

Implementing and Developing Cloud Computing Applications

DAVID E.Y. SARNA



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business
AN AUERBACH BOOK

Auerbach Publications
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2011 by Taylor and Francis Group, LLC
Auerbach Publications is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed in the United States of America on acid-free paper
10 9 8 7 6 5 4 3 2 1

International Standard Book Number: 978-1-4398-3082-6 (Hardback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, micro-filming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data

Sarna, David E. Y.

Implementing and developing cloud computing applications / by David E.Y. Sarna.
p. cm.

Summary: "Major enterprises and small start-ups are beginning to embrace cloud computing for the scalability and reliability that cloud vendors can provide. This book demonstrates how to implement robust and highly scalable cloud computing applications. Filled with comparative charts and decision trees to help navigate the many implementation alternatives, the author describes the major available commercial offerings and guides readers in choosing the best combination of platform, tools, and services for a small, growing start-up or an established enterprise. Aimed at software developers and their managers, the text details cloud development environments, lifecycles, and project management"-- Provided by publisher.

Includes bibliographical references and index.

ISBN 978-1-4398-3082-6 (hardback)

1. Cloud computing. 2. Business--Data processing. I. Title.

QA76.585.S37 2010

006.7'8--dc22

2010037120

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the Auerbach Web site at
<http://www.auerbach-publications.com>

Contents

| | |
|---|--------------|
| About the Author | xix |
| Preface | xxi |
| Author's Acknowledgements | xxiii |
| Executive Summary | xxv |
| Cloud Computing is a True Paradigm Shift | xxv |
| From Do It Yourself to Public Cloud—A Continuum | xxv |
| Cloud Computing: Is It Old Mainframe Bess in a New Dress? | xxv |
| Moving Into and Around the Clouds and Efforts at Standardization | xxvi |
| Cloud Economics and Capacity Management | xxvii |
| Demystifying the Cloud: A Case Study Using Amazon's Cloud Services (AWS) | xxvii |
| Virtualization: Open Source and VMware | xxvii |
| Securing the Cloud: Reliability, Availability, and Security | xxviii |
| Scale and Reuse: Standing on the Shoulders of Giants | xxviii |
| Windows Azure | xxix |
| Google in the Cloud | xxix |
| Enterprise Cloud Vendors | xxx |
| Cloud Service Providers | xxxi |
| Practice Fusion Case Study | xxxi |
| Support and Reference Materials | xxxi |
| | |
| Chapter 1 Cloud Computing is a True Paradigm Shift | 1 |
| Chapter Overview | 1 |
| 1.1 Introduction | 1 |
| 1.2 What is Cloud Computing? | 2 |
| 1.3 We're Using Cloud Computing Already | 4 |

| | | |
|------------------|--|-----------|
| 1.3.1 | Electronic Faxing | 4 |
| 1.3.2 | Voice in the Cloud | 4 |
| 1.3.3 | Commerce in the Cloud | 5 |
| 1.3.4 | Distributed Hosting in the Cloud | 5 |
| 1.3.5 | Accounting and Online Banking in the Cloud | 5 |
| 1.4 | New in the Cloud | 6 |
| 1.5 | Other Cloud Applications | 8 |
| 1.6 | What about the Enterprise? | 8 |
| 1.7 | More to Come | 9 |
| | Summary | 9 |
| Chapter 2 | From Do It Yourself to Public Cloud—A Continuum | 11 |
| | Chapter Objectives | 11 |
| 2.1 | A Brief History | 11 |
| 2.2 | Virtualization | 11 |
| 2.3 | Remote Hosting | 13 |
| 2.4 | Hosting Services | 14 |
| 2.5 | Cloud Computing Defined | 15 |
| 2.5.1 | Essential Characteristics | 15 |
| 2.5.2 | Cloud Service Models | 16 |
| 2.5.3 | Deployment Models | 17 |
| 2.5.4 | Cloud Software | 17 |
| 2.5.5 | Advantages of Cloud Computing | 17 |
| 2.6 | The Divisive Issue of Multitenancy | 18 |
| 2.7 | Advantages of Cloud Hosting Over Remote Hosting | 19 |
| 2.8 | The Battle Over Public and Private Clouds | 20 |
| 2.9 | Then Came the Internet | 22 |
| 2.10 | The Argument for Private Clouds | 25 |
| 2.11 | Hybrid Solutions | 25 |
| 2.11.1 | Hybrid Cloud—Not Really | 25 |
| 2.11.2 | The Hybrid Cloud Model | 26 |
| 2.12 | Cloud Computing for Development | 26 |
| 2.13 | Eucalyptus—Open Source Software Supporting Hybrid Solutions | 27 |
| | Eucalyptus Features and Benefits | 28 |
| 2.14 | Microsoft Also Endorses the Hybrid Model | 30 |
| | Summary | 30 |

