

Wiley Series on Parallel and Distributed Computing • Albert Y. Zomaya, Series Editor

CLOUD COMPUTING

Principles and Paradigms

Edited by

Rajkumar Buyya
James Broberg
Andrzej Goscinski

CLOUD COMPUTING

Principles and Paradigms

Edited by

Rajkumar Buyya

The University of Melbourne and Manjrasoft Pty Ltd., Australia

James Broberg

The University of Melbourne, Australia

Andrzej Goscinski

Deakin University, Australia



A JOHN WILEY & SONS, INC., PUBLICATION

Copyright © 2011 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey
Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Cloud computing : principles and paradigms / edited by Rajkumar Buyya, James Broberg, Andrzej Goscinski.

p. ; cm.

Includes bibliographical references and index.

ISBN 978-0-470-88799-8 (hardback)

1. Cloud computing. I. Buyya, Rajkumar, 1970— II. Broberg, James. III. Goscinski, Andrzej.

QA76.585.C58 2011

004.67'8—dc22

2010046367

Printed in Singapore

10 9 8 7 6 5 4 3 2 1

CLOUD COMPUTING

PREFACE

Cloud computing has recently emerged as one of the buzzwords in the ICT industry. Numerous IT vendors are promising to offer computation, storage, and application hosting services and to provide coverage in several continents, offering service-level agreements (SLA)-backed performance and uptime promises for their services. While these “clouds” are the natural evolution of traditional data centers, they are distinguished by exposing resources (computation, data/storage, and applications) as standards-based Web services and following a “utility” pricing model where customers are charged based on their utilization of computational resources, storage, and transfer of data. They offer subscription-based access to infrastructure, platforms, and applications that are popularly referred to as IaaS (Infrastructure as a Service), PaaS (Platform as a Service), and SaaS (Software as a Service). While these emerging services have increased interoperability and usability and reduced the cost of computation, application hosting, and content storage and delivery by several orders of magnitude, there is significant complexity involved in ensuring that applications and services can scale as needed to achieve consistent and reliable operation under peak loads.

Currently, expert developers are required to implement cloud services. Cloud vendors, researchers, and practitioners alike are working to ensure that potential users are educated about the benefits of cloud computing and the best way to harness the full potential of the cloud. 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 illuminates these issues by introducing the reader with the cloud computing paradigm. The book provides case studies of numerous existing compute, storage, and application cloud services and illustrates capabilities and limitations of current providers of cloud computing services. This allows the reader to understand the mechanisms needed to harness cloud computing in their own respective endeavors. Finally, many open research problems 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. We believe the book to serve as a reference for larger audience such as systems architects, practitioners, developers, new researchers, and graduate-level students. This book also comes with an associated Web site (hosted at <http://www.manjrasoft.com/CloudBook/>) containing pointers to advanced on-line resources.

ORGANIZATION OF THE BOOK

This book contains chapters authored by several leading experts in the field of cloud computing. The book is presented in a coordinated and integrated manner starting with the fundamentals and followed by the technologies that implement them.

The content of the book is organized into six parts:

- I. Foundations
- II. Infrastructure as a Service (IaaS)
- III. Platform and Software as a Service (PaaS/SaaS)
- IV. Monitoring and Management
- V. Applications
- VI. Governance and Case Studies

Part I presents fundamental concepts of cloud computing, charting their evolution from mainframe, cluster, grid, and utility computing. Delivery models such as Infrastructure as a Service, Platform as a Service, and Software as a Service are detailed, as well as deployment models such as Public, Private, and Hybrid Clouds. It also presents models for migrating applications to cloud environments.

Part II covers Infrastructure as a Service (IaaS), from enabling technologies such as virtual machines and virtualized storage, to sophisticated mechanisms for securely storing data in the cloud and managing virtual clusters.

Part III introduces Platform and Software as a Service (PaaS/IaaS), detailing the delivery of cloud hosted software and applications. The design and operation of sophisticated, auto-scaling applications and environments are explored.

Part IV presents monitoring and management mechanisms for cloud computing, which becomes critical as cloud environments become more complex and interoperable. Architectures for federating cloud computing resources are explored, as well as service level agreement (SLA) management and performance prediction.

Part V details some novel applications that have been made possible by the rapid emergence of cloud computing resources. Best practices for architecting cloud applications are covered, describing how to harness the power of loosely coupled cloud resources. The design and execution of applications that leverage

cloud resources such as massively multiplayer online game hosting, content delivery and mashups are explored.

