

Mastering Cloud Computing

About The Authors

Dr. Rajkumar Buyya is Professor of Computer Science and Software Engineering; and Director of the Cloud Computing and Distributed Systems (CLOUDS) Laboratory at the University of Melbourne, Australia. He is the founder and CEO of Manjrasoft, a spin-off company of the University of Melbourne, commercializing its innovations in Cloud Computing. He has authored over 425 publications and four textbooks including *Mastering C++ and Object Oriented Programming with Java*. He is one of the highly cited authors in computer science and software engineering worldwide. The Microsoft Academic Search Index rates Dr Buyya as one of the top five authors in the world during the last decade (2001–2012) and the number one author in the world during the last five years (2007–2012) in the area of Distributed and Parallel Computing.

Software technologies for Grid and Cloud computing developed under Dr Buyya's leadership have gained rapid acceptance and are in use at several academic institutions and commercial enterprises in 40 countries around the world. Dr Buyya has led the establishment and development of key community activities, including serving as foundation Chair of the IEEE Technical Committee on Scalable Computing (TCSC) and several IEEE/ACM conferences. These contributions and international research leadership of Dr Buyya are recognized through the award of '2009 IEEE Medal for Excellence in Scalable Computing'. Manjrasoft's Aneka Cloud Technology developed under his leadership has received the '2010 Frost & Sullivan New Product Innovation Award' and the '2011 Telstra Innovation Challenge, People's Choice Award'. He is recently appointed as the foundation Editor-in-Chief (EIC) of IEEE Transactions on Cloud Computing.

Dr. Christian Vecchiola is Research Scientist at IBM Research Australia and Honorary Research Fellow at the Department of Computing and Information Systems, The University of Melbourne, Australia. He is currently working in the development of emergency management application solutions based on Cloud computing. Prior to joining IBM, Dr. Vecchiola served as Post-Doctoral Research Fellow in the Cloud Computing and Distributed Systems (CLOUDS) Laboratory at the University of Melbourne. He has designed, re-engineered, and implemented the several components of Aneka.

Dr. S Thamarai Selvi is Professor and Dean of Madras Institute of Technology (MIT) at Anna University, Chennai, India; and Principal Investigator of the Centre for Advance Computing Research and Education (CARE) supported by the Department of Information Technology (DIT), Ministry of Information and Communication Technology of the Government of India. She has authored over 150 publications and four textbooks including *Object Oriented Programming with Java* published by Tata McGraw Hill, India, 2009.

Dr Selvi has over 25 years of experience in teaching and research; and has coordinated and completed several research projects funded by various funding agencies including the Defence Research & Development Organisation (DRDO). She received an Indian patent titled "A system and a method to assess the trustworthiness of a resource provider" for her research on trust management in Grid computing. The CARE Resource Broker (CRB) developed for managing Grid and Cloud computing under her leadership has gained more popularity in India. She led the establishment of International conference on Advance Computing (ICoAC). For her outstanding research contributions, she has received the '2009 IBM Faculty Award' and 'Active Researcher Award' for the year 2011 from Anna University.

Mastering Cloud Computing

Rajkumar Buyya

*The University of Melbourne and Manjrasoft Pvt Ltd,
Melbourne, Australia*

Christian Vecchiola

*The University of Melbourne and IBM Research,
Melbourne, Australia*

S Thamarai Selvi

*Madras Institute of Technology,
Anna University,
Chennai, India*



Tata McGraw Hill Education Private Limited

NEW DELHI

McGraw-Hill Offices

New Delhi New York St Louis San Francisco Auckland Bogotá Caracas
Kuala Lumpur Lisbon London Madrid Mexico City Milan Montreal
San Juan Santiago Singapore Sydney Tokyo Toronto



Tata McGraw-Hill

Published by Tata McGraw Hill Education Private Limited,
7 West Patel Nagar, New Delhi 110 008

Mastering Cloud Computing

Copyright © 2013, by Tata McGraw Hill Education Private Limited.

No part of this publication may be reproduced or distributed in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise or stored in a database or retrieval system without the prior written permission of the publishers. The program listings (if any) may be entered, stored and executed in a computer system, but they may not be reproduced for publication.

