Computer Communications 91–92 (2016) 44–61

Contents lists available at ScienceDirect

Computer Communications

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

PredCloud: Providing predictable network performance in large-scale OpenFlow-enabled cloud platforms through trust-based allocation of resources

Daniel S. Marcon, Miguel C. Neves, Rodrigo R. Oliveira, Luciano P. Gaspary, Marinho P. Barcellos*

Institute of Informatics, Federal University of Rio Grande do Sul, Av. Bento Gonçalves, 9500 – 91.501–970 – Porto Alegre, RS, Postal Code 15 064, Brazil

ARTICLE INFO

Article history:
Received 29 February 2016
Revised 19 July 2016
Accepted 20 July 2016
Available online 21 July 2016

Keywords:
Performance interference
Resource allocation
Intra-cloud network sharing

ABSTRACT

Cloud computing allows tenants to run a wide range of applications without any upfront capital investment. However, providers lack mechanisms to provide fair and predictable bandwidth sharing among allocated applications, enabling selfish and malicious tenants to cause performance interference in the network (and denial of service in an extreme case). Such *interference* results in poor and unpredictable network performance for well-behaved applications. Recent research has proposed techniques that (*i*) cannot protect tenants against interference; (*ii*) result in under utilization of resources; or (*iii*) add substantial management overhead. In this paper, we describe a resource allocation strategy that aims at providing predictable network performance (i.e., minimizing performance interference) with bandwidth guarantees for tenant applications, while maintaining high network utilization and low management overhead. These benefits are achieved by grouping applications from mutually trusting users into logically isolated domains (virtual infrastructures - VIs) with bandwidth guarantees, while also considering the amount of traffic generated by applications. Despite the benefits, grouping may lead to fragmentation (i.e., available resources are dispersed among VIs and some requests may be unnecessarily declined). Therefore, we also study the associated trade-off (grouping to increase isolation versus resource fragmentation). To illustrate the feasibility of grouping applications inside VIs, we develop PredCloud, a system that implements the proposed strategy on SDN/OpenFlow-enabled networks. Through an extensive evaluation, we show that PredCloud significantly reduces performance interference and application exposure to attacks, while maintaining low resource fragmentation. Furthermore, provider revenue can be increased by efficiently managing and charging network resources.

1. Introduction

Cloud Computing has become the platform of choice for the delivery and consumption of IT resources. It offers several advantages, such as pay-as-you-go pricing model, on-demand self-service, broad network access and rapid elasticity [1]. In this model, providers avoid allocating physically isolated resources for each tenant. Instead, they implement cloud data centers as highly multiplexed shared environments, with different applications coexisting on the same set of physical resources [2]. Therefore, they can increase resource utilization, reduce operational costs and, thus, achieve economies of scale.

However, the intra-cloud network is typically shared in a best-effort manner, without a robust mechanism to provide fair and predictable bandwidth sharing among applications [3,4]. This enables users to perform insider attacks [5,6]. We distinguish two types of attacks: (*a*) selfish, characterized by the intentional consumption of an unfair share of the network (causing performance interference [7,8]); and (*b*) malicious, characterized by denial of service (DoS) in the network, an extreme case of performance interference.[1]

* Corresponding author.
E-mail addresses: daniel.stefani@inf.ufrgs.br (D.S. Marcon), mcneves@inf.ufrgs.br (M.C. Neves), ruas.oliveira@inf.ufrgs.br (R.R. Oliveira), paschoal@inf.ufrgs.br (L.P. Gaspary), marinho@inf.ufrgs.br (M.P. Barcellos).

[1] Note that the lack of robust network sharing mechanisms allows tenants to perform other kinds of attacks (e.g., extraction of confidential information) [9]. In this paper, we focus on performance interference and DoS and defer a detailed study of other attacks to future work.

http://dx.doi.org/10.1016/j.comcom.2016.07.010
0140-3664/© 2016 Elsevier B.V. All rights reserved.

The lack of network guarantees hurts both well-behaved tenants[2] and providers. Tenants have unpredictable costs, because performance interference and DoS result in poor and unpredictable network performance for applications [10]. Providers, in turn, lose revenue, because performance interference reduces datacenter throughput and DoS affects network availability [11].

In this paper, we propose a resource allocation strategy for Infrastructure as a Service (IaaS) providers. Unlike previous work [7,8,10,12,13], the strategy is based on grouping applications in virtual infrastructures[3] (VIs) with bandwidth guarantees, in order to maintain high resource utilization and low management overhead (two common goals of providers). Grouping applications into VIs has both benefits and drawbacks, discussed as follows.

