

SuperCloud: Economical Cloud Service on Multiple Vendors

Qin Jia, Robbert Van Renesse, and Hakim Weatherspoon

1 Introduction

Today, Infrastructure-as-a-Service (IaaS) cloud providers such as Amazon's Elastic Compute Engine (EC2), Google's Compute Engine, and Microsoft's Azure offer elastic and isolated compute resources via virtualization and users often choose one of these providers based on price, locality, performance, and features. Typically, a user will choose the same provider for computation and storage to minimize latency and networking costs. Unfortunately, it can be difficult to switch providers once one is selected due to vendor lock-in [2].

We are challenging the status quo and investigating a new multi-provider model called a "supercloud" [2]. A supercloud supports applications launched across multiple public and private clouds: Virtual machines (VMs), storage, and even virtual network topologies can be migrated across many *heterogeneous* cloud providers. The enabling technology for superclouds is a uniform virtualization layer provided by nested virtualization, Xen-Blanket [1]. Moreover, control of resources and increased flexibility are achieved in the supercloud—placement of VMs, oversubscription of VMs, page sharing amongst VMs, etc.

Heterogeneity is an asset. VMs and storage can be placed intelligently and dynamically to optimize for price, latency, reliability, security, or other special features that are offered by the underlying public or private clouds and specified by the users. Meanwhile, oversubscription can be used to minimize the total compute resource used and save money.

But heterogeneity is also a challenge. Different cloud providers provide different interfaces and distinct hypervisors such as KVM and Xen. In addition, a supercloud has to take the networking overheads caused by migration into consideration.

Copyright © 2013 by the Association for Computing Machinery, Inc. (ACM). Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SoCC'13, 1–3 Oct. 2013, Santa Clara, California, USA.
ACM 978-1-4503-2428-1. <http://dx.doi.org/10.1145/2523616.2525966>

2 Design

To overcome the lack of homogeneity in cloud providers, Xen-blanket [1] could be used to virtualize the compute resources provided by public and private cloud vendors. Xen-blanket therefore offers a uniform VM abstraction to the upper layer—the supercloud layer. Consequently, the supercloud offers both traditional and improved IaaS, e.g security groups, VM placement, etc.

Our architecture contains 5 layers. Physical machines form the bottom layer. Xen or KVM hypervisors form the second layer. The next layer is the cloud layer. In addition to commercial cloud providers, this layer could also contain private clouds powered by OpenStack, Eucalyptus, etc. The fourth layer is the Xen-blanket layer, which homogenizes compute resources into uniform Xen VMs. It could be considered as a virtualized Xen hypervisor. Eventually, the topmost layer is the supercloud layer, which is yet another cloud layer but provides the most flexible and money-saving service across cloud vendors. Furthermore, since the supercloud VMs are placed in the VMs of the commercial cloud, oversubscription is possible—supercloud VMs that are not likely to peak together are placed into one Xen VM.

In order to leverage price variation, a VM placement decision technique is needed. Auto-regression techniques can be used to predict price trends. A supercloud places VMs on cloud providers with prices that are predicted to remain low. If costs improve elsewhere, superclouds are able to migrate VMs without loss of performance. As for the oversubscription, heuristic packing algorithms are able to pack as many supercloud VMs whose resource usages will not exceed the resource limit into one Xen VM as possible based on their observed and estimated usage. Therefore, the resource utilization is improved and more money is saved.

In summary, we are investigating a new ecosystem enabled by superclouds where multiple heterogeneous providers are used and resources economically consumed (VMs are intelligently placed or migrated and oversubscribed).

References

- [1] D. Williams, H. Jamjoom, and H. Weatherspoon. The Xen-Blanket: Virtualize once, run everywhere. In *the 7th ACM European Conference on Computer Systems (EuroSys)*, pages 113–126. ACM, 2012.
- [2] D. Williams, H. Jamjoom, and H. Weatherspoon. Plug into the supercloud. *IEEE Internet Computing*, 17(2):28–34, 2013.