Sales territories: India, Pakistan, Nepal, Bangladesh, Sri Lanka and Bhutan
Tata McGraw Hill Education Private Limited.

ISBN (13 digits): 978-1-25-902995-0

ISBN (10 digits): 1-25-902995-6

Vice President and Managing Director: *Ajay Shukla*

Head—Higher Education (Publishing and Marketing): *Vibha Mahajan*

Publishing Manager (SEM & Tech. Ed.): *Shalini Jha*

Assistant Sponsoring Editor: *Smruti Snigdha*

Editorial Researcher: *Amiya Mahapatra*

Manager—Production Systems: *Satinder S Baveja*

Assistant Manager—Editorial Services: *Sohini Mukherjee*

Production Executive: *Anuj K Shrivastava*

Assistant General Manager (Marketing)—Higher Education: *Vijay Sarathi*

Senior Product Specialist: *Tina Jajoriya*

Senior Graphic Designer—Cover: *Meenu Raghav*

General Manager—Production: *Rajender P Ghansela*

Manager—Production: *Reji Kumar*

Information contained in this work has been obtained by Tata McGraw-Hill, from sources believed to be reliable. However, neither Tata McGraw-Hill nor its authors guarantee the accuracy or completeness of any information published herein, and neither Tata McGraw-Hill nor its authors shall be responsible for any errors, omissions, or damages arising out of use of this information. This work is published with the understanding that Tata McGraw-Hill and its authors are supplying information but are not attempting to render engineering or other professional services. If such services are required, the assistance of an appropriate professional should be sought.

Typeset at BeSpoke Integrated Solutions, Puducherry, India 605 008 and printed at Magic International Pvt. Ltd., Plot No. 26E, Sector-31(INDUSTRIAL), Site-IV, Greater Noida - 201306

Cover Printer: Magic International Pvt.

RZZQCRDODDBQD



Contents

Preface	<i>xi</i>
Acknowledgements	<i>xv</i>
Chapter 1—Introduction	1.1
1.1 Cloud Computing at a Glance	1.1
1.1.1 The Vision of Cloud Computing	1.2
1.1.2 Defining a Cloud	1.4
1.1.3 A Closer Look	1.6
1.1.4 Cloud Computing Reference Model	1.8
1.1.5 Characteristics and Benefits	1.9
1.1.6 Challenges Ahead	1.10
1.2 Historical Developments	1.11
1.2.1 Distributed Systems	1.12
1.2.2 Virtualization	1.13
1.2.3 Web 2.0	1.14
1.2.4 Service-Oriented Computing	1.15
1.2.5 Utility-Oriented Computing	1.16
1.3 Building Cloud Computing Environments	1.17
1.3.1 Application Development	1.17
1.3.2 Infrastructure and System Development	1.17
1.4 Computing Platforms and Technologies	1.18
1.4.1 Amazon Web Services (AWS)	1.18
1.4.2 Google AppEngine	1.18
1.4.3 Microsoft Azure	1.19
1.4.4 Hadoop	1.19
1.4.5 Force.com and Salesforce.com	1.19
1.4.6 Manjrasoft Aneka	1.19
<i>Summary</i>	1.20
<i>Review Questions</i>	1.21
Chapter 2—Principles of Parallel and Distributed Computing	2.1
2.1 Eras of Computing	2.1
2.2 Parallel vs. Distributed Computing	2.2
2.3 Elements of Parallel Computing	2.2
2.3.1 What is Parallel Processing?	2.3
2.3.2 Hardware Architectures for Parallel Processing	2.3
2.3.3 Approaches to Parallel Programming	2.6
2.3.4 Levels of Parallelism	2.7
2.3.5 Laws of Caution	2.8
2.4 Elements of Distributed Computing	2.8