Benefits of grouping. There are three benefits. The first one is related to security: grouping can provide isolation among applications from mutually untrusted tenants, reducing application exposure to attacks in the network. In other words, applications from mutually untrusted tenants are allocated in different VIs and do not compete for bandwidth (i.e., are isolated). Consequently, the system becomes more resilient against tenants that would try to cause disruption in the network or use a disproportionate share of resources (seeking to reduce costs at the expense of other applications). The second benefit regards performance, since grouping allows cloud platforms to provide performance isolation among applications allocated in different VIs (reducing performance interference and, thus, improving network predictability and guarantees). The third benefit is related to management overhead in the network: unlike related work that manages network resources at flow level [7,14] (there are millions of flows per second in the network [15]) and at application level [13,16] (thousands of applications), our strategy only needs to manage network resources at the virtual infrastructure level (a few number of VIs).

Drawbacks of grouping. In contrast, there are two shortcomings. First, the number of groups created is pragmatically limited by the overhead of the virtualization technology. Second, groups may lead to resource fragmentation among VIs when allocating requests (that is, available resources are dispersed among VIs and, in some occasions, no VI alone may have the necessary amount of free resources to hold an incoming application). Therefore, we study different strategies to group tenants based on their mutual trust and on the network requirements (bandwidth)[4] of their applications.

PredCloud. To show the benefits and drawbacks of grouping applications inside VIs, we introduce PredCloud, a system that implements the proposed strategy on top of SDN/OpenFlow-enabled networks [18,19]. OpenFlow is the most accepted implementation of SDN by both the industry and academia and is being widely deployed in current network devices [20]. It enables the separation between the control and data planes and allows providers to dynamically configure and manage the network [21], presenting an efficient way of allocating and ensuring bandwidth for applications. Using SDN/OpenFlow, PredCloud embeds VIs onto the cloud substrate[5] according to the demands (needed resources and performance requirements), maps applications inside VIs (in accordance with trust relationships between tenants, in order to minimize application exposure to attacks) and ensures bandwidth guarantees for both intra- and inter-application communication (enabling applications to achieve predictable network performance).

Application-level granularity. Note that, unlike other approaches such as Rimal and El-Refaey [22], PredCloud allows the granularity of application instead of tenant because of two reasons. First, per-application granularity is preferred to address performance interference: applications of the same tenant may have different network requirements and could impact on the network performance of one another if not isolated [23]. Second, tenants may desire different levels of security for different applications (e.g., one that manipulates confidential information and one that requires only publicly disclosed information).

Contributions. In comparison to our previous work [24], here we present a substantially improved version of the proposed strategy. Overall, the major contributions of this paper are:

- Two novel algorithms: one to efficiently embed VIs onto the cloud substrate (as opposed to a time-consuming integer linear programming algorithm), which is needed as public cloud platforms present high rates of request arrival and departure; and other that leverages dynamic programming to find optimal placements for applications, without requiring complex processing (rather than a simple constructive heuristic);
- PredCloud, a system that takes advantage of Software-Defined Networking (SDN) to implement the proposed strategy to provide increased security and predictable network performance with bandwidth guarantees;
- An extensive evaluation of PredCloud, considering several factors to analyze its benefits, overheads and technical feasibility. We show that PredCloud minimizes performance interference in the network, reduces application exposure to attacks, presents low resource fragmentation among VIs and maximizes provider revenue.

The remainder of this paper is outlined as follows. Section 2 provides the background on the current resource allocation scheme employed by public cloud providers, on legacy algorithms used for network design, on legacy technologies used to implement network virtualization and on attacks performed in the intra-cloud network. Sections 3 and 4 define the basis in which our approach was developed, that is, the threat model considered and basic formulations related to the problem, which are later used throughout the paper. Section 5 describes our resource allocation strategy, while Section 6 introduces PredCloud, a system that implements the proposed strategy on top of SDN/OpenFlow-enabled networks. Section 7 presents the results of an extensive evaluation of PredCloud, quantifying its benefits, overheads and technical feasibility in large clouds. Section 8 discusses the generality and limitations of PredCloud, and Section 9 examines related work and differentiates it from PredCloud. Finally, Section 10 concludes the paper.

2. Background

In this section, we discuss (*i*) the current allocation strategy employed by public cloud providers; (*ii*) legacy algorithms used for network design; (*iii*) legacy technologies used to implement network virtualization; and (*iv*) attacks performed in the intra-cloud network.

Allocation strategy employed by public cloud providers. Providers typically use resource allocation strategies focused on computing resources (i.e., they assume that the network is shared in a best-effort manner among all allocated applications) [25]. Such strategies use round-robin across servers or across racks, taking into account only the amount of available resources in servers.

[2] Well-behaved tenants are users that do not launch attacks (performance interference and DoS) in the network.

[3] The term virtual infrastructure is used to represent a set of virtual machines as well as the virtual network interconnecting them. This concept is formally defined in Section 4.2.

[4] There are other connectivity requirements defined by RFC 7297 [17]. However, in this paper, we focus specifically on bandwidth because the challenge being addressed (performance interference) happens due to the lack of bandwidth guarantees in cloud datacenters.

[5] We use the terms “physical infrastructure” and “physical substrate” interchangeably, to refer to the set of physical resources of the cloud platform.