Towards a Network Marketplace in a Cloud

Rodrigo Fonseca
Brown University

Orran Krieger
Boston University

November 20, 2015

Cloud Computing, to this day, has been dominated by a limited number of cloud providers that, despite providing public service to tenants, are still vertically integrated and in control of their own infrastructure. As an example of why this is problematic, while it is possible for anyone to rent virtual machines from Amazon and offer a MapReduce service to clients, such an offering would be at a perpetual disadvantage vis-à-vis Amazon's own Elastic MapReduce offering. In this case, Amazon knows the location of data blocks, the network topology and allocation, and can place computation for its own service where it is most advantageous. The single-provider model of today's offerings eliminates competition at the infrastructure level, creates vendor lock-in, and artificial non-market-based pricing for resources [1].

In the Massachusetts Open Cloud [5] project, we are working on towards the vision of an Open Cloud Exchange (OCX) [2], a truly open, multi-provider cloud environment, in which unprivileged service providers compete in a marketplace, with offerings at all levels of the stack, including the hardware level. The only privileged infrastructure, beyond power, cooling, and basic connectivity, is a hardware-as-a-service allocation layer, coupled with a set of exchanges where offers from providers are matched with service requests from clients.

In the evolving OCX architecture, a provider can physically deploy machines in the common datacenter, and advertise access to these machines to clients. These clients can, in turn, offer an OpenStack service on top of such leased physical machines, or use them to run Spark and offer this as a service to higher-level clients of their own. There is nothing preventing other providers from also deploying physical machines and competing for features and prices with other providers.

At the lowest layers of this stack, the notion of multiple competing providers applies fairly naturally to storage and compute It does not immediately apply, however, to networking inside the datacenter. **One of our goals is to examine whether it makes sense for multiple providers to compete for networking services within an open cloud datacenter, including at the physical layer.** The prevailing view is, rather, that the datacenter network is a utility, a common substrate, and if it is sufficiently provisioned, should not be a cause for concern. While it is possible that this is the case, this view hinders innovation and differentiation at the network level, and our goal is to have an architecture that does not preclude this.

As examples, we want to enable a provider to connect a subset of the datacenter with 100Gbps Ethernet, or Infiniband, even if it is not cost effective to do this for the entire datacenter; another provider to offer a set of paths along which one can enable Ethernet flow control, use pFabric switches, or allow control of the switches' ECN parameters to enable DCTCP. Yet another provider might want to provide 60GHz connectivity among a few racks of the datacenter, or simply offer tenants access to high priority queues along specific paths.

The Internet offers a powerful analogy: an organization can today physically connect to an IXP, and from there choose services from several different transit providers that compete on capacity, reliability, connectivity, and cost. This architecture allows co-existence the 18+ low-latency providers between the New York and Chicago exchanges [4], or the multitude of cables that cross the Pacific connecting Los Angeles to Asia. In the OCX, a set of racks controlled by a single hardware provider would be analogous to an ISP, and the top of the rack switches would roughly play the role of IXPs, to which the different network providers would connect and offer "transit" service to other racks.

There are several challenges to realize this vision, including path specification and discovery, how to allow tenants to see different options and negotiate parameters, and how to implement the market decisions to steering subsets of tenants' traffic along the right paths. Some of these pieces have been solved for other domains, and can serve as inspiration, such as the work on software-defined exchanges (SDX) [3], or OpenVirtex. Other challenges include how to efficiently map overlay networks onto this now complex underlay, and how to integrate the market and provisioning of the network with the markets for other resources such as computation and storage.

The OCX model heavily depends on programmatic allocation, provisioning, isolation, and virtualization of resources – computation, storage, networking – and of more complex services that providers build from these basic elements. All of these services are exposed by providers in a marketplace, where the result of negotiations is then automatically realized, metered, and enforced in the infrastructure. This entire orchestration of resources can be seen as a complex set of logically centralized control planes, and is truly a software-defined infrastructure.

References

[1] O. Agmon Ben-Yehuda, M. Ben-Yehuda, A. Schuster, and D. Tsafrir. Deconstructing amazon ec2 spot instance pricing. In *Proceedings of the 2011 IEEE Third International Conference on Cloud Computing Technology and Science*, CLOUDCOM '11, pages 304–311, Washington, DC, USA, 2011. IEEE Computer Society.

[2] A. Bestavros and O. Krieger. Toward an open cloud marketplace: Vision and first steps. *IEEE Internet Computing*, 18(1):72–77, Jan. 2014.

[3] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett. SDX: A software defined internet exchange. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, SIGCOMM '14, pages 551–562, New York, NY, USA, 2014. ACM.

[4] G. Laughlin, A. Aguirre, and J. Grundfest. Information transmission between financial markets in chicago and new york. Working Paper 442, Stanford Law and Economics Olin, November 2012.

[5] Massachusetts open cloud. `http://www.bu.edu/hic/research/massachusetts-open-cloud/`.

About

Rodrigo Fonseca is an assistant professor at Brown University's Computer Science department, interested in networking, distributed systems and operating systems. Prior to Brown he was a post-doctoral researcher at Yahoo! Labs. He has done extensive research in resource management and performance tuning of large-scale distributed systems, and in networking architectures for wireless sensor networks. In the field of software-defined networking, he has done research in novel ways to integrate application knowledge into the network configuration, as well as in debugging of SDNs. Specifically in the MOC/OCX project, he has been focusing on the multi-provider network architecture. More details at `http://www.cs.brown.edu/people/rfonseca`.

Orran Krieger is the lead on the Massachusetts Open Cloud, Founding Director for the Cloud Computing Initiative (CCI) at BU, Resident Fellow of the Hariri Institute for Computing and Computational Science & Engineering, and a Professor of the practice at the Department of Electrical and Computer Engineering Boston University. Before coming to BU, he spent five years at VMware starting and working on vCloud. Prior to that he was a researcher and manager at IBM T. J. Watson, leading the Advanced Operating System Research Department. Orran did his PhD and MASc in Electrical Engineering at the University of Toronto.