2.4.1 General Concepts and Definitions	2.8
2.4.2 Components of a Distributed System	2.9
2.4.3 Architectural Styles for Distributed Computing	2.11
2.4.4 Models for Inter-Process Communication	2.19
2.5 Technologies for Distributed Computing	2.22
2.5.1 Remote Procedure Call	2.22
2.5.2 Distributed Object Frameworks	2.23
2.5.3 Service Oriented Computing	2.28
<i>Summary</i>	2.34
<i>Review Questions</i>	2.35
Chapter 3—Virtualization	3.1
3.1 Introduction	3.1
3.2 Characteristics of Virtualized Environments	3.3
3.3 Taxonomy of Virtualization Techniques	3.6
3.3.1 Execution Virtualization	3.7
3.3.2 Other Types of Virtualization	3.16
3.4 Virtualization and Cloud Computing	3.17
3.5 Pros and Cons of Virtualization	3.19
3.6 Technology Examples	3.21
3.6.1 Xen: Paravirtualization	3.21
3.6.2 VMware: Full Virtualization	3.22
3.6.3 Microsoft Hyper-V	3.27
<i>Summary</i>	3.31
<i>Review Questions</i>	3.31
Chapter 4—Cloud Computing Architecture	4.1
4.1 Introduction	4.1
4.2 Cloud Reference Model	4.2
4.2.1 Architecture	4.2
4.2.2 Infrastructure / Hardware as a Service	4.4
4.2.3 Platform as a Service	4.6
4.2.4 Software as a Service	4.9
4.3 Types of Clouds	4.12
4.3.1 Public Clouds	4.12
4.3.2 Private Clouds	4.13
4.3.3 Hybrid Clouds	4.15
4.3.4 Community Clouds	4.17
4.4 Economics of the Cloud	4.19
4.5 Open Challenges	4.21
4.5.1 Cloud Definition	4.21
4.5.2 Cloud Interoperability and Standards	4.22
4.5.3 Scalability and Fault Tolerance	4.23

4.5.4 Security, Trust, and Privacy	4.23
4.5.5 Organizational Aspects	4.23
<i>Summary</i>	4.24
<i>Review Questions</i>	4.25
Chapter 5—Aneka: Cloud Application Platform	5.1
5.1 Framework Overview	5.2
5.2 Anatomy of the Aneka Container	5.4
5.2.1 From the Ground Up: Platform Abstraction Layer	5.4
5.2.2 Fabric Services	5.5
5.2.3 Foundation Services	5.6
5.2.4 Application Services	5.9
5.3 Building Aneka Clouds	5.11
5.3.1 Infrastructure Organization	5.11
5.3.2 Logical Organization	5.12
5.3.3 Private Cloud Deployment Mode	5.13
5.3.4 Public Cloud Deployment Mode	5.14
5.3.5 Hybrid Cloud Deployment Mode	5.15
5.4 Cloud Programming and Management	5.16
5.4.1 Aneka SDK	5.17
5.4.2 Management Tools	5.20
<i>Summary</i>	5.21
<i>Review Questions</i>	5.22
Chapter 6—Concurrent Computing: Thread Programming	6.1
6.1 Introducing Parallelism for Single Machine Computation	6.1
6.2 Programming Applications with Threads	6.3
6.2.1 What is a Thread?	6.3
6.2.2 Thread APIs	6.5
6.2.3 Techniques for Parallel Computation with Threads	6.6
6.3 Multithreading with Aneka	6.19
6.3.1 Introducing the Thread Programming Model	6.20
6.3.2 Aneka Thread vs. Common Threads	6.21
6.4 Programming Applications with Aneka Threads	6.24
6.4.1 Aneka Threads Application Model	6.24
6.4.2 Domain Decomposition: Matrix Multiplication	6.26
6.4.3 Functional Decomposition: Sine, Cosine, and Tangent	6.33
<i>Summary</i>	6.39
<i>Review Questions</i>	6.40
Chapter 7—High-Throughput Computing: Task Programming	7.1
7.1 Task Computing	7.1
7.1.1 Characterizing a Task	7.2
7.1.2 Computing Categories	7.3