Part VI outlines the organizational, structural, regulatory and legal issues that are commonly encountered in cloud computing environments. Details on how companies can successfully prepare and transition to cloud environments are explored, as well as achieving production readiness once such a transition is completed. Data security and legal concerns are explored in detail, as users reconcile moving their sensitive data and computation to cloud computing providers.

Rajkumar Buyya

The University of Melbourne and Manjrasoft Pty Ltd., Australia

James Broberg

The University of Melbourne, Australia

Andrzej Goscinski

Deakin University, Australia

ACKNOWLEDGMENTS

First and foremost, we are grateful to all the contributing authors for their time, effort, and understanding during the preparation of the book.

We thank Professor Albert Zomaya, editor of the Wiley book series on parallel and distributed computing, for his enthusiastic support and guidance during the preparation of book and enabling us to easily navigate through Wiley's publication process.

We would like to thank members of the book Editorial Advisory Board for their guidance during the preparation of the book. The board members are: Dr. Geng Lin (CISCO Systems, USA), Prof. Manish Parashar (Rutgers: The State University of New Jersey, USA), Dr. Wolfgang Gentzsch (Max-Planck-Gesellschaft, München, Germany), Prof. Omer Rana (Cardiff University, UK), Prof. Hai Jin (Huazhong University of Science and Technology, China), Dr. Simon See (Sun Microsystems, Singapore), Dr. Greg Pfister (IBM, USA (retired)), Prof. Ignacio M. Llorente (Universidad Complutense de Madrid, Spain), Prof. Geoffrey Fox (Indiana University, USA), and Dr. Walfredo Cirne (Google, USA).

All chapters were reviewed and authors have updated their chapters to address review comments. We thank members of the Melbourne CLOUDS Lab for their time and effort in peer reviewing of chapters.

Raj would like to thank his family members, especially Smrithi, Soumya, and Radha Buyya, for their love, understanding, and support during the preparation of the book. James would like to thank his wife, Amy, for her love and support. Andrzej would like to thank his wife, Teresa, for her love and support.

Finally, we would like to thank the staff at Wiley, particularly, Simone Taylor (Senior Editor, Wiley), Michael Christian (Editorial Assistant, Wiley), and S. Nalini (MPS Limited, a Macmillan Company, Chennai, India). They were wonderful to work with!

R.B.

J.B.

A.G.

CONTRIBUTORS

MATTHIAS ASSEL, High Performance Computing Center Stuttgart (HLRS),
University of Stuttgart, 70550 Stuttgart, Germany

ROCCO AVERSA, Department of Information Engineering, Second University of
Naples, 81031 Aversa (CE), Italy

SUMIT BOSE, Unisys Research Center, Bangalore, India - 560025

JANINE ANTHONY BOWEN, ESQ., McKenna Long & Aldridge LLP, Atlanta, GA
30308, USA

DAVID BREITGAND, IBM Haifa Research Lab, Haifa University Campus, 31095,
Haifa, Israel

JAMES BROBERG, Department of Computer Science and Software Engineering,
The University of Melbourne, Parkville, Melbourne, VIC 3010, Australia

MICHAEL BROCK, School of Information Technology, Deakin University,
Geelong, Victoria 3217, Australia

RAJKUMAR BUYYA, Department of Computer Science and Software Engineering,
The University of Melbourne, Parkville, Melbourne, VIC 3010, Australia

HAIJUN CAO, School of Computer Science and Technology, Huazhong Uni-
versity of Science and Technology, Wuhan, 430074, China

WAI-KIT CHEAH, Advanced Customer Services, Oracle Corporation (S) Pte
Ltd., Singapore 038986

YU CHEN, Department of Electrical and Computer Engineering, State Uni-
versity of New York—Binghamton, Binghamton, NY 13902

XINGCHEN CHU, Department of Computer Science and Software Engineering,
The University of Melbourne, Parkville, Melbourne, VIC 3010, Australia

BENIAMINO DI MARTINO, Department of Information Engineering, Second
University of Naples, 81031 Aversa (CE), Italy

TARIQ ELLAHI, SAP Research Belfast, BT3 9DT, Belfast, United Kingdom

MOHAMED A. EL-REFAEY, Arab Academy for Science, Technology and Maritime Transport, College of Computing and Information Technology, Cairo, Egypt

