosen set of common cloud simulators, giving significant attention to a few simulators that contain interesting implementation concepts. This simulator comparison is then consolidated in Section 5. We then proceed to discuss pros and cons of building a new simulator as opposed to extending a current one. Finally, an illustration of some of the research challenges and future work is presented and the paper is concluded.

2. The activities of a cloud simulator

To better understand expectations of the varying simulator users, a clear perspective of the cloud simulator role needs to be defined. This can be done by depicting the potential cloud service planning activities a simulator can serve in a cloud environment. Each activity can be matched with a use case where using a cloud simulator can achieve the required impact. These activities can be summarized in the following:

1. **Define**: Develop greater understanding of process details. Service deployment options, green data center policies and high availability policies are examples of aspects that are affected by several conditions that work simultaneously. Understanding their interaction on the lowest level is a must if efficient resource management is to be achieved.

2. **Pinpoint**: Identify problem areas or bottlenecks in a process that affect execution speed or increase the solution cost.

3. **Maneuver**: Test different “What if?” scenarios to better predict how a real life problem evolves under specific conditions.

4. **Analyze**: Evaluate the effect of system or process changes such as demand, supply, resources specifications, and constraints

5. **Decide**: Compare the impact of alternative policies to determine their points of strength. A simulator is an accurate way to quantify the advantages and disadvantages of newly developed policies.

6. **Scale**: Run real life scenarios on different scales as needed and repeat them as many times as the verification and validation process requires.

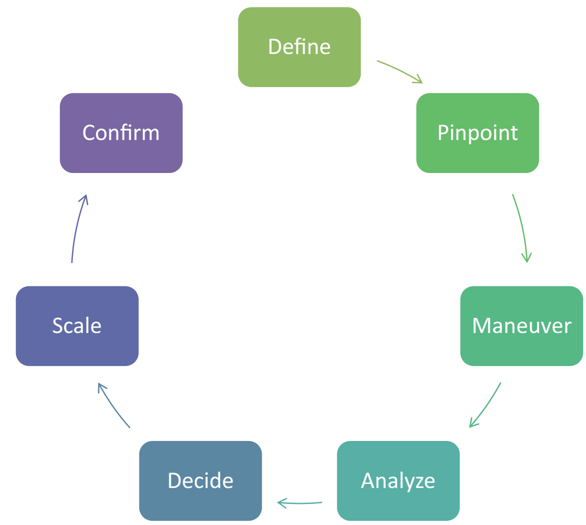

Fig. 1. The cloud service planning activities where a cloud simulator has a role to play.

7. **Confirm**: Use it as a last step to confirm the behavior of proposed solutions when a failed solution has high risk or high cost associations. This includes tuning performance bottlenecks before deployment. The simulator roles are illustrated in Fig. 1.

An example of a use case in which the simulator would play these roles is a scenario where the architects are developing a multi layered resource allocation policy for a cloud data center. This use case constitutes developing a solution that performs the resource allocation for the cloud in order to minimize user request latency to comply with service level agreements (SLAs). The cloud simulator first would help us define the problem by answering questions like: Which elements are involved in this experiment (environment)? Which resources are affected? When do we apply the resource allocation algorithm? Then, it would help us pinpoint potential bottlenecks that would cause this solution to underperform: Is it the VM placement? Is it the resource allocation for the VMs? Is it the chosen topology? Is it the network resources? Then, it would help us maneuver by testing multiple “What if?” scenarios in terms of testing different data sets or use cases or edge cases. Now, we analyze the effect of every factor on the resource allocation process. For example, we might notice that the algorithm performs with sparse heavy requests better than numerous light requests or that it performs better with high-CPU VMs. Based on that, we can gauge and decide which factors have more importance and give them more weight in our scheduling policy. Next step is to scale by testing our scheduling algorithm for a very large data center, large internal network, heavy communication and a large number of requests and see where/when it breaks. Finally, it can be used for final tuning before the algorithm is tested on a real test bed to schedule real requests.

The potential users taking advantage of these simulator roles include:

1. Cloud providers and solution architects: Naturally, cloud providers represent the main user as they would use this to develop, evaluate and improve their solutions.

2. Cloud clients: Typical cloud client list include large companies that have the capabilities to run their own clouds. A cloud simulator would be beneficial to compare different providers, evaluate currently deployed solutions or for studies of the client workload and to support decisions regarding private vs. public clouds.

3. The research community: simulators are a critical preliminary step for researchers who are developing new cloud technology before testing on a real setup.