7.1.3 Frameworks for Task Computing	7.4
7.2 Task-based Application Models	7.5
7.2.1 Embarrassingly Parallel Applications	7.5
7.2.2 Parameter Sweep Applications	7.6
7.2.3 MPI Applications	7.8
7.2.4 Workflow Applications with Task Dependencies	7.10
7.3 Aneka Task-Based Programming	7.13
7.3.1 Task Programming Model	7.13
7.3.2 Developing Applications with the Task Model	7.14
7.3.3 Developing Parameter Sweep Application	7.30
7.3.4 Managing Workflows	7.34
<i>Summary</i>	7.36
<i>Review Questions</i>	7.37
Chapter 8—Data Intensive Computing: Map-Reduce Programming	8.1
8.1 What is Data-Intensive Computing?	8.1
8.1.1 Characterizing Data-Intensive Computations	8.2
8.1.2 Challenges Ahead	8.2
8.1.3 Historical Perspective	8.3
8.2 Technologies for Data-Intensive Computing	8.7
8.2.1 Storage Systems	8.7
8.2.2 Programming Platforms	8.14
8.3 Aneka MapReduce Programming	8.20
8.3.1 Introducing the MapReduce Programming Model	8.21
8.3.2 Example Application	8.44
<i>Summary</i>	8.56
<i>Review Questions</i>	8.56
Chapter 9—Cloud Platforms in Industry	9.1
9.1 Amazon Web Services	9.1
9.1.1 Compute Services	9.2
9.1.2 Storage Services	9.7
9.1.3 Communication Services	9.14
9.1.4 Additional Services	9.15
9.1.5 Summary	9.16
9.2 Google AppEngine	9.16
9.2.1 Architecture and Core Concepts	9.16
9.2.2 Application Life-Cycle	9.21
9.2.3 Cost Model	9.23
9.2.4 Observations	9.23
9.3 Microsoft Azure	9.24
9.3.1 Azure Core Concepts	9.24
9.3.2 SQL Azure	9.28

9.3.3 Windows Azure Platform Appliance	9.30
9.3.4 Summary	9.30
9.4 Observations	9.31
<i>Review Questions</i>	9.31
Chapter 10—Cloud Applications	10.1
10.1 Scientific Applications	10.1
10.1.1 Healthcare: ECG Analysis in the Cloud	10.1
10.1.2 Biology: Protein Structure Prediction	10.3
10.1.3 Biology: Gene Expression Data Analysis for Cancer Diagnosis	10.4
10.1.4 Geoscience: Satellite Image Processing	10.5
10.2 Business and Consumer Applications	10.6
10.2.1 CRM and ERP	10.6
10.2.2 Productivity	10.8
10.2.3 Social Networking	10.11
10.2.4 Media Applications	10.12
10.2.5 Multiplayer Online Gaming	10.15
<i>Summary</i>	10.16
<i>Review Questions</i>	10.17
Chapter 11 —Advanced Topics in Cloud Computing	11.1
11.1 Energy Efficiency in Clouds	11.1
11.1.1 Energy-Efficient and Green Cloud Computing Architecture	11.3
11.2 Market Based Management of Clouds	11.5
11.2.1 Market-Oriented Cloud Computing	11.5
11.2.2 A Reference Model for MOCC	11.6
11.2.3 Technologies and Initiatives Supporting MOCC	11.11
11.2.4 Observations	11.16
11.3 Federated Clouds / InterCloud	11.16
11.3.1 Characterization and Definition	11.16
11.3.2 Cloud Federation Stack	11.17
11.3.3 Aspects of Interest	11.22
11.3.4 Technologies for Cloud Federations	11.38
11.3.5 Observations	11.41
11.4 Third Party Cloud Services	11.42
11.4.1 MetaCDN	11.42
11.4.2 SpotCloud	11.43
<i>Summary</i>	11.44
<i>Review Questions</i>	11.45
References	R.1
Index	I.1

Preface

The growing popularity of the Internet and the Web along with the availability of powerful hand-held computing, mobile and sensing devices are 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 has been around for several years, its 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 enabling millions of users to make use of software simultaneously.

