

Network-Aware Service Placement in a Distributed Cloud Environment

Moritz Steiner, Bob Gaglianella, Vijay Gurbani,
Volker Hilt, W. D. Roome, Michael Scharf, and Thomas Voith
Bell Labs, Alcatel-Lucent
{first.last}@alcatel-lucent.com

ABSTRACT

We consider a system of compute and storage resources geographically distributed over a large number of locations connected via a wide-area network. By distributing the resources, latency to users can be decreased, bandwidth costs reduced and availability increased. The challenge is to distribute services with varying characteristics among the data centers optimally. Some services are very latency sensitive, others need vast amounts of storage, and yet others are computationally complex but do not require hard deadlines on execution. We propose efficient algorithms for the placement of services to get the maximum benefit from a distributed cloud systems. The algorithms need input on the status of the network, compute resources and data resources, which are matched to application requirements.

This demonstration shows how a network-aware cloud can combine all three resource types - computation, storage, and network connectivity - in distributed cloud environments. Our dynamic service placement algorithm monitors the network and data center resources in real-time. Our prototype uses the information gathered to place or migrate services to provide the best user experience for a service.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Network operating systems*

Keywords

Cloud, service placement

1. INTRODUCTION

Cloud computing has seen sizeable growth in the past few years and analysts predict the trend is going to accelerate [3]. Today, public cloud providers operate a small number of large data centers, offering compute and storage resources. The cloud operators do not own nor do they operate the network that connects their data centers to the public Internet or to their customer's sites. Network operators, on the other hand, operate the network but are not aware of the cloud services using that network. We argue that many new opportunities arise by combining the data center and network operations and by adding compute and storage resources at multiple locations within the network.

One such opportunity is to run services closer to users, thereby running a service at many locations in parallel. This facilitates

offering of latency-sensitive services from the cloud, and provide performance as good – or better – than users currently receive from their local computers. Examples include remote storage, virtual desktops, or action games. Integrated network and cloud providers can offer services in high quality to end users, since they control all resources involved in the delivery of the service. Finally, bandwidth costs can be reduced by running services on data-centers close to the users.

2. RESEARCH CHALLENGES

To obtain the maximum benefit from a distributed cloud system, we need efficient service-placement algorithms that can accommodate different, sometimes contradicting, goals [1]. A service provider may, for example, want to minimize communication, compute and storage costs, maximize link and data-center utilization and guarantee the user's quality of experience. This is a nontrivial, multi-dimensional optimization problem.

The placement and resource allocation algorithms need input from the data centers and from the network describing available resources and some normalized notion of the cost of accessing that resource. A challenge is to summarize and abstract this data to a level that is meaningful for solving the optimization and without revealing too many details about the communication network and data centers. Besides information about the resources, the algorithms need to know the application's profile: is it latency sensitive? does it require a large amount of storage? how do the different components of an application communicate? which components are user-facing? are there any legal restrictions involved in executing services for a given user outside a defined country, etc.? There are, of course, existing research and standardization activities that can be leveraged, but regardless, some extensions are needed for distributed cloud environments.

A further challenge is to offer an environment that enables service developers to rapidly deploy new services that are distributed, elastic, close to the users, and resilient. This requires a management system optimized for distributed clouds.

3. DEMONSTRATION

We have implemented a solution for network-aware service placement in a distributed cloud, which will be shown during the demonstration. The solution has several components. First, we utilize CloudStack [2], an open source cloud management system. By using such a system, we show that our solution does not depend on proprietary protocols. Second, we implemented our own Web-based frontend "dashboard" that communicates with the other components, and provides a central interface for monitoring and configuring the system. Finally, we use the ALTO protocol [4], developed within the Internet Engineering Task Force (IETF) as a standard in-

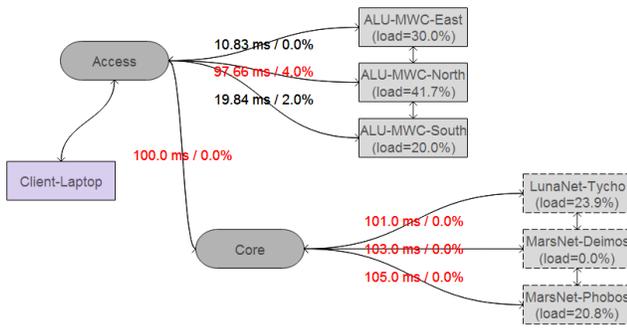


Figure 1: The demonstration network topology, including live measurement data.

interface to obtain information about the communications network. We developed our own ALTO server that collects network topology data and provides an abstract view of that topology and costs. The networking information used by the ALTO server is gathered from live measurements in the demonstration network and includes dynamic loss and delay information. The linux commands “tc” and “netem”, included in current distributions, are used to introduce latency and loss into the system.