| | | |
|------------------|---|-----------|
| Chapter 3 | Cloud Computing: Is It Old Mainframe Bess in a New Dress? | 33 |
| | Chapter Overview | 33 |
| | 3.1 Déjà Vu? | 33 |
| | 3.2 Not Remote Hosting | 34 |
| | 3.3 Cloud Computing is Maturing Quickly | 34 |
| | Cloud Computing is Not a New Concept | 35 |
| | 3.4 Vision of Computer Utility | 36 |
| | 3.5 Desktop Virtualization | 40 |
| | 3.6 PaaS: Platform as a Service | 41 |
| | 3.7 SaaS Applications | 42 |
| | 3.8 Force.com and Standing on Tall Shoulders | 43 |
| | 3.9 Other Popular SaaS Applications | 45 |
| | 3.10 The Holy Grail of Computing | 46 |
| | 3.11 SaaS 2.0 | 46 |
| | Summary | 47 |
| | | |
| Chapter 4 | Moving Into and Around the Clouds and Efforts at Standardization | 49 |
| | 4.1 Portable Software | 49 |
| | 4.2 Openness, Linux, and Apache | 50 |
| | 4.3 Closed Architectures | 51 |
| | 4.4 Legacy Applications and Migration to the Cloud | 51 |
| | 4.5 Preventing Vendor Lock-In as You Migrate to the Cloud | 52 |
| | 4.5.1 What to do? | 52 |
| | 4.5.2 More Questions | 53 |
| | 4.5.3 Comparing Costs | 55 |
| | 4.6 Narrowing the Choices | 55 |
| | 4.7 Scripting Languages | 55 |
| | Microsoft Visual Studio and Other Development Environments | 56 |
| | 4.8 Cloud Software | 57 |
| | Eucalyptus Enterprise Edition | 57 |
| | 4.9 Cloud-Optimized Linux | 58 |
| | 4.9.1 CloudLinux | 58 |
| | 4.9.2 Peppermint | 60 |
| | 4.9.3 Ubuntu's Cloud Strategy | 61 |
| | 4.10 CohesiveFT | 63 |
| | 4.10.1 Elastic Server | 63 |

| | | |
|------------------|---|-----------|
| 4.10.2 | VPN-Cubed IPsec to Cloud for Hybrid and Cloud-to-Cloud Applications | 65 |
| 4.11 | Zend | 66 |
| 4.12 | Abiquo | 68 |
| 4.12.1 | Abiquo's Vision | 68 |
| 4.12.2 | Management Benefits | 70 |
| 4.12.3 | Other Abiquo Benefits | 71 |
| 4.13 | 3Tera | 72 |
| 4.14 | Elastra | 73 |
| | Elastra for Amazon Web Services | 75 |
| 4.15 | RightScale | 75 |
| | ServerTemplates versus Machine Images | 77 |
| 4.16 | Today is Like 1973 | 78 |
| 4.17 | Interclouding, Standards, and VMware's Focus on Open PaaS | 79 |
| 4.18 | DMTF | 81 |
| | OCSI Standardization Efforts | 82 |
| 4.19 | The Problem of Metering | 83 |
| 4.20 | Remember the Dodo Bird | 83 |
| 4.21 | Cloud Broker | 86 |
| | Interclouding, DTMF and OVFS | 88 |
| 4.22 | Product Offerings | 89 |
| | Summary | 89 |
| Chapter 5 | Cloud Economics and Capacity Management | 91 |
| | Overview | 91 |
| | Choices Abound | 91 |
| 5.1 | Capacity Planning: A Play in Three Acts | 92 |
| | Capacity Management: An Old-New Technique | 93 |
| 5.2 | Queueing Theory | 94 |
| 5.3 | Queueing and Response Time | 94 |
| 5.4 | Historical Note on Computer Capacity Management | 97 |
| 5.5 | Evidence-Based Decision Making | 98 |
| 5.6 | Instrumentation (Measuring Resource Consumption) | 98 |
| 5.6.1 | First, Get Your Business Needs Down Clearly | 99 |
| 5.6.2 | What Technologists Must Know to Manage Performance and Capacity | 100 |