Computing is being transformed to a model consisting of commoditised services 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 akin to 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 one of the buzzwords in the IT industry. Several IT vendors are promising to offer storage, computation and application-hosting services, and 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). Whilst these emerging services have reduced the cost of computation and application hosting by several orders of magnitude, there 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 manpower with new skill sets in Cloud computing. Universities play an important role in this regard by training the next generation of IT professionals and equipping them with the necessary tools and knowledge to tackle these challenges. 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 ensure potential users are educated about the benefits of Cloud computing and the best way to harness its full potential. However, being a new and popular paradigm, the very definition of Cloud computing depends on which computing expert is asked. So, while the realization of true utility computing appears closer than ever, its acceptance is currently restricted to Cloud experts due to 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 paradigms. It discusses the concepts of virtualization technologies along with the architectural models of Cloud computing. It presents prominent Cloud computing technologies available in the marketplace. It 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 leverage 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 uptake 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. By starting from the fundamentals, the book drives 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 both have a solid understanding of the core concepts characterizing the phenomenon and a practical grasp of how to design and implement Cloud computing applications and systems. 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, more advanced topics presented at the end of the manuscript, make the book an interesting reading for researchers in the field of Cloud computing that are can get an overview of the next challenges in Cloud computing for 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 recognised 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 uptake 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 post-graduate 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 for the introduction of a **laboratory subject** that offers hands-on experience. The lab 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 programming exercises noted in the **Review Questions** section of these chapters as lab 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!

Acknowledgments

First and foremost, we are grateful to all researchers and industrial developers worldwide for their contributions to various concepts and technologies discussed in the book. Our special thanks to all the members/consultants of Manjrasoft, CLOUDS (Cloud Computing and Distributed Systems) Lab from the University of Melbourne, and Melbourne Ventures who contributed towards the development of Aneka Cloud Application Platform, the preparation of associated application demonstrators and documents, and/or the commercialization of the Aneka technology. They include Chu Xingchen, Srikumar Venugopal, Krishna Nadiminti, Dileban Karunamoorthy, Chao Jin, Rodrigo Calheiros, Michael Mattess, Jessie Wei, Enayat Masoumi, Ivan Mellado, Richard Day, Wolfgang Gentsch, Laurence Liew, David Sinclair, Suraj Pandey, Abhi Shekar, Dexter Duncan, Murali Sathya, Karthik Sukumar, and Sita Venkatraman.

We are grateful to the Australian Research Council (ARC) and the Department of Innovation, Industry, Science and Research (DIISR) for supporting our research and commercialization endeavors.

We thank all of our colleagues at the University of Melbourne especially Professors Rao Kotagiri, Iven Mareels, and Glyn Davis for their mentorship and positive support for our research and knowledge transfer efforts.

A note of appreciation is due to all colleagues and users of the Aneka technology for their direct/indirect contributions towards application case studies reported in the book. Our special thanks to Raghavendra from ADRIN/ISRO for his enthusiastic efforts in creating satellite image-processing application using Aneka and publishing articles in this area. Srinivasa Iyengar from MSRIT deserves a special mention for creating data-mining applications using Aneka and demonstrating the power of Aneka to academics from the early days of Cloud computing.

We thank the members of the CLOUDS Lab for proofreading one or more chapters. They include Rodrigo Calheiros, Nikolay Grozev, Amir Vahid, Anton Beloglazov, Adel Toosi, Deepak Poola, Mohamed AlRokayan, Atefeh Khosravi, Sareh Piraghaj, and Yaser Mansouri.

Dr Buyya is grateful to his family members, including Smrithi Buyya, Soumya Buyya, and Radha Buyya, for their love and understanding during the preparation of the book.

We sincerely acknowledge external reviewers K Chandra Sekaran of NIT Shraathkal, Sateesh K Peddoju of IIT Roorkee, M P Sebastian of IIM Kozhikode and a few others for their critical comments and suggestions on enhancing the presentation and organization of many chapters at a finer level. This has greatly helped us in improving the quality of the book.

Finally, we would like to thank the staff at Tata McGraw Hill Education for their enthusiastic support and guidance during the preparation of the book. In particular, Vibha Mahajan inspired us to take up this project and set the publication process in motion; Smruti Snigdha and Amiya Mahapatra for coordinating the publication process. They were wonderful to work with!



Guided Tour



Covers all important subject areas

4 Cloud-Computing Architecture

The term "Cloud computing" is a wide umbrella encompassing many different things. Lately, it has become a buzzword, easily misused to revamp existing technologies and ideas for the public. What makes it so interesting to IT stakeholders and research practitioners? How does it introduce innovation into the field of distributed computing? This chapter addresses all these questions and characterizes the phenomenon. It provides a reference model, which serves as a basis for discussion on Cloud-computing technologies.