The overall function of the system is as follows: The cloud dashboard queries the ALTO server for the network metrics and the cloud management for the data center load. Using this input and the user location, the dashboard runs a placement decision algorithm, which computes a score for each data center and recommends the data center with the best overall score.

The demonstration utilizes four laptops and an Ethernet switch to simulate multiple data centers connected to and controlled by a cloud management system. Each of three laptops form a mini data center. The fourth laptop hosts the cloud management software, the dashboard, the measurement logic and the ALTO server. An additional laptop runs the remote desktop client, and another laptop or tablet displays the network topology and the dashboard of the cloud operations center.

4. SCENARIOS

Our demonstration uses a virtual desktop as a sample application. Access to a remote desktop is a promising use case for outsourcing enterprise infrastructure. It uses the standard RDP protocol [5] to connect the user to the host of the virtual desktop. Because it is very latency sensitive, the virtual desktop service highlights the advantages of having many distributed data centers over having one or a few data centers.

The demonstration starts by running a high definition video on the laptop serving as a client. Note that the virtual desktop, running in the remote data-center, renders the video and writes the video bit-stream to the virtual screen. RDP sends those screen updates over the network to the client computer, and the client displays them on the user’s physical screen. Initially the video plays smoothly, and the audio is synchronized with the video. This is the experience the user expects.

Then we degrade the user’s experience by introducing latency and loss between the client computer and the data-center hosting the virtual desktop. In the real world, this could happen on an overloaded link, or when the user disconnects from the network and resumes after reconnecting from a different location. The video becomes choppy and out of sync with the audio. In the demonstration, we allow the user to “look behind the curtain” to see the cause of

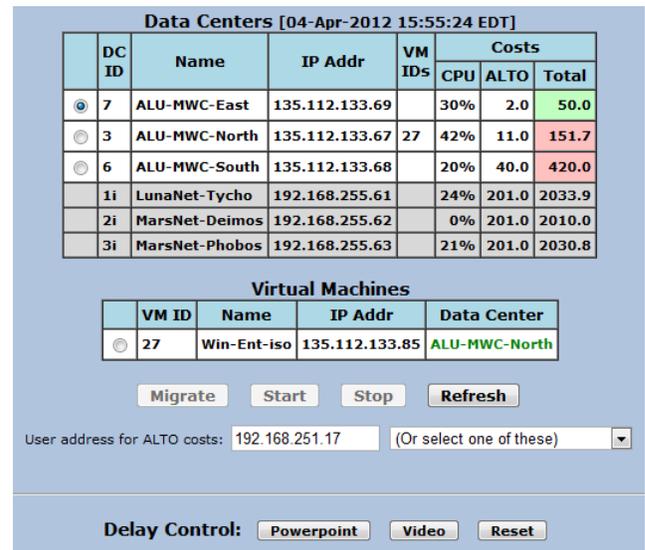


Figure 2: The cloud control dashboard showing the data centers and the virtual machines.

the performance problems: we display the network topology and live measurements of link and data center load (Figure 1). Provisioning of both network and compute performance highlights that selecting the right data center is a non-trivial optimization problem.

Next we show how our placement algorithm uses this data to calculate an overall score for each data center and suggest a data-center to host the user’s virtual desktop. Then we use our dashboard to migrate the virtual desktop to the new data-center (Figure 2). This is a manual step in our demonstration; in reality, the system would do this automatically without user intervention. Finally, we show how the video quality – the user’s experience – is restored to what it was at the beginning.

We further show a user working on a Powerpoint presentation with the virtual desktop. At first, Powerpoint responds smoothly to the user’s edits. But as we increase the network latency, the user finds it difficult, and finally impossible, to work effectively. Again, we use our dashboard and network monitoring tools to select a better data-center and migrate the desktop to it, thus restoring the user’s productivity.

In summary, our demonstration shows the tight coupling of the network with a distributed cloud solution using a system that integrates several of our own software components addressing individual challenges. We demonstrate that in case of network or data center problems migrating a service to another data center can mitigate the performance issues.

5. REFERENCES

- [1] M. Alicherry and T. V. Lakshman. Network Aware Resource Allocation in Distributed Clouds. In *INFOCOM*, 2012.
- [2] cloudstack - open source cloud computing. <http://www.cloudstack.org/>.
- [3] Gartner Cloud Market Prediction. <http://tiny.cc/x696bw>.
- [4] IETF ALTO working group. <http://datatracker.ietf.org/wg/alto/>.
- [5] Windows Remote Desktop Services. <http://tiny.cc/zt05bw>.