

Chapter 86

Performance Degradation Detection of Virtual Machines Via Passive Measurement and Machine Learning

Toshiaki Hayashi

Hokuriku Computer Service Co., Ltd., Japan

Satoru Ohta

Toyama Prefectural University, Japan

ABSTRACT

Virtualization is commonly used for efficient operation of servers in datacenters. The autonomic management of virtual machines enhances the advantages of virtualization. Therefore, for the development of such management, it is important to establish a method to accurately detect the performance degradation in virtual machines. This paper proposes a method that detects degradation via passive measurement of traffic exchanged by virtual machines. Using passive traffic measurement is advantageous because it is robust against heavy loads, non-intrusive to the managed machines, and independent of hardware/software platforms. From the measured traffic metrics, performance state is determined by a machine learning technique that algorithmically determines the complex relationships between traffic metrics and performance degradation from training data. The feasibility and effectiveness of the proposed method are confirmed experimentally.

INTRODUCTION

Virtualization (Sahoo, Mohapatra, & Lath, 2010) is an indispensable technique used for the efficient operation of servers in datacenters and for providing cloud services. In this technique, multiple virtual machines run on a physical host computer while sharing the host's resources. Virtual machines must be suitably managed to provide adequate quality of service that conforms to the service level agreement (SLA). However, if excessive load is experienced by a virtual machine and the resources of its physical

DOI: 10.4018/978-1-5225-1759-7.ch086

Performance Degradation Detection of Virtual Machines

host machine are exhausted, the SLA might be violated due to performance degradation. If this happens, the performance will be restored by moving the deteriorated virtual machine to another host by using a live migration technique (Clark et al., 2005).

To ensure good quality of service to clients and users with the least operational expenditure, autonomic management of virtual machines is required. That is, the mapping between virtual machines and physical hosts should be dynamically rearranged depending on the load offered to each machine. This type of management minimizes the number of operating physical hosts while providing adequate quality of service. To this end, the management system must (1) detect performance degradation in virtual machines, (2) determine the virtual machine to be migrated to an alternative host, and (3) execute a live migration. This paper focuses on detecting performance degradation in virtual machines.

Autonomic management of virtual machines has been studied in (Bobroff, Kochut, & Beaty, 2007; Xu & Fortes, 2011; Andreolini, Casolari, Colajanni, & Messori, 2010). The method proposed by (Bobroff, et al., 2007) determines the virtual machines that must be migrated by analyzing changes in CPU utilization. In the method proposed by (Xu & Fortes, 2011), information on resource consumption is obtained for both virtual machines and physical hosts. The resource utilization metrics used in their method include CPU utilization, disk I/O, network utilization, and temperature. The method presented by (Andreolini, et al., 2010) determines the virtual machine to be migrated by considering the state change characteristic of a single load metric, such as CPU utilization. In all these methods, performance state is determined using metrics obtained through the operating system (OS). Thus, degradation detection is intrusive. In addition, the software that gathers the resource consumption information depends on the machine platform. It is also unclear if the employed metrics accurately represent the degradation experienced by users. For example, performance might degrade at very low CPU utilization for some service content and at high CPU utilization for other content. It is not certain that the existing method can successfully handle this difficulty.

Several technical problems arise in detecting performance degradation in a virtual machine. First, performance is degraded by the exhaustion of different resources, such as the CPU, disks, and network interfaces, depending on differences in the service content requested by users. The utilization of packet queuing buffers or setting of the server program parameters also affects performance. Thus, it is necessary to monitor multiple metrics that reflect broader utilization of these resources. Moreover, the relationship between resource utilization and quality of service experienced by users is not necessarily clear. Conflicts among virtual machines on the same physical host complicate the problem further. Obtaining information on resource utilization through the virtual machine or physical host OS is also problematic. That is, if the performance of a machine is greatly degraded, it becomes difficult to successfully extract resource information through the OS. Compatibility is another problem; different measurement software must be developed for different OSs. This increases development time and cost. In addition, the measurement of resource utilization through the OS consumes computational resources of both the virtual machines and their physical hosts. This might lead to a decrease in machine capacity.

To avoid these problems, performance degradation in virtual machines should be detected in a manner that is non-intrusive to both virtual machines and physical hosts. The metrics measured for degradation detection should include information that reflects the consumption of various performance related resources. It is also necessary to clarify a method to detect degradation signs among these metrics.

The goal of this paper is to develop virtual machine management for use when virtual machines are operated as servers that provide network services. In this case, the paper investigates a method in which performance degradation is detected through several metrics related to traffic sent and received by virtual

17 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:
www.igi-global.com/chapter/performance-degradation-detection-of-virtual-machines-via-passive-measurement-and-machine-learning/173416

Related Content

Integration of Different Analytical Concepts on Multimedia Contents in Service of Intelligent Knowledge Extraction

Goran Klepac (2017). *Artificial Intelligence: Concepts, Methodologies, Tools, and Applications* (pp. 2493-2522).

www.irma-international.org/chapter/integration-of-different-analytical-concepts-on-multimedia-contents-in-service-of-intelligent-knowledge-extraction/173432

Using a Knapsack Model to Optimize Continuous Building of a Hybrid Intelligent Tutoring System: Application to Information Technology Professionals

Maha Khemaja (2018). *Intelligent Systems: Concepts, Methodologies, Tools, and Applications* (pp. 1488-1506).

www.irma-international.org/chapter/using-a-knapsack-model-to-optimize-continuous-building-of-a-hybrid-intelligent-tutoring-system/205844

A Novel Fuzzy Associative Memory Architecture for Stock Market Prediction and Trading

Chai Quek, Zaiyi Guo and Douglas L. Maskell (2011). *International Journal of Fuzzy System Applications* (pp. 61-78).

www.irma-international.org/article/novel-fuzzy-associative-memory-architecture/52055

Transformative Power of Artificial Intelligence in Decision-Making, Automation, and Customer Engagement

R. Nalini (2024). *Complex AI Dynamics and Interactions in Management* (pp. 189-208).

www.irma-international.org/chapter/transformative-power-of-artificial-intelligence-in-decision-making-automation-and-customer-engagement/339748

Symbol Grounding Problem

Angelo Loula and João Queiroz (2009). *Encyclopedia of Artificial Intelligence* (pp. 1543-1548).

www.irma-international.org/chapter/symbol-grounding-problem/10443