4.1 INTRODUCTION

Utility-oriented data centers are the first outcome of Cloud-computing, and they serve as the infrastructure through which the services are implemented and delivered. Any Cloud service, whether it is virtual hardware, development platform, or application software, relies on a distributed infrastructure owned by the provider or rented from a third party. As it can be noticed from the previous discussion, the characterization of a Cloud is quite general: it can be implemented by using a distributed collection of clusters, or a heterogeneous distributed system composed by desktop PCs, workstations, and servers. Commonly, Clouds are built by relying on one or more datacenters. In most of the cases, hardware resources are virtualized to provide isolation of workloads and to exploit, at best, the infrastructure. According to the specific service delivered to the end user, different layers can be defined on top of the virtual infrastructure.

As noted, this is a broad definition. Cloud computing offers several different options for building enterprise Cloud computing applications or computing technologies to integrate and extend existing industrial applications. An important Cloud computing platform and a brief description of the type of service offered are shown in Table 9.1. A Cloud computing system can be developed by using either a vendor or a combination of them.

This chapter presents some of the representative Cloud computing solutions offered as Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) services in market. It provides a practical overview of the architecture of the major Cloud computing technology offerings.

5 Aneka: Cloud-Application Platform

Aneka is Manjrasoft's solution for developing, deploying, and managing Cloud applications. It consists of a scalable Cloud middleware that can be deployed on top of heterogeneous computing resources. It offers an extensible collection of services coordinating the execution of applications, helping administrators to monitor the status of the Cloud, and providing integration with existing Cloud technologies. One of the key advantages of Aneka is its extensible set of APIs associated with different types of programming models—such as Task, Thread, and MapReduce—used for developing distributed applications. Integrating new capabilities into the Cloud, and supporting different types of Cloud deployment models: public, private, and hybrid (see Fig. 5.1). These features differentiate Aneka from infrastructure management software and characterize it as a platform for developing, deploying and managing execution of applications on various types of Clouds.

This chapter provides a complete overview of the framework by firstly describing the architecture of the system. It introduces the component and the fundamental services building up the Aneka Cloud and some common deployment scenarios.

9 Cloud Platforms in Industry

9.1 AMAZON WEB SERVICES

Amazon Web Services (AWS) is a platform allowing the development of flexible applications for elastic infrastructure scalability, messaging, and file and data storage accessible through SOAP or RESTful Web service interfaces and provides a Web based users can administrate and monitor the resources required as well as their expenses pay as you go basis.

Table 9.1. Some Examples of Cloud-Computing Offerings.

Vendor/ Product	Service Type	Description
Amazon Web Services	IaaS, PaaS, SaaS	Amazon Web Services (AWS) is a collection of services providing developers with compute more advanced services. AWS is mostly provided as Infrastructure-as-a-Service (IaaS) services and primarily for its elastic computing capabilities.

10 Cloud Applications

Cloud computing has gained huge popularity in industry due to its ability to host applications whose services can be delivered to consumers rapidly at minimal cost. This chapter discusses various application case studies detailing their architecture and how they leveraged various Cloud technologies. Applications from a range of domains—from scientific to engineering, gaming, to social networking—are considered.

10.1 SCIENTIFIC APPLICATIONS

Scientific applications are a sector that is increasingly using Cloud computing systems and technologies. The immediate beneficiaries by researchers and academics is the potentially infinite availability of computing resources and storage at sustainable prices. If compared to a complete in-house deployment, Cloud computing systems meet the needs of different types of applications in the scientific domain: High Performance Computing (HPC) applications, High Throughput Computing (HTC) applications, and data-intensive applications. The opportunity for using Cloud resources is even more appealing since minimal changes need to be done to existing applications in order to leverage Cloud resources.