JUN FENG, Department of Electrical and Computer Engineering, State University of New York—Binghamton, Binghamton, NY 13902

FERMÍN GALÁN, Telefónica I + D, Emilio Vargas, 6. 28043 Madrid, Spain

ALEX GALIS, University College London, Department of Electronic and Electrical Engineering, Torrington Place, London WC1E 7JE, United Kingdom

ANDRZEJ GOSCINSKI, School of Information Technology, Deakin University, Geelong, Victoria 3217, Australia

DAVID HADAS, IBM Haifa Research Lab, Haifa University Campus, 31095, Haifa, Israel

BENOIT HUDZIA, SAP Research Belfast, BT3 9DT, Belfast, United Kingdom

SHADI IBRAHIM, School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan, 430074, China

ALEXANDRU IOSUP, Electrical Engineering, Mathematics and Computer Science Department, Delft University of Technology, 2628 CD, Delft, The Netherlands

SHANTENU JHA, Center for Computation and Technology and Department of Computer Science, Louisiana State University, Baton Rouge, LA 70803

HAI JIN, School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan, 430074, China

DILEBAN KARUNAMOORTHY, Department of Computer Science and Software Engineering, The University of Melbourne, Parkville, Melbourne, VIC 3010, Australia

HENRY KASIM, HPC and Cloud Computing Center, Oracle Corporation (S) Pte Ltd, #18-01 Suntec Tower Four, Singapore 038986

DANIEL S. KATZ, Computation Institute, University of Chicago, Chicago, Illinois 60637

HYUNJOO KIM, Department of Electrical and Computer Engineering, Rutgers, The State University of New Jersey, New Brunswick, NJ

ALEXANDER KIPP, High Performance Computing Center Stuttgart (HLRS), University of Stuttgart, 70550 Stuttgart, Germany

WEI-SHINN KU, Department of Computer Science and Software Engineering, Auburn University, AL 36849

ROBERT LAM, School of Information and Communication Technologies SAIT
Polytechnic, Calgary, Canada T2M 0L4

LARS LARSSON, Department of Computing Science, University Umea, Sweden

ELIEZER LEVY, SAP Research SRC Ra'anana, Ra'anana 43665; Israel

HUI LI, SAP Research Karlsruhe, Vincenz-Priessnitz-Strasse, 176131 Karlsruhe, Germany

MAIK A. LINDNER, SAP Research Belfast, BT3 9DT, Belfast, United Kingdom

PU LIU, IBM Endicott Center, New York, NY

IGNACIO M. LLORENTE, Distributed Systems Architecture Research Group,
Departamento de Arquitectura de Computadores y Automática, Facultad de
Informática, Universidad Complutense de Madrid, 28040 Madrid, Spain

ANDRE LUCKOW, Center for Computation and Technology, Louisiana State
University, Baton Rouge, LA, 70803

GANESAN MALAIYANDISAMY, SETLabs, Infosys Technologies Limited, Electronics City, Bangalore, India, 560100

ALESSANDRO MARASCHINI, ElsasDatamat spa, Rome, Italy

PHILIPPE MASSONET, CETIC, B-6041 Charleroi, Belgium

MICHAEL MATTESS, Department of Computer Science and Software Engineering,
The University of Melbourne, Parkville, Melbourne, VIC 3010,
Australia

ANDRE MERZKY, Center for Computation and Technology, Louisiana State
University, Baton Rouge, LA, 70803

T. S. MOHAN, Infosys Technologies Limited, Electronics City, Bangalore,
India, 560100

RUBÉN S. MONTERO, Distributed Systems Architecture Research Group,
Departamento de Arquitectura de Computadores, y Automática, Facultad de
Informática, Universidad Complutense de Madrid, 28040 Madrid, Spain

SUSAN MORROW, Avoco Secure, London W1S 2LQ, United Kingdom

SRIDHAR MURTHY, Infosys Technologies Limited, Electronics City, Bangalore,
India, 560100

VLAD NAE, Institute of Computer Science, University of Innsbruck, Technikerstraße 21a, A-6020 Innsbruck, Austria

KENNETH NAGIN, IBM Haifa Research Lab, Haifa University Campus, 31095,
Haifa, Israel

SURAJ PANDEY, Department of Computer Science and Software Engineering,
The University of Melbourne, Parkville, Melbourne, VIC 3010, Australia

MANISH PARASHAR, Department of Electrical and Computer Engineering,
Rutgers, The State University of New Jersey, New Jersey, USA.

