A Highly-Virtualising Cloud Resource Broker

Josef Spillner*, Andrey Brito†, Francisco Brasileiro†, Alexander Schill*
* Faculty of Computer Science
Technische Universität Dresden
01062 Dresden, Germany
Email: {josef.spillner,alexander.schill}@tu-dresden.de
† Laboratório de Sistemas Distribuídos
Universidade Federal de Campina Grande
CEP 58429-900, Campina Grande - PB, Brazil
Email: {andrey,fubica}@dsc.ufcg.edu.br

***Abstract*—Cloud computing infrastructure services encompass temporally, spatially and structurally coarse-grained access to computing resources. This prevents a genuine pay-per-use model for the consumer and effectively leads to over-reservation and inadequate tariffs especially for short-term, small-scale computations. We propose an economically motivated compensation approach to increase the granularity and utility of reserved computation and storage services. The outcome is a highly-virtualising cloud resource broker. It consists of a fiduciary resource service marketplace and a consumer-configurable virtual machine for resource sharing. The system supports hierarchically nested virtualisation with dynamically adjustable resource limits for fine-grained structural, temporal and vertical-spatial scalability. We argue that cloud computing providers should embrace the recently developed hierarchical virtualisation and vertical scalability technologies in order to improve the utility and quality of experience for on-demand infrastructure service consumers.**

I. Background

Infrastructure-as-a-Service (IaaS) is a popular service category in the context of Cloud Computing. Computing, storage and communication resources delivered through it are provided on demand and billed according to pay-per-use schemes. The scalability of software in the cloud is high when both horizontal (scale-out) and vertical (scale-up) mechanisms are supported by the infrastructure services, although in practice most providers limit their offers to scale-out despite the higher management costs [1]. Both commercial users and researchers have adopted IaaS systems for their respective work. From a research perspective, the idea to deliver resources on demand sounds compelling, but the implemented systems are not yet very convincing. One drawback is that the resource reservations are still causing overhead which seems to be negligible for small remote calculations but add up considerably over time, causing too much overhead and overprovisioning from the consumer side due to the lack of fine-grained reservation and control facilities. The limitations include CPU billing by the full hour, arbitrary RAM sizes and inability to sub-allocate resources to external users.

In this paper, we present an intermediate cloud resource broker which works as an agent on the user's behalf across multiple providers, which besides offering vendor-neutral management procedures also establishes highly virtualised allocations from which consumers can repurpose and resell surplus slices to other prospective consumers.

II. Solution Roadmap

Our solution approach is to subdivide resource reservations into either serial or parallel segments (or slices) in a way that the use of either a segment or an entire reservation is transparent to the consumer. Instrumental to such a division are nested virtualisation and vertical scaling on the operating system and hypervisor levels beneath the cloud computing stack. Fig. 1 shows a sketch of both parallel and serial slicing of a cloud resource allocation and the expected overhead factors for both. Without nested virtualisation, parallel slices are insufficiently isolated from each other; and without vertical scaling, the slices cannot be reconfigured on the fly to match the effective demand.

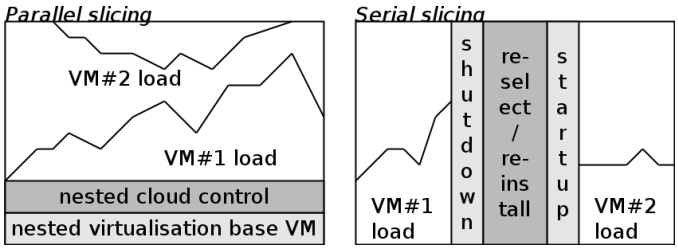

Fig. 1. Conceptual parallel and serial slicing of a cloud resource allocation

On our solution roadmap, we first have a look at nested virtualisation capabilities in current operating systems. Documentation on this topic is still sparse and many factors including hardware architecture and software versions contribute to feasible combinations. Then, for all successful combinations we determine the overhead in time and in resources: CPU usage, memory consumption, disk space. Afterwards, we translate this technical overhead into an economic one by evaluating current infrastructure service terms and conditions. As a result, we will determine the potential gain of using our approach from which we then derive the net gain by subtracting the operational cost. Finally, we anticipate an architecture for a highly-virtualising cloud resource broker with a user-friendly marketplace frontend and a cloud execution backend.

III. First Results and Discussion

Our experimental findings suggest the deployment of a KVM hypervisor with low-overhead nesting support of around $o = 6\%$ for the first and second level of virtualisation, respectively [2]. In economic terms, unused computing minutes u of an hour costing c translate to $U = (u/60) * (1 - o) * c$. To give an example, a process running for 70 minutes on an Amazon EC2 Large US-East instance ($c = \$9.76$, $u = 50$) opens a reselling potential of $U = \$7.64$ for the remaining time under the assumption of comparable process load and excluding further resource costs such as data transmission.

The broker design requires that users entrust it the provider login credentials to let it request the desired resource allocations on the user's behalf. The overall design, applicable to compute and storage services, is shown in Fig. 2.