The most relevant option is Infrastructure-as-a-Service solutions, which offer the optimal environment for running bag-of-tasks applications and workflows. Virtual machine instances are opportunistically customized to host the required software stack for running such applications and coordinated together by distributed computing middleware capable of interacting with Cloud-based infrastructures. Platform-as-a-Service solutions have been also considered. They allow scientists to explore new programming models for tackling computationally challenging problems. Applications have been redesigned and implemented on top of Cloud programming application models and platforms to leverage their unique capabilities. For instance, MapReduce programming model provides scientists with a very simple and effective model for building applications that need to process large datasets. Therefore, it has been widely used to develop data-intensive scientific applications. Problems that require a higher degree of flexibility in terms of structuring their computation model can leverage platforms such as Aneka, which supports MapReduce and other programming models. We now discuss some interesting case studies in which Aneka has been used.

10.1.1 Healthcare: ECG Analysis in the Cloud

Healthcare is a domain where computer technology has found diverse applications: from supporting the business functions to assisting scientists in developing solutions to cure diseases. An



- More than 100 illustrations and diagrams are present to enhance the concepts.
- Full-page figures add greater clarity to the subject.

Fig. 1.3. A Bird's Eye View of Cloud Computing

This diagram illustrates the relationship between different cloud models. At the top, 'Public Clouds' (Public Internet Clouds), 'Private Clouds' (Private Enterprise Clouds), and 'Hybrid Clouds' (Hybrid/Enter Clouds) are shown. Below them, 'Other Cloud Services' and 'Govt. Cloud Services' are also depicted. A central 'Cloud Manager' oversees 'Clients', 'Private Cloud', and 'Public Cloud'. The 'Public Cloud' contains 'Development and Runtime Platform' and 'Compute' components, which are linked to 'Storage' and 'Applications'. A 'Hybrid Cloud' is shown as a 'Mixed usage of private and public clouds, Leasing public cloud services when private cloud capacity is insufficient'.

Fig. 1.4. Log Level Series Distribution

This pie chart shows the distribution of log levels: INFO (22%), DEBUG (12%), WARN (9%), ERROR (4%), and SKIPPED (62%).

Fig. 1.5. MapReduce Scheduling Service Architecture

This diagram shows the architecture of the MapReduce Scheduling Service. It includes 'MapReduceApplication' and 'MapReduceExecutionService' components. The 'MapReduceScheduler' manages 'MapReduceExecutorService' and 'MapReduceSchedulerService'. The architecture is divided into 'DFS Implementation' and 'MapReduce Scheduler' sections.

Fig. 1.6. RESORVER Cloud Deployment

This diagram illustrates the deployment of the RESORVER cloud. It shows 'Service Applications' (Service Application 1, 2, 3) running on 'Virtualizer' and 'Computational Resource' components. These are managed by 'VEE Manager (VEEM)' and 'VEE Hosts' across multiple 'Sites' (Site A, Site B, Site C). The 'Service Provider' and 'Service Manager' are also shown in the deployment structure.



Each chapter has an extensive **Summary** for quick recapitulation of the concepts discussed.

Summary

In this chapter, we have introduced task-based programming and provided an overview of the technologies supporting the development of distributed applications based on the concept of tasks. Task-based programming constitutes the most intuitive approach for distributing the computation of an application over a set of nodes. The main abstraction of task-based programming is the concept of *task*, which represents a group of operations that can be isolated and executed as a single unit. A task can be a simple program that is executed through the shell or a more complex piece of code requiring a specific runtime environment to execute. Quite often, tasks require input files for their execution and produce output files as a result. According to this model, an application is expressed as a collection of tasks; the way in which these tasks are interrelated and their specific nature and characteristics differentiate the different models that are an expression of task-based programming.

Traditionally, task-based programming model has been successfully used in the development of distributed applications in many areas. We identified three major computing categories where task model can be utilized. *High Performance Computing (HPC)* refers to the use of distributed computing facilities for solving problems needing large computing power. Common HPC applications feature a large collection of compute intensive tasks whose duration is relatively short. *High Throughput Computing (HTC)* identifies scenarios where distributed computing facilities are used to support the execution of applications needing large computing power for a long period of time. Tasks may not be numerous, but have a long duration, and infrastructure reliability becomes fundamental. *Many Task Computing (MTC)* is the latest emergent trend, and identifies a heterogeneous set of applications and requirements for applications, which fill the gap between HPC and HTC.

