Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 110 (2017) 510–515

Procedia Computer Science

www.elsevier.com/locate/procedia

2nd International Workshop on Applications of Software-Defined Networking in Cloud Computing (SDNCC)

Automated Anomaly Detection in Virtualized Services Using Deep Packet Inspection

Marcel Wallschläger[a], Anton Gulenko[a], Florian Schmidt[a], Odej Kao[a], Feng Liu[b]

[a]*Technische Universität Berlin (TU Berlin, Complex and Distributed IT Systems (CIT)), 10587 Berlin, Germany*
[b]*Huawei European Research Center, Huawei Technologies Co., Ltd., 80992 Munich, Germany*

Abstract

Virtualization technologies have proven to be important drivers for the fast and cost-efficient development and deployment of services. While the benefits are tremendous, there are many challenges to be faced when developing or porting services to virtualized infrastructure. Especially critical applications like Virtualized Network Functions must meet high requirements in terms of reliability and resilience. An important tool when meeting such requirements is detecting anomalous system components and recovering the anomaly before it turns into a fault and subsequently into a failure visible to the client.

Anomaly detection for virtualized services relies on collecting system metrics that represent the normal operation state of every component and allow the usage of machine learning algorithms to automatically build models representing such state. This paper presents an approach for collecting service-layer metrics while treating services as black-boxes. This allows service providers to implement anomaly detection on the application layer without the need to modify third-party software. Deep Packet Inspection is used to analyse the traffic of virtual machines on the hypervisor layer, producing both generic and protocol-specific communication metrics. An evaluation shows that the resulting metrics represent the normal operation state of an example Virtualized Network Function and are therefore a valuable contribution to automatic anomaly detection in virtualized services.

© 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the Conference Program Chairs.

Keywords: Cloud, Network, Deep Packet Inspection, Network Function Virtualization, Anomaly Detection

1. Introduction

As the number of services running on virtualized infrastructure continues to increase, even critical applications from the telecommunication sector are ported from traditional appliances onto cloud deployments.

Driven by high customer expectations, such critical services have an especially high demand for reliability and continuous service delivery.

* Marcel Wallschläger. Tel.: +49-30-314-78592 ; fax: +0-000-000-0000 .
E-mail address: marcel.wallschlaeger@tu-berlin.de

1877-0509 © 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the Conference Program Chairs.
10.1016/j.procs.2017.06.137

Traditional reactive fault management is not sufficient to guarantee a constantly high service availability. Cost effective private and public clouds rely on commodity hardware which cannot guarantee failover latencies short enough to hide faults from clients.

Since anomalies often precede faults, anomaly detection mechanisms provide an early warning system that enables proactive fault management[1]. Anomaly detection can operate across all layers of the cloud system by collecting various operative time series metrics. Unsupervised machine learning techniques analyse the collected data to detect abnormal patterns and outliers.

The anomaly detection accuracy mainly depends on the input data used to build the underlying machine learning models. In the context of virtualized services, many indicators for normal operation can be found in resource usage data like the utilization of CPU, memory, network and disk. This data is available on all system layers and does not require specific knowledge about the monitored service. Treating services as black boxes makes this anomaly detection approach easily usable in arbitrary productive environments. Furthermore, it allows Infrastructure-as-a-Service (IaaS) providers to offer an anomaly detection service to customers while maintaining minimal interference with customer machines.

However, mere resource usage fails to accurately reflect the communication patterns of the observed services and therefore does not cover all potential anomalies. This paper presents a mechanism for collecting service communication metrics while still remaining largely service agnostic, as long as the services rely on standardized application layer protocols. Deep Packet Inspection (DPI) on the hypervisor level allows real-time analysis of the inter-service communication with low interference. Services do not need to be customized or extended to obtain this information, but knowledge about the used protocols can reduce the performance overhead of packet inspection. The second contribution of this paper is an evaluation of the presented data collection mechanism. We show that a number of simulated anomalies manifest themselves in the resulting metrics.

The remainder of the paper is organized as follows. The following section presents related research in the field of deep package inspection and anomaly detection. Section 3 describes the presented approach in detail, while section 4 presents an evaluation thereof.

2. Related Work

Much related work has been conducted in the field of deep package inspection and automated protocol analysis.

M. Danelutto et al. show in their work[2] how package inspection is possible on commodity server hardware using a skeleton-based parallel programming library targeting efficient streaming on multi-core architectures. Using their framework the authors show that package inspection for packets with 60 byte payloads can be performed with a commodity Intel 10 Gbit network card.

Anat Bremler-Barr et al. identify DPI as a common task for middleboxes that inspect application layer packets many times during their route from sender to the final destination[3]. Since most middleboxes perform similar analyses to implement traffic control and QoS measurements, the authors propose a DPI-as-a-service infrastructure to reduce the total overhead of repeated packet inspection. The proposed infrastructure can lead to improved performance, scalability and robustness, showing that DPI has become mature and fast enough to be executed even as a standalone service.

Further research shows that the application layer protocol used by packets can be identified in real-time using DPI[4,5,6].

3. Approach

An exemplary infrastructure illustrates the proposed data collection approach and follows a typical cloud deployment used both by IaaS providers and in private clouds hosting Virtualized Network Function services. The services run on virtualized hardware and network resources provided by the cloud operating system OpenStack[7] and complementary Software Defined Networking (SDN) components. Open vSwitch[8] manages the virtual network connecting all VMs.

Figure 1 shows the basic structure of the network architecture inside a physical compute node managed by OpenStack. It illustrates two virtual machines within one hypervisor connected to various Open vSwitch bridges. A dedicated monitoring agent implements the data collection by sniffing packets on the tap interface of every relevant

Download English Version:

https://daneshyari.com/en/article/4960855

Download Persian Version:

https://daneshyari.com/article/4960855

Daneshyari.com