

Preface

Each day sees the growing popularity of the Internet and the Web. Along with the availability of powerful hand-held computing, mobile, and sensing devices the Internet is changing the way we interact, manage our lives, conduct business, and access or deliver services. The lowering costs of computation and communication are driving the focus from personal to Data Center-centric computing. Although parallel and distributed computing have been around for several years, their new forms, Multicore and Cloud computing, have brought about a sweeping change in the industry. These trends are pushing the industry focus from developing applications for PCs to Cloud Data Centers. This is enabling millions of users to make use of software simultaneously.

Computing is being transformed to a model having commoditised services. Such services are delivered in a manner similar to utilities such as water, electricity, gas, and telephony. As a result, IT (Information Technology) services are billed and delivered as “computing utilities” over shared delivery networks just as the water, electricity, gas and telephony services. In such a model, users access services based on their requirements regardless of where they are hosted. Several computing paradigms have promised to deliver this utility-computing vision. Cloud computing is the most recent emerging paradigm, promising to turn the vision of “computing utilities” into a reality.

Cloud computing has become a buzzword in the IT industry. Several IT vendors are promising to offer storage, computation and application-hosting services. They will also provide coverage in several continents, offering Service-Level Agreements (SLA) backed performance and uptime promises for their services. They offer subscription-based access to infrastructure, platforms, and applications, popularly termed IaaS (Infrastructure as a Service), PaaS (Platform as a Service), and SaaS (Software as a Service). These emerging services have reduced the cost of computation and application hosting by several steps. Still, is a significant complexity involved in the development and delivery of applications and their services in a seamless, scalable, and reliable manner.

There exist several Cloud technologies and platforms in the market. To mention a few: Google AppEngine, Microsoft Azure, and Manjrasoft Aneka. Google AppEngine provides an extensible runtime environment for Web-based applications, which leverage huge Google IT infrastructure. Microsoft Azure provides a wide array of Windows-based services for developing and deploying Windows-based applications on the Cloud. Manjrasoft Aneka provides a flexible model for creating Cloud applications and deploying them on a wide variety of infrastructures including public Clouds such as Amazon EC2.

With this sweeping shift from developing applications from PCs to Data Centres, there is a huge demand for people with new skill sets in Cloud computing. Universities play an important role in this regard. They can train the next generation of IT professionals and equip them with the necessary tools and knowledge to tackle these challenges. For this, they need to be able to set up a Cloud computing environment for teaching and learning with a minimal investment.

Currently, expert developers are required to create Cloud applications and services. Cloud researchers, practitioners, and vendors alike are working to educate potential users about the benefits of Cloud computing and the best way to make full use of it. However, being a new and popular concept, the very definition of Cloud computing depends on which computing expert is asked. So, while the dream of true utility computing appears closer to the truth than ever, its acceptance is currently restricted to Cloud experts. This is because of the perceived complexities of interacting with Cloud computing providers. This book aims to change the game by simplifying and imparting Cloud computing foundations, technologies, and programming skills to readers such that even the average programmers and software engineers are able to develop Cloud applications easily.

Salient Features

- Introduction to Cloud Computing, Cloud Architecture, Cloud Applications, Programming of Clouds, and Cloud Platforms
- Focuses on the platforms and technologies essential to Cloud Computing like Google AppEngine, Microsoft Azure, and Manjrasoft Aneka
- Dedicated chapter on Aneka: A Software Platform for .NET-based Cloud Computing
- Detailed coverage to Software as a Service (SaaS), Platform as a Service (PaaS), Infrastructure as a Service (IaaS), Server and Desktop Virtualization
- Pedagogy:
 - 120 Illustrations
 - 191 Review Questions
 - Footnotes interspersed within chapter contents

The Book at a Glance

This book introduces the fundamental principles of Cloud computing and its related concepts. It discusses virtualization technologies along with the architectural models of Cloud computing. It presents prominent Cloud computing technologies available in the marketplace, and contains dedicated chapters on concurrent, high-throughput and data-intensive computing paradigms and their use in programming Cloud applications. Various application case studies from domains such as science, engineering, gaming, and social networking are introduced along with their architecture and how they make use of various Cloud technologies. This allows the reader to understand the mechanisms needed to harness Cloud computing in their own respective endeavors. Finally, many open research problems and opportunities that have arisen from the rapid incorporation of Cloud computing are detailed. We hope that this motivates the reader to address these in their own future research and development.