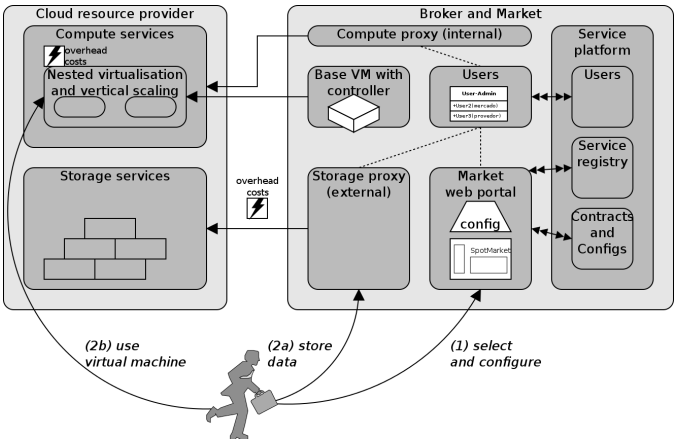

Fig. 2. Abstract architecture for a broker and market-in-the-middle for compute and storage resource services

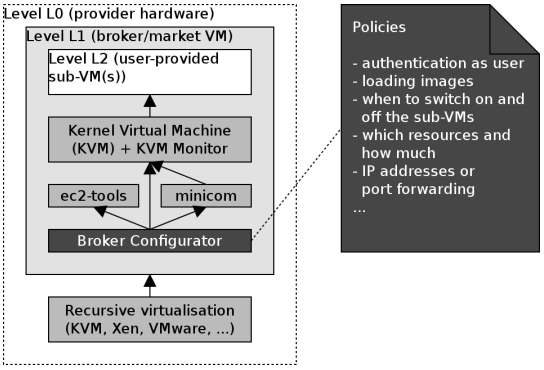

Fig. 3. Nested cloud with base virtual machine

A base virtual machine, shown in Fig. 3, is dedicated to enabling the nested cloud with other VMs referred to as sub-VMs running at a higher virtualisation level. The nested cloud VM is to be deployed by the broker and offers control facilities through a configurator which turn it into a light-weight infrastructure manager.

The facilities are bound to a service preparation tool running at the broker. It asks the potential user of a virtual machine instance about the utilisation and sharing of otherwise unused resource capacities: (1) No reuse; (2) Offer at the market for a cost; (3) Contribute to a grid or volunteer computing project in return for some incentive; (4) Reserve for another user, for free or at a cost; (5) Reserve for proper use by oneself.

We have implemented the design on top of SPACE, the Service Platform Architecture for Contracting and Execution [3]. For experiments, SPACE is integrated with a bootable operating system (Debian 6 with Linux kernel 2.6.32), a connector for VM hosting at cloud stacks (SPACE-Cloud extension using Eucalyptus 2.0.3) and service ontologies into the live demonstrator SPACEflight. To achieve a highly-virtualising broker, the system was made bootable with a custom Linux 3.5 kernel, extended with the nested cloud VM, and enhanced with the sharing variation on resource services. Furthermore, a cloud resource spot market was added as user-facing broker frontend and linked with the SPACE service registry. Its web interface is shown in Fig. 4.

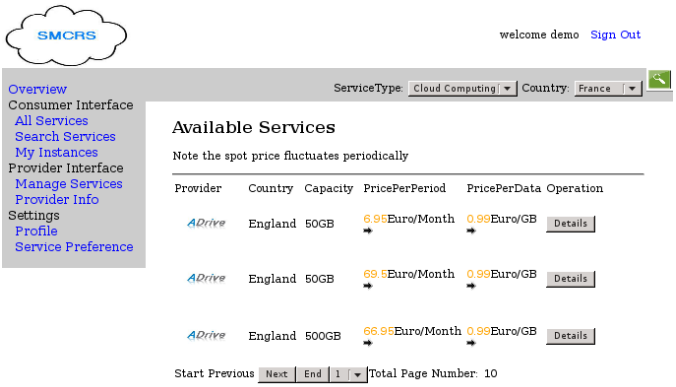

Fig. 4. Spot market for cloud resource services, with superfluous allocation slices offered by consumers (screenshot)

Acknowledgements

This work has received funding under project number 080949277 by means of the European Regional Development Fund (ERDF), the European Social Fund (ESF) and the German Free State of Saxony. It has also been funded by a scholarship from the Brazilian National Council for Scientific and Technological Development - CNPq (PDJ 159716/2011-0).

References

[1] M. Michael, J. E. Moreira, D. Shiloach, and R. W. Wisniewski, "Scale-up x Scale-out: A Case Study using Nutch/Lucene," in *21st IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, March 2007, pp. 1–8, Long Beach, California, USA.

[2] M. Ben-Yehuda, M. D. Day, Z. Dubitzky, M. Factor, N. Har'El, A. Gordon, A. Liguori, O. Wasserman, and B.-A. Yassour, "The Turtles Project: Design and Implementation of Nested Virtualization," in *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation (OSDI)*, October 2010, pp. 423–436, Vancouver, British Columbia, Canada.

[3] J. Spillner, "SPACEflight - A Versatile Live Demonstrator and Teaching System for Advanced Service-Oriented Technologies," in *21st International Crimean Conference on Microwave and Telecommunication Technology (CriMiCo/КрыМиКо)*, September 2011, pp. 455–456, Sevastopol, Ukraine. [Online]. Available: http://serviceplatform.org/