Future Generation Computer Systems 28 (2012) 85–93

Contents lists available at SciVerse ScienceDirect

Future Generation Computer Systems

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

Virtual machine provisioning through satellite communications in federated Cloud environments

Antonio Celesti *, Maria Fazio, Massimo Villari, Antonio Puliafito

Department of Mathematics, Faculty of Engineering, University of Messina, Contrada di Dio, S. Agata, 98166, Messina, Italy

ARTICLE INFO

Article history:
Received 17 February 2011
Received in revised form
11 May 2011
Accepted 28 May 2011
Available online 12 June 2011

Keywords:
Cloud computing
Federation
Satellite communications
Service delivery
Distributed Cloud service

ABSTRACT

Cloud federation offers plenty of new services and business opportunities. However, many advanced services cannot be implemented in the real Cloud market due to several issues that have not been overcome yet. One of these concerns is the transfer of huge amount of data among federated Clouds. This paper aims to overcome such a limitation proposing an approach based on satellite communications. By comparing performance in data delivery on the Internet and satellite systems, it is evident that satellite technologies are enough ripe to be competitive against systems with a wired infrastructure. Thus, we propose to make use of satellite transmission to implement fast delivery of huge amount of data. Through the discussion of a use case, where a WEB TV company offers a streaming service, we show how to practically apply the proposed strategy in a real scenario, specifying the involvement of Cloud providers, Cloud users, satellite companies and end-user clients.

1. Introduction

Cloud computing offers new business opportunities for both service providers and their clients (e.g. organizations, enterprises, and end users), by means of an architecture for delivering Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). A Cloud encloses the IaaS, PaaS, and/or SaaS inside its own virtualization infrastructure, in order to carry out an abstraction from its underlying physical assets. Typically, the virtualization of a service implies the aggregation of several proprietary processes collected in a virtual environment, called Virtual Machine (VM) [1].

Often, Clouds are also spread over distributed virtualization infrastructure covering larger geographical areas (for example, let us think about RESERVOIR (an European project facing the Cloud Computing IaaS topic [2]), Amazon [3], and Azure [4]). In addition, the perspective of Cloud federation [5,6], where Cloud providers use virtualization infrastructures of other federated Clouds, opens toward new scenarios in which more and more types of new services can be supplied. In fact, Clouds exploiting distributed virtualization infrastructures are able to provide new types of "Distributed IaaS, PaaS, and SaaS".

However, the potential growth of Cloud environments in terms of the extension of the covered geographical area, business opportunities, and types of services can be delayed by the limits of the communication infrastructure, which is provided by the Internet. Therefore, in Cloud infrastructures distributed over a Wide Area Network (WAN), the provisioning and deployment of services raise several issues due to the latency of the Internet. For example, provisioning implies the displacement of several VMs composing a distributed Cloud service, from the Cloud that receive the service allocation request to the other federated Clouds, where the service must be deployed.

Since we believe in Cloud computing as the future service provisioning model, we propose a new strategy to improve the quality of service provisioning, based on the employment of satellite technologies for long range and fast data delivery. Satellite systems are becoming very popular for high-speed data transmission [7]. As confirmed by the effort devoted from many business companies to propose ADSL Broadband Services and their investments to launch new satellites into orbit. In our strategy, satellite links can be used to implement fast delivery of huge amount of data, by using opportunistic burst transmission, such as SMS transmission in cellular networks. In fact, satellite systems are the best way for broadcasting media content, without additional costs for cabling and regardless geographical position of transmitters and receivers. In comparison with wired networks, satellite systems allow to perform multi-casting in a very easy way, providing a high improvement in throughput while reducing the effective load on nodes in the network.

In this paper, we analyze the potentiality of federated Cloud architectures [2] focusing on distributed Cloud services and underlying which new market perspective they can open. We

* Corresponding author.
E-mail addresses: acelesti@unime.it (A. Celesti), mfazio@unime.it (M. Fazio), mvillari@unime.it (M. Villari), apuliafito@unime.it (A. Puliafito).

0167-739X/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2011.05.021

discuss the limits of the provisioning of distributed Cloud services using the Internet, proposing an alternative and more efficient Cloud service provisioning strategy based on satellite communications. In order, to highlight the advantages of such a new scenario, we consider the use case of a WEB TV company, which uses a distributed IaaS in a Cloud federation exploiting the satellite service.

The paper is organized as follows. Section 2 provides an overview regarding the delivery of VMs in federated Cloud environments. In Section 3, we introduce the new idea of distributed Cloud services available in federated Cloud environments highlighting the limits of the delivery due to the Internet infrastructure. Then, we focus on the use case of a WEB TV company performing live streaming, which uses a distributed IaaS. In Section 4, we give a brief overview of the modern satellite systems and presents a comparison between performances of satellite and wired links in terms of coverage and data rate. Then, in Section 5, we discuss a new approach to deliver services in federated Clouds based on satellite communications. In Section 6, we show how satellite systems are useful with reference to the use case of a WEB TV company. Conclusion and highlights to future works are presented in Section 7.