| | | |
|------------------|---|------------|
| 5.7 | Managers Are from Mars, Technologists Are from Venus | 102 |
| 5.8 | Bottlenecks | 102 |
| 5.9 | Getting the Facts | 103 |
| 5.10 | Strategies for Capacity Planning | 104 |
| 5.11 | Critical Success Factors (CSF) and Best Practices | 104 |
| 5.12 | Key Volume Indicators | 107 |
| 5.12.1 | Types of Workloads | 109 |
| 5.12.2 | Determining KVLs for an Application | 109 |
| 5.12.3 | Monitoring and Improving Forecastability | 109 |
| 5.12.4 | Standard Costs | 109 |
| 5.12.5 | Determining Whether Resources are Adequate for Projected Demand | 110 |
| 5.12.6 | New Applications | 110 |
| 5.12.7 | Accuracy of Forecasts | 110 |
| 5.12.8 | Queueing Models | 111 |
| 5.12.9 | Make or Buy a Cloud | 112 |
| | Summary | 113 |
| | | |
| Chapter 6 | Demystifying the Cloud: A Case Study Using Amazon's Cloud Services (AWS) | 115 |
| 6.1 | Why Amazon? | 116 |
| 6.1.1 | Amazon is Just an Illustrative Example | 117 |
| 6.1.2 | Let's Do It Now | 118 |
| 6.1.3 | Amazon S3 Functionality | 118 |
| 6.2 | Using Amazon S3 | 121 |
| 6.3 | Gladinet Puts a Desktop Face on S3 | 122 |
| 6.3.1 | Use Cases for Using Virtual Drives | 123 |
| 6.3.2 | Beyond One-on-One: Use a Cloud Gateway | 125 |
| 6.3.3 | Benefits of Using Cloud Gateway | 125 |
| 6.4 | Moving A Simple Application to the Cloud | 126 |
| 6.5 | Step One: Move Static Content to S3 | 126 |
| 6.5.1 | Using CloudFront | 127 |
| 6.5.2 | Other Tools for Moving Content to S3 | 127 |
| 6.5.3 | Using Amazon S3 with Firefox S3Fox | 128 |
| 6.6 | Step Two: Move Web Servers and Backend Servers to EC2 | 130 |
| | The Web Servers | 130 |
| 6.7 | Moving The Database | 136 |

| | | |
|---|---|------------|
| 6.8 | Using EBS for MySQL | 137 |
| 6.9 | Accessing Public Data | 139 |
| 6.10 | Crawl, Walk, Run | 139 |
| 6.11 | Scaling and Monitoring: Taking Advantage of Cloud Services | 140 |
| | Monitoring | 140 |
| 6.12 | Eucalyptus Enterprise Edition | 142 |
| | Key Features and Functionality | 143 |
| 6.13 | Nimbula—Roll Your Own Private EC2 | 143 |
| | Summary | 145 |
| Chapter 7 Virtualization: Open Source and VMware | | 147 |
| | Overview | 147 |
| | Virtualization Is an Old Story | 148 |
| 7.1 | The Hypervisor is the Secret Sauce | 148 |
| 7.2 | KVM | 149 |
| 7.3 | Xen | 149 |
| 7.4 | QEMU | 150 |
| 7.5 | Comparing KVM and Xen | 150 |
| 7.6 | Comparing KVM and QEMU | 151 |
| 7.7 | Parallels | 151 |
| 7.8 | A Unique Hypervisor: Microsoft Azure and Hyper-V | 151 |
| | 7.8.1 Managing a Virtualized Infrastructure | 152 |
| | 7.8.2 Monitoring and Management | 153 |
| | 7.8.3 Commercial Virtualization Offerings | 154 |
| | 7.8.4 Citrix | 154 |
| | 7.8.5 VMware | 155 |
| 7.9 | EMC's VPLEX and VMware | 160 |
| 7.10 | VMware Partners with Salesforce.com and Google | 162 |
| 7.11 | VMforce | 163 |
| 7.12 | VMware and Google | 164 |
| | 7.12.1 Spring for AppEngine | 165 |
| | 7.12.2 Spring Insight and Google Speed Tracer | 166 |
| 7.13 | Eucalyptus and VMware | 166 |
| | Recent VM Acquisitions | 167 |
| 7.14 | OpenStack | 168 |
| | Summary | 169 |