ANJANEYULU PASALA, SETLabs, Infosys Technologies Limited, Electronics
City, Bangalore, India, 560100

MICHAEL PAULY, T-Systems, Aachen, Germany

RADU PRODAN, Institute of Computer Science, University of Innsbruck, A-6020
Innsbruck, Austria

LI QI, School of Computer Science and Technology, Huazhong University of
Science and Technology, Wuhan, 430074, China

DHEEPAK R A, SETLabs, Infosys Technologies Limited, Electronics City,
Bangalore, India, 560100

PETHURU RAJ, Robert Bosch India, Bangalore 560068, India

MASSIMILIANO RAK, Department of Information Engineering, Second University
of Naples, 81031 Aversa (CE), Italy

PHILIP ROBINSON, SAP Research Belfast, BT3 9DT, Belfast, United Kingdom

BENNY ROCHWERGER, IBM Haifa Research Lab, Haifa University Campus,
31095, Haifa, Israel

LUTZ SCHUBERT, High Performance Computing Center Stuttgart (HLRS),
University of Stuttgart, 70550 Stuttgart, Germany

XUANHUA SHI, School of Computer Science and Technology, Huazhong
University of Science and Technology, Wuhan, 430074, China

BORJA SOTOMAYOR, Department of Computer Science, University of Chicago,
Chicago, IL

KATERINA STAMOU, Department of Computer Science, Louisiana State
University, Baton Rouge, LA, 70803

ZHOU SU, Department of Computer Science, Graduate School of Science and
Engineering, Waseda University, Japan

JINESH VARIA, Amazon Web Services, Seattle, WA 98109

CONSTANTINO VÁZQUEZ, Facultad de Informática, Universidad Complutense
de Madrid, 28040 Madrid, Spain

CHRISTIAN VECCHIOLA, Department of Computer Science and Software
Engineering, The University of Melbourne, Parkville, Melbourne,
VIC 3010, Australia

SALVATORE VENTICINQUE, Department of Information Engineering, Second University of Naples, 81031 Aversa (CE), Italy

UMBERTO VILLANO, Department of Engineering, University of Sannio, 82100 Benevento, Italy

MASSIMO VILLARI, Department. of Mathematics Faculty of Engineering, University of Messina, 98166 Messina, Italy

WILLIAM VOORSLUYS, Department of Computer Science and Software Engineering, The University of Melbourne, Parkville, Melbourne, VIC 3010, Australia

STEFAN WESNER, High Performance Computing Center Stuttgart (HLRS), University of Stuttgart, 70550 Stuttgart, Germany

YARON WOLFSTHAL, IBM Haifa Research Lab, Haifa University Campus, 31095, Haifa, Israel

SONG WU, School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan, 430074, China

CONTENTS

PREFACE	XV
ACKNOWLEDGMENTS	XIX
CONTRIBUTORS	XXI

PART I FOUNDATIONS	1
--------------------	---

1 Introduction to Cloud Computing	3
--	----------

William Voorsluys, James Broberg, and Rajkumar Buyya

- 1.1 Cloud Computing in a Nutshell / 3
- 1.2 Roots of Cloud Computing / 5
- 1.3 Layers and Types of Clouds / 13
- 1.4 Desired Features of a Cloud / 16
- 1.5 Cloud Infrastructure Management / 17
- 1.6 Infrastructure as a Service Providers / 26
- 1.7 Platform as a Service Providers / 31
- 1.8 Challenges and Risks / 34
- 1.9 Summary / 37
- References / 37

2 Migrating into a Cloud	43
---------------------------------	-----------

T. S. Mohan

- 2.1 Introduction / 43
- 2.2 Broad Approaches to Migrating into the Cloud / 48
- 2.3 The Seven-Step Model of Migration into a Cloud / 51
- 2.4 Conclusions / 54
- Acknowledgments / 55
- References / 55

3 Enriching the ‘Integration as a Service’ Paradigm for the Cloud Era **57**

Pethuru Raj

- 3.1 An Introduction / 57
- 3.2 The Onset of Knowledge Era / 59
- 3.3 The Evolution of SaaS / 59
- 3.4 The Challenges of SaaS Paradigm / 61
- 3.5 Approaching the SaaS Integration Enigma / 63
- 3.6 New Integration Scenarios / 67
- 3.7 The Integration Methodologies / 69
- 3.8 SaaS Integration Products and Platforms / 72
- 3.9 SaaS Integration Services / 80
- 3.10 Businesses-to-Business Integration (B2Bi) Services / 84
- 3.11 A Framework of Sensor—Cloud Integration [3] / 89
- 3.12 SaaS Integration Appliances / 94
- 3.13 Conclusion / 95
 - References / 95

