Cloud Computing Tutorial

Dan C. Marinescu
Computer Science Division
Department of Electrical Engineering and Computer Science
University of Central Florida, Orlando, FL 32816, USA
Email:dcm@cs.ucf.edu

January 15, 2013

1 Motivation

The idea that computation may be organized as a public utility, like water and electricity, was formulated in 1960s by John McCarthy, a visionary computer scientist who championed mathematical logic in artificial intelligence. Four decades later, utility computing was embraced by major IT companies such as Amazon, Apple, Google, HP, IBM, Microsoft, and Oracle.

Cloud computing is a movement started sometime during the middle of the first decade of the new millennium; the movement is motivated by the idea that information processing can be done more efficiently on large farms of computing and storage systems accessible via the Internet. Computer clouds support a paradigm shift from local to network-centric computing and network-centric content, when computing and storage resources are provided by distant datacenters. Scientific and engineering applications, data mining, computational financing, gaming and social networking, as well as many other computational and data-intensive activities can benefit from cloud computing. Storing information on the cloud has significant advantages and was embraced by cloud service providers. Content previously confined to personal devices such as workstations, laptops, tablets, and smart phones need no longer be stored locally, can be shared by all these devices, and it is accessible whenever a device is connected to the Internet.

The appeal of cloud computing is that it offers scalable and elastic computing and storage services. The resources used for these services can be metered and the users can be charged only for the resources they used. Cloud computing is a business reality, as a large number of organizations are adopting this paradigm.

Cloud computing reinforces the idea that computing and communication are deeply intertwined. Advances in one field are critical for the other. Indeed, cloud computing could not emerge as a feasible alternative to the traditional paradigms for high-performance computing before the Internet was able to support high-bandwidth, low-latency, reliable, low-cost communication; at the same time, modern networks could not function without powerful computing systems to manage the network.

There are virtually no bounds on composition of digital systems controlled by software, so we are tempted to build increasingly more complex systems. The behavior and the properties of such systems are not always well understood; thus, we should not be surprised that computing clouds will occasionally exhibit an unexpected behavior and system failures. The architecture, the coordination algorithms, the design methodology, and the analysis techniques for large-scale complex systems

like computing clouds will evolve in response to changes in technology, the environment, and the social impact of cloud computing.

This tutorial attempts to provide a snapshot of the state of the art in a dynamic field likely to experience significant developments in the near future. It starts with a view of cloud computing from the perspective of an application developer and gradually moves towards the inner working of a cloud and research issues in this exciting field.

The tutorial is based on the book "Cloud Computing: Theory and Practice" to be published by Morgan Kaufmann in May 2014.

2 Outline

1. Introduction
    - Network-centric computing and network-centric content
    - Cloud delivery models and services
    - Cloud vulnerabilities
    - Ethical issues in cloud computing
    - Major challenges
2. Cloud Infrastructure
    - Cloud computing at Amazon and Google
    - Open-source platforms
    - Energy use and ecological impact of large-scale data centers
    - Service level agreements
    - Software licensing
3. Applications and paradigms
    - Existing and new application opportunities
    - Architectural style of cloud applications
    - Coordination models - the Zookeeper
    - MapReduce programming model
    - Clouds for science and engineering
    - Cloud computing for biology research
    - Social computing, digital content and cloud computing
4. Resource Virtualization
    - Resource virtualization
    - Virtual machine monitors
    - Virtual machines
    - Performance and security isolation
    - Hardware support for virtualization; full and paravirtualization

- Xen
- The darker side of virtualization

5. Resource management and Scheduling
    - Polices and mechanisms
    - Resource bundling; combinatorial auctioning algorithms
    - Scheduling algorithms
    - Resource management and dynamic application scaling
6. Networking Support
    - The transformation of the Internet
    - Storage area networks
    - Content delivery networks
    - Overlay networks; scale-free and small worlds networks
7. Storage Systems
    - Google File System
    - Apache Hadoop
    - Chubby
    - Transaction processing and NoSQL databases
    - Bigtable
    - Megastrore
8. Cloud Security
    - Cloud security risks
    - Privacy and trust
    - OS and virtual machine security
    - Security of virtualization
    - Security risks posed by shared images
9. Complex Systems and Self-organization
    - Complex systems
    - Composability bounds and self-organization
    - Complexity of computing and communication systems
    - Systems of systems

3 Bio

Dan Cristian Marinescu was a Professor of Computer Science at Purdue University in West Lafayette Indiana from 1984 till 2001. Since August 2001 he is a Professor of Computer Science and Provost Research Professor at University of Central Florida. He has held visiting faculty positions at IBM T. J. Watson Research Center, (Yorktown Heights, NY); the Institute of Information Sciences and Tsinghua University, in Beijing, China; the Scalable Systems Division of Intel Corporation; Deutsche Telecom and GSI-Darmstadt in Germany; INRIA Rocquancourt in France; and Universidad Tecnica Federico Santa Maria (UTFSM) in Valparaiso, Chile.

His research interests are: scientific computing, process coordination and distributed computing, and quantum information processing; he has published more than 210 papers in professional journals and referred conference proceedings. He published several books: Internet-based Workflow Management (Wiley-2002), Process Coordination and Ubiquitous Computing (CRC Press -2002); Approaching Quantum Computing (co-authored with Gabriela M. Marinescu, Prentice Hall - 2005); Classical and Quantum Information, (co-authored with Gabriela M. Marinescu, Academic Press in 2011).