| | |
|---|------------|
| Chapter 8 Securing the Cloud: Reliability, Availability, and Security | 171 |
| Overview | 171 |
| 8.1 The FUDD Factor | 171 |
| 8.2 Leakage | 172 |
| 8.3 Not All Threats Are External | 172 |
| 8.4 Virtualization Is Inherently More Secure | 172 |
| 8.5 Virtualization is Not Enough | 173 |
| 8.6 The Best Security May Be Unavailable for (In-House) Private Clouds | 173 |
| 8.7 Providers Make Security Their Business | 173 |
| 8.8 Cloud Security Providers Employ a Hierarchy of Containment Strategies | 174 |
| 8.9 How a Denial of Service Attack Is Carried Out | 176 |
| 8.10 Cloud Computing Offers Enhanced Defenses for Thwarting DoS Attacks | 177 |
| 8.11 Who's Responsible? Amazon's AWS EC2 and Salesforce.com Compared | 178 |
| 8.12 VMForce.com | 178 |
| 8.13 Azure and Security | 179 |
| 8.14 OASIS and SPLM | 179 |
| 8.15 Trust, but Verify | 179 |
| 8.16 Independent Third-Party Validation is a Prerequisite | 180 |
| 8.17 Standards and Vendor Selection | 180 |
| 8.17.1 ISO 27001 | 180 |
| 8.17.2 SAS 70 (Statement on Auditing Standards No. 70): Service Organizations | 182 |
| 8.17.3 Type I and Type II Audits | 183 |
| 8.18 SAS 70 and Cloud Computing | 184 |
| 8.19 Cloud Security Alliance | 186 |
| 8.20 SysTrust Certification | 186 |
| 8.21 Cloud Security Alliance Working Toward Cloud-Specific Certifications | 187 |
| CSA Goes Beyond SAS 70 and ISO 27001 | 189 |
| 8.22 Customers Demand Better Proof | 190 |
| 8.23 CloudAudit | 191 |
| Summary | 192 |
| | |
| Chapter 9 Scale and Reuse: Standing on the Shoulders of Giants | 193 |
| 9.1 Objectives | 193 |

| | | |
|-------------------|---|------------|
| 9.2 | Cloud Computing on One Foot | 193 |
| 9.3 | Just Make the Call; Let Google Do It | 194 |
| 9.4 | Hardware Reuse | 195 |
| 9.5 | Scale and Reuse (Use it or Lose it) | 196 |
| 9.6 | Service-Oriented Architecture | 197 |
| 9.7 | Web 2.0 | 199 |
| | Summary | 199 |
| Chapter 10 | Windows Azure | 201 |
| | Chapter Objectives | 201 |
| 10.1 | Back to the Future | 201 |
| 10.2 | But Windows had not kept pace. | 202 |
| 10.3 | Billionaire's Agita | 203 |
| 10.4 | Prologue to Windows Azure | 203 |
| 10.5 | Introducing Windows Azure | 204 |
| 10.6 | What is Windows Azure? | 205 |
| 10.7 | Microsoft's Secret Datacenter | 206 |
| 10.8 | Azure is an Open Platform | 207 |
| 10.9 | How does the Windows Azure SDK for PHP fit in? | 208 |
| 10.10 | Deployment Scenarios | 208 |
| 10.11 | Recent Enhancements | 209 |
| 10.12 | Open Source Embraced | 210 |
| 10.13 | Azure: IaaS or PaaS? | 211 |
| 10.14 | Competition with Salesforce.com | 212 |
| 10.15 | Salesforce.com is Microsoft's Real Concern | 212 |
| 10.16 | Preparing for Midori | 213 |
| 10.17 | F# and Midori | 214 |
| 10.18 | An Azure Tie-In-to Midori? | 215 |
| 10.19 | Azure Pricing | 216 |
| 10.20 | Microsoft Intune: A New SaaS-based Service | 216 |
| 10.21 | Advanced Management Tools | 218 |
| 10.22 | Intune is Microsoft-Centric | 219 |
| 10.23 | Microsoft Resources | 219 |
| | Summary | 220 |
| Chapter 11 | Google in the Cloud | 221 |
| | Overview | 221 |
| 11.1 | Free is Good | 222 |
| 11.2 | Reaching Out to the Development Community | 222 |
| 11.3 | App Engine Cost Structure | 223 |