4 The Enterprise Cloud Computing Paradigm **97**

Tariq Ellahi, Benoit Hudzia, Hui Li, Maik A. Lindner, and Philip Robinson

- 4.1 Introduction / 97
- 4.2 Background / 98
- 4.3 Issues for Enterprise Applications on the Cloud / 103
- 4.4 Transition Challenges / 106
- 4.5 Enterprise Cloud Technology and Market Evolution / 108
- 4.6 Business Drivers Toward a Marketplace for Enterprise Cloud Computing / 112
- 4.7 The Cloud Supply Chain / 115
- 4.8 Summary / 117
 - Acknowledgments / 117
 - References / 118

PART II INFRASTRUCTURE AS A SERVICE (IAAS) **121**

5 Virtual Machines Provisioning and Migration Services **123**

Mohamed El-Refaey

- 5.1 Introduction and Inspiration / 123

5.2	Background and Related Work /	124
5.3	Virtual Machines Provisioning and Manageability /	130
5.4	Virtual Machine Migration Services /	132
5.5	VM Provisioning and Migration in Action /	136
5.6	Provisioning in the Cloud Context /	145
5.7	Future Research Directions /	151
5.8	Conclusion /	154
	References /	154

6 On the Management of Virtual Machines for Cloud Infrastructures **157**

Ignacio M. Llorente, Rubén S. Montero, Borja Sotomayor, David Breitgand, Alessandro Maraschini, Eliezer Levy, and Benny Rochwerger

6.1	The Anatomy of Cloud Infrastructures /	158
6.2	Distributed Management of Virtual Infrastructures /	161
6.3	Scheduling Techniques for Advance Reservation of Capacity /	166
6.4	Capacity Management to meet SLA Commitments /	172
6.5	Conclusions and Future Work /	185
	Acknowledgments /	186
	References /	187

7 Enhancing Cloud Computing Environments Using a Cluster as a Service **193**

Michael Brock and Andrzej Goscinski

7.1	Introduction /	193
7.2	Related Work /	194
7.3	RVWS Design /	197
7.4	Cluster as a Service: The Logical Design /	202
7.5	Proof of Concept /	212
7.6	Future Research Directions /	218
7.7	Conclusion /	219
	References /	219

8 Secure Distributed Data Storage in Cloud Computing **221**

Yu Chen, Wei-Shinn Ku, Jun Feng, Pu Liu, and Zhou Su

8.1	Introduction /	221
8.2	Cloud Storage: from LANs TO WANs /	222
8.3	Technologies for Data Security in Cloud Computing /	232

- 8.4 Open Questions and Challenges / 242
- 8.5 Summary / 246
- References / 246

PART III PLATFORM AND SOFTWARE AS A SERVICE (PaaS/IaaS) 249

9 Aneka—Integration of Private and Public Clouds 251

Christian Vecchiola, Xingchen Chu, Michael Mattess, and Rajkumar Buyya

- 9.1 Introduction / 251
- 9.2 Technologies and Tools for Cloud Computing / 254
- 9.3 Aneka Cloud Platform / 257
- 9.4 Aneka Resource Provisioning Service / 259
- 9.5 Hybrid Cloud Implementation / 262
- 9.6 Visionary thoughts for Practitioners / 269
- 9.7 Summary and Conclusions / 271
- Acknowledgments / 272
- References / 273

10 CometCloud: An Autonomic Cloud Engine 275

Hyunjoo Kim and Manish Parashar

- 10.1 Introduction / 275
- 10.2 CometCloud Architecture / 276
- 10.3 Autonomic Behavior of CometCloud / 280
- 10.4 Overview of CometCloud-based Applications / 286
- 10.5 Implementation and Evaluation / 287
- 10.6 Conclusion and Future Research Directions / 295
- Acknowledgments / 295
- References / 296

11 T-Systems' Cloud-Based Solutions for Business Applications 299

Michael Pauly

- 11.1 Introduction / 299
- 11.2 What Enterprises Demand of Cloud Computing / 300
- 11.3 Dynamic ICT Services / 302
- 11.4 Importance of Quality and Security in Clouds / 305