The book contains 11 chapters, which can be organized into three major parts:

Part I: Foundations

- Chapter 1—Introduction
- Chapter 2—Principles of Parallel and Distributed Computing
- Chapter 3—Virtualization
- Chapter 4—Cloud Computing Architecture

Part II: Cloud Application Programming and the Aneka Platform

- Chapter 5—Aneka: Cloud Application Platform
- Chapter 6—Concurrent Computing: Thread Programming
- Chapter 7—High-Throughput Computing: Task Programming
- Chapter 8—Data Intensive Computing: Map-Reduce Programming

Part III: Industrial Platforms and New Developments

- Chapter 9—Cloud Platforms in Industry
- Chapter 10—Cloud Applications
- Chapter 11—Advanced Topics in Cloud Computing

The book serves as a perfect guide into the world of Cloud computing. Starting from the fundamentals, the book guides students and professionals through the practical use of these concepts by hands-on sessions on how to develop Cloud applications by using Aneka, Amazon Web Services, Google AppEngine, and Microsoft Azure. The last part introduces real applications, identifies emerging trends and offers future directions of Cloud computing.

Online Learning Center

The book also comes with an associated Web site (hosted at <http://www.mhhe.com/buyya/mcc1> and <http://www.buyya.com/MasteringClouds>) containing pointers to additional online resources, PowerPoint slides, and research papers.

Benefits and Readership

Given the rapid emergence of Cloud computing as a mainstream computing paradigm, it is essential to know-how have a solid understanding of the core concepts characterizing the subject and a practical of how to design and implement Cloud computing applications and systems as well. This set of skills is already fundamental today for software architects, engineers, and developers as many applications are being moved to the Cloud. It will become even more important in the future when this technology matures further. This book provides a perfect blend of background information, theory, and practical Cloud computing development, expressed in a language that is accessible to a wide range of readers: from graduate-level students to practitioners, developers, and engineers who want to, or need to, design and implement Cloud computing solutions. Moreover, advanced topics presented at the end of the manuscript make the book an interesting read for researchers in the field of Cloud computing who can get an overview of the next challenges in Cloud computing in the coming years.

This book is a timely contribution to the Cloud computing field that is gaining considerable commercial interest and momentum. The book is targeted at graduate students and IT professionals such as system architects, practitioners, software engineers, and application programmers. As Cloud computing is recognized as one of the top five emerging technologies that will have a major impact on the quality of science and society over the next 20 years, its knowledge will help position our readers at the forefront of the field.

Directions for Adoption: Theory, Laboratories, and Projects

Given the importance of the Cloud computing paradigm and its rapid acceptance in industry, universities/ educational institutions need to upgrade their curriculum by introducing one or more subjects in the area of Cloud computing and related topics such as parallel computing and distributed systems. We recommend that they offer at least one subject on Cloud computing as part of their undergraduate and postgraduate degree programs such as BE/B.Tech./BSc (Hons) in Computer Science and related areas; and Masters including the MCA (Master of Computer Applications). We believe that this book will serve as an excellent textbook for such subjects. If the students have already had exposure to the concepts of parallel and distributed computing, Chapter 2 can be skipped.

For those aiming to make their curriculum rich with Cloud computing, we recommend to offer two subjects: “Introduction to Cloud Computing” and “Advanced Cloud Computing” in two different semesters. This book has sufficient content to cater to both of them. The first subject can be based on Chapters 1 to 6 and the second one based on Chapters 7 to 11.

In addition to theory, we strongly recommend the introduction of a **laboratory subject** that offers hands-on experience. The laboratory exercises and assignments can focus on creating high-performance Cloud applications and assignments on a range of topics including parallel execution of mathematical functions, sorting of large data in parallel, image processing, and data mining. By using Cloud software systems, institutions can easily set up a private/enterprise Cloud computing facility by utilising existing LAN-connected PCs running Windows. Students can make use of this facility to learn about various Cloud application programming models and interfaces discussed in Chapter 6 (Thread Programming), Chapter 7 (Task Programming), and Chapter 8 (MapReduce Programming). Students need to learn various programming examples discussed in these chapters and execute them on Cloud facility. We encourage students to take up some of the programming exercises noted in the **Review Questions** section of these chapters as laboratory assignments and develop their own solutions.

Students can also carry out their final-year projects focused on developing Cloud applications solving real-world problems. For example, students can work with academics/researchers/experts from other science and engineering disciplines such as Life and Medical Sciences or Civil and Mechanical Engineering and develop suitable applications that can harness the power of Cloud computing. For inspiration, please read various application case studies presented in Chapter 11.

Rajkumar Buyya
Christian Vecchiola
S Thamarai Selvi

Publisher's Note

Do you have any further request or a suggestion? We are always open to new ideas (the best ones come from you!). You may send your comments to tmh.csefeedback@gmail.com
Piracy-related issues may also be reported!