Cloud Service Engineering

Stefan Tai
Karlsruhe Institute of Technology (KIT)
Karlsruhe Service Research Institute (KSRI)
Englerstr. 11, 76131 Karlsruhe, Germany
stefan.tai@kit.edu

Jens Nimis
FZI Research Center for Information Technology
Haid-und-Neu-Str. 10-14, 76131 Karlsruhe, Germany
nimis@fzi.de

Alexander Lenk
FZI Research Center for Information Technology
Haid-und-Neu-Str. 10-14, 76131 Karlsruhe, Germany
lenk@fzi.de

Markus Klems
Karlsruhe Institute of Technology (KIT)
Karlsruhe Service Research Institute (KSRI)
Englerstr. 11, 76131 Karlsruhe, Germany
markus.klems@kit.edu

ABSTRACT

Building on compute and storage virtualization, Cloud Computing provides scalable, network-centric, abstracted IT infrastructure, platforms, and applications as on-demand services that are billed by consumption. Cloud Service Engineering is the application of a systematic approach to leverage Cloud Computing in the context of the Internet in its combined role as a platform for technical, economic, organizational and social networks. This tutorial introduces concepts and technology of Cloud Computing and Cloud Service Engineering, providing an overview of state-of-the-art in research and practice.

We demonstrate how to set up a private cloud that delivers Infrastructure-as-a-Service (IaaS). Eucalyptus and OpenNebula are popular open source software frameworks for creating on-premise clouds. Promises, challenges and solutions for integrating services of a private cloud with public cloud services such as Amazon EC2 and SQS are discussed. We show how the best of both worlds – private and public clouds – can be combined to build scalable and secure systems.

Categories and Subject Descriptors

C.2.4 [**Computer Systems Organization**]: Distributed Systems; C.0 [**Computer Systems Organization**]: General—*System Architectures*; D.2.11 [**Software Engineering**]: Software Architectures

ICSE '10, May 2-8 2010, Cape Town, South Africa

Keywords

Cloud Computing, Cloud Service Engineering

1. CLOUD SERVICE ENGINEERING

Cloud Computing is receiving enormous attention in the industry, mostly due to business-driven promises and expectations, including significantly lower upfront IT costs, a faster time to market, and opportunities for creating new business models and sources of value. From a technology perspective, Cloud Computing provides scalable, network-centric, abstracted IT infrastructure, middleware platforms, and applications as online, on-demand services. A good understanding of both the business drivers and technology enablers, hence, go hand in hand.

We define and propose *Cloud Service Engineering* as a discipline that applies a systematic approach to create value-add services on top of a Cloud Computing infrastructure. We argue that the tremendous potential of Clouds lies in making effective use of Cloud Computing as a distributed computing model in a business context, and that therefore the design of Cloud services must incorporate valuation methods. A rich ecosystem of Cloud services (and corresponding providers and consumers) enables compositions of Cloud services into a service-oriented business value network (SVN).

Hence, Cloud Service Engineering in particular addresses the following three new aspects:

1. Everything is (or can be) a service
2. Services have costs and value
3. Services constitute Value Networks

These three aspects describe open research challenges that require focused experimentation. We believe that establishing Cloud Service Engineering as a discipline is of critical

importance to the industry and for advancing Cloud Computing from the stage of expectations to the stage of demonstrated benefits and productivity.

1.1 Everything is (can be) a service

There is virtually no limit in what can be provided and consumed as an online, on-demand service. This includes infrastructure (such as physical resources, compute power, storage capacity, and bandwidth), middleware and platforms (programming and execution environments), software (basic and composite applications, ranging from office to social networking applications), and human intelligence [3]. However, how do we compose services of such diverse nature, from a programming model, from a systems runtime, and from a business perspective? Existing approaches to process-driven (Web) service composition or situational application development in the form of mashups take a restricted, confined view on services. The diversity of Cloud services already available today and the complexity of potential compositions of these however call for new composition and aggregation models.

1.2 Services have costs and value

The simple starting premise is that consumers of cloud services are charged based on actual service usage. But what are appropriate costs, pricing and billing models both from a consumer and a service provider viewpoint? How do we calculate costs, compare alternatives, estimate risks and determine value of traditional versus Cloud Computing solutions? Our observation is that we are still at a very early stage of fine-grained modeling and understanding of Cloud Computing costs and benefits. In [2], we devise a basic framework for estimating costs and benefits from Cloud Computing as an alternative to conventional IT infrastructure. We propose to calculate a "cost metric" and to estimate the value in terms of opportunity costs.

The importance of Cloud Computing economics is also discussed in [1]. The authors note that expected average and peak resource utilization must be carefully examined, and operational as well as power, cooling and physical plant costs must be determined. The difficulty in determining and predicting such data often presents major obstacles to adopting Cloud Computing.

1.3 Services constitute Value Networks

We expect many Cloud Computing environments to be of a hybrid nature, integrating private and public resources and composing value-add services from multiple providers, re-sellers (intermediaries), and integrators. The development of applications as compositions of "everything-as-a-service" (Research Challenge 1), where each service has a cost and value model (Research Challenge 2), naturally results in service-oriented business value networks (SVNs).

The formation and transformation of SVNs is driven by both technical criteria, including quality aspects (such as availability, reliability, trustworthiness), as well as business criteria including costs, reputation, policy and legal aspects of Cloud services and providers. Further, SVNs are likely to be subject to dynamic changes. Hence, we see the need for models, algorithms, tools, and platforms in support of the dynamic (trans-)formation of SVNs, including effective SLA management, change management, and failure management.

2. TUTORIAL AGENDA

Our tutorial on Cloud Service Engineering held at the ICSE 2010 conference consists of four major parts:

PART 1

- Introduction to Cloud Service Engineering
- Cloud Services and Software Engineering
- Cloud Architectures

PART 2

- Introduction to Amazon Web Services
- Programming Amazon Web Services with Java
- Cloud Service Recipes

PART 3

- Introduction to Open Source Cloud Systems: Eucalyptus and OpenNebula
- Building a Private Cloud

PART 4

- Cloud Computing Business Cases
- Open Issues in Cloud Research

3. CONCLUSION

Cloud computing is widely perceived as a disruptive technology that has the potential to fundamentally change the way services are built and delivered, and hence the way businesses operate. In this tutorial, we argue for Cloud Service Engineering as a discipline that takes a focused perspective on Cloud Computing and an integrative approach of technology and economics. We describe three aspects as open research challenges that require particular attention.

This tutorial provides the technical foundations to explore the field of cloud service engineering. Further references can be found at http://www.eOrganization.de.

4. REFERENCES

[1] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia. Above the clouds: A berkeley view of cloud computing. Technical Report UCB/EECS-2009-28, EECS Department, University of California, Berkeley, Feb 2009.

[2] M. Klems, J. Nimis, and S. Tai. Do clouds compute? a framework for estimating the value of cloud computing. In *Proceedings of the 7th Workshop on e-Business (WeB)*, 2008.

[3] A. Lenk, T. Sandholm, M. Klems, J. Nimis, and S. Tai. What's inside the cloud? an architectural map of the cloud landscape. In *Proceedings of the ICSE 2009 Workshop on Software Engineering Challenges in Cloud Computing*, 2009.