2. Related works and background

Over the years, there has been an ever increasing request of services provisioned through the Internet. The client/server model was the first adopted way for the delivery of multimedia contents. But with the growth of both network complexity and number of clients, this model has demonstrated several weaknesses. Thus in the 90s, with Napster, OpenNAP and IRC serving channels appeared the first peer-to-peer (P2P) file sharing systems, where each peer was able to supply and consume resources in contrast with the traditional client/server model. At the same time, in order to both reduce the exchanged traffic of P2P networks and enhance the Quality of Service (QoS), the NEXUS International Broadcasting Association developed the first Content Delivery Network (CDN). A CDN is a system of computers containing copies of data (e.g., media files, softwares, documents, etc.), placed in several sites in a network in order to maximize the bandwidth utilization. A CDN using protocols just like Internet Content Adaptation Protocol (ICAP) and Open Pluggable Edge Services (OPES) can use several algorithms (e.g., Global Server Load Balancing, DNS-based request routing, Dynamic metafile generation, HTML rewriting, and anycasting) in order to enhance the network throughput. Examples of free and commercial CDNs include Coral Content Distribution Network, FreeCast, PeerCast, Akamai, Bitgravity and CacheFly. Nowadays Cloud computing brings new scenarios of content delivery especially in federated environments. Considering Cloud federation, new terms such as InterCloud [8] or Cross-Cloud [9] are becoming very debated in the scientific community. As discussed in [5], Cloud federation brings many business advantages for the enhancement of providers' profit. In fact it has been proposed as a new paradigm allowing providers to avoid the limitation of owning only a restricted amount of resources. In [10], highlighting the limits of several open source Cloud platforms, a three-phase model is proposed for the establishment of a federation among different Cloud architectures, where Clouds are able to use the virtualization infrastructure of other federated Clouds in order to increase their computational and storage capabilities. Cloud federation raises several issues and one of these is represented by the delivery of large amount of data over a WAN. However, on this topic, there have not been many available related works yet. Regarding the provisioning and delivery of VMs among federated Clouds, a relocatable storage access mechanism capable of rapidly relocating VMs, with a limited impact on I/O performance of the migrant VMs, is presented in [11,12]. Instead a model able to reduce the cost of the VM migration by means of a disk image composition technique is presented in [13]. In particular, the authors demonstrate how their approach is valid for scenarios composed of hundreds of federated Clouds. As discussed in Section 3.1, we think that in a federated Cloud environment, with a high level of flexibility, approaches based on P2P, CDNs, or relocatable access storage over the Internet are not suitable. Therefore, in our opinion we need alternatives and more powerful approaches for the composition, delivery and upgrade of Cloud services.

3. Distributed Cloud service provisioning in federated Cloud environments

Cloud providers rely their computational capabilities on the concept of "virtualization". Virtualization technologies aim to hide the underlying infrastructure by introducing a logical layer between the physical infrastructure itself and the computational processes. Through Virtualization Machine Monitors (e.g., [14], KVM [15], VMware [16]), commonly known as "hypervisors", each Cloud is able to control and emulate several processing environments (i.e. VMs), each running its own "guest" software, typically an operating system. Commonly, a Cloud uses its own virtualization infrastructure (i.e., a set of servers running hypervisors) to instantiate and aggregate one or more VMs and, thus, deliver to its clients IaaS, PaaS, and/or SaaS. Fig. 1 depicts an example of IaaS composed of 5 VMs hosted within a Cloud virtualization infrastructure composed of 16 servers. In our opinion, according to the size of the Clouds' virtualization infrastructure we can distinguish between small, medium, and large Clouds. Small and medium Clouds are held by small and medium size companies, instead large Clouds are held by large size companies (e.g., Amazon Europe or RESERVOIR). The latter are able to provide new types of Cloud services named "Distributed IaaS, PaaS, and SaaS" (generalizing "D*aaS"). With D*aaS, we indicate *"a distributed Cloud service composed of a set of VMs spread over a wide geographical area, orchestrated in order to achieve a target purpose, and which is provided on demand to a client to meet his business needs"*. Also to give an opportunity to both small and medium Clouds to offer D*aaS, a federation of Clouds can be constituted. In a federation, small and medium Clouds can enlarge their virtualization capabilities using the virtualization infrastructures of other federated Clouds (small, medium, or large) for a given business purpose, becoming as competitive as large Clouds. This business model permits on one hand to elastically increase the virtualization capabilities of Clouds and, on the other hand to enable Clouds to rent their computational and storage capabilities when their virtualization infrastructures are partially or totally unused. A Cloud might choose to establish a federation relationship with other Clouds for many reasons. For example: (1) it needs extra resources because the capabilities of its virtualization infrastructure are run out; (2) it needs to have some resources placed in a given geographical location; (3) it wants to save money reducing the electricity consumption allocating services in other Clouds and so on. Fig. 2 depicts an example of Cloud federation where Cloud D provides a DIaaS to a client, which uses the VMs hosted in other federated. Federated Clouds cover wide geographical areas, so allowing D*aaS for massive distribution.

The deployment of VMs inside the federation can be performed in two possible ways, which are as follows.

- *One-time* deployment: it is characterized by an a priory arrangement of resources. If the allocation of VMs does not match real needs of service users, resources can remain underused or even totally unused.