We have briefly reviewed common models related to task programming. *Embarassingly parallel* applications are composed of a collection of tasks which do not relate to each other, can be executed in any order, and do not require co-allocation. *Parameter sweep* applications are a special instance of *embarassingly parallel* model. They are characterized by a collection of independent tasks which are automatically generated from a template task by varying the combination of parameter values. In this case, the task executed is the same in terms of computation logic, but operates on different data. Therefore, parameter sweep application can also be considered an expression of the *(Single Program Multiple Data/SPMD)* model. *MPI* applications are characterized by a collection of tasks that need to be executed all together and which exchange data by message passing. Even though the program executed by an MPI application might be the same, it is quite common to provide an implementation logic that differentiates the behavior of each task according to its rank. *Workflow* applications are characterized by a collection of tasks whose dependencies can be expressed in terms of a directed acyclic graph. Dependencies are mostly represented by files, which are produced as output of a specific task, and are required for the computation of the dependent tasks. The nature of the tasks and kind of computation performed by each task is, in general, different.

Review Questions

1. What is the innovative characteristic of Cloud computing?
2. Which are the technologies that Cloud computing relies on?
3. Provide a brief characterization of a distributed system.
4. Define Cloud computing and identify its core features.
5. What are the major distributed computing technologies that led to Cloud computing?
6. What is virtualization?
7. What is the major revolution introduced by Web 2.0?
8. Give some examples of Web 2.0 applications.
9. Describe the main characteristics of service orientation.
10. What is utility computing?
11. Describe the vision introduced by Cloud computing.
12. Briefly summarize the Cloud computing reference model.
13. What is the major advantage of Cloud computing?
14. Briefly summarize the challenges still open in Cloud computing.
15. How does Cloud development differentiate from traditional software development?

Review questions are given in each chapter to test the student's subjective grasp on the topics, terms and definitions, and revision of concepts



A **Reference list** is provided at the end of the book to help students find books and journals for further reading.

References

- [1] A.S. Tanenbaum and M. Van Steen, *Distributed Systems: Principles and Paradigm*, Prentice Hall PTR, Upper Saddle River, NJ, USA, 2001.
- [2] G. Coulouris, J. Dalmonte, and T. Khabib, *Distributed Systems: Concepts and Design*, 4th Ed., Addison Wesley, 2005.
- [3] G. Plister, *In Search of Clusters*, 2nd Ed., Prentice Hall PTR, NJ, 1996.
- [4] R. Buyya, *High-Performance Cluster Computing: Architecture and Systems*, Prentice Hall PTR, NJ, 1999.
- [5] D. Thain, T. Tannenbaum, and M. Livny, *Distributed Computing in Practice: The Concor Experience, Concurrency and Computation: Practice and Experience*, 17(2-4), pp. 323-356, 2005.
- [6] U.S. Sunderam, *PVM: A Framework for Parallel Distributed Computing, Concurrency, Practice and Experience*, 2(4), pp. 315-339, 1990.
- [7] W. Gropp, E. Lusk, N. Doss, and A. Skjellum, *A High-Performance, Portable Implementation of the MPI Message Passing Interface*, *Parallel Computing*, 22(6), pp. 789-828, 1996.
- [8] I. Foster and C. Kesselman, *The Grid Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishing, 1999.
- [9] I. Foster, C. Kesselman, and S. Tuecke, *The Anatomy of the Grid: Enabling Scalable Virtual Organizations*, *International Journal of High Performance Computing Applications*, 15(3), 2001.
- [10] J. Broberg, S. Venugopal, and R. Buyya, *Market Oriented Grids and Utility Computing: The State of the Art and Future Directions*, *Journal of Grid Computing*, 6(3), Springer Netherlands, pp. 255-276, 2008.
- [11] R. Buyya and K. Babendorfer, *Market-Oriented Grid Computing and Utility Computing*, Wiley, 2006.
- [12] I. Foster, *Globus Toolkit Version 4: Software for Service-Oriented Systems*, IFIP International Conference on Network and Parallel Computing, Lecture Notes in Computer Science (LNCS) 3779, Springer-Verlag, pp. 2-13, 2005.
- [13] W. Gombusch, *Sun Grid Engine: Towards Creating a Compute Power Grid*, Proceedings of the 1st IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2001), IEEE Computer Society, Brisbane, Australia, pp. 35-36, 2001.

