

Implementing and Developing Cloud Computing Applications

DAVID E.Y. SARNA



CRC Press
Taylor & Francis Group

AN AUERBACH BOOK

Implementing and Developing Cloud Computing Applications

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-13: 978-1-4398-3083-3 (Ebook-PDF)

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, microfilming, 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.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>
and the Auerbach Web site at
<http://www.auerbach-publications.com>

Dedication

*For my loving wife, Dr. Rachel C. Sarna
All that is Mine and Yours is Hers
B. Talmud Nedarim 50*

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 KVs 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 Nimble—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 |

| | | |
|-------|--|-----|
| 11.4 | Google Web Toolkit™ | 223 |
| 11.5 | Google Cloud Applications Built on GWT | 224 |
| 11.6 | Google Gears R.I.P. | 224 |
| 11.7 | Google Apps Script | 225 |
| 11.8 | What Is Google App Engine? | 226 |
| 11.9 | Google App Engine for Business | 228 |
| 11.10 | Collaboration with VMware | 229 |
| | Summary | 230 |

Chapter 12 Enterprise Cloud Vendors 233

| | | |
|--------|---|-----|
| | Overview | 233 |
| 12.1 | IBM | 233 |
| 12.1.1 | Consistent Development Experience | 234 |
| 12.1.2 | Eleven Cloud Computing Competency Centers Worldwide | 235 |
| 12.1.3 | IBM Cloud Resources | 239 |
| 12.1.4 | Recent IBM Cloud Press Releases | 239 |
| 12.2 | Amazon AWS | 239 |
| 12.2.1 | Amazon RDS | 239 |
| 12.2.2 | Amazon CloudWatch | 240 |
| 12.3 | Hewlett Packard | 241 |
| 12.3.1 | HP's SaaS Offerings | 242 |
| 12.3.2 | HP Business Service Management | 242 |
| 12.3.3 | HP's View of the Hybrid World | 243 |
| 12.3.4 | HP BSM 9.0 | 243 |
| 12.3.5 | HP Business Availability Center | 244 |
| 12.3.6 | HP's Test Data Management | 245 |
| 12.3.7 | HP Partners with Enomaly | 246 |
| 12.3.8 | HP's Alliance With Microsoft | 247 |
| 12.3.9 | HP Resources | 247 |
| 12.4 | Oracle (Sun) | 248 |
| 12.4.1 | Oracle and Sun | 251 |
| 12.4.2 | Oracle Resources | 253 |
| 12.5 | CA Technologies | 253 |
| 12.5.1 | Partnership with NetApp | 256 |
| 12.5.2 | CA Resources | 257 |
| 12.6 | Unisys | 257 |
| 12.6.1 | Unisys Stealth Security | 257 |
| 12.6.2 | Unisys Secure Cloud Solution | 258 |
| 12.6.3 | Unisys Secure Private Cloud Solution | 259 |
| 12.6.4 | Unisys ClearPath Cloud Solutions | 259 |

| | | |
|--------|--------------------------------------|-----|
| 12.6.5 | Unisys Cloud Transformation Services | 260 |
| 12.6.6 | Unisys Resources | 260 |
| 12.6.7 | Unisys Partnerships | 260 |
| 12.7 | Cloud Research | 261 |
| | Summary | 261 |

Chapter 13 Cloud Service Providers 263

| | | |
|------|---------------------------------------|-----|
| | Overview | 263 |
| 13.1 | Comprehensive Cloud Service Providers | 263 |
| | Joyent | 263 |
| 13.2 | IaaS Providers | 264 |
| | 13.2.1 Rackspace | 264 |
| | 13.2.2 GoGrid | 264 |
| | 13.2.3 ElasticHosts | 264 |
| | 13.2.4 SymetriQ | 265 |
| 13.3 | PaaS Providers | 266 |
| | 13.3.1 AT&T | 266 |
| | 13.3.2 Terremark | 266 |
| | 13.3.3 EngineYard | 267 |
| 13.4 | SaaS Providers | 268 |
| | 13.4.1 NetSuite | 268 |
| | 13.4.2 Intuit | 268 |
| | 13.4.3 Intacct | 268 |
| | 13.4.4 FinancialForce.com | 269 |
| | 13.4.5 Coupa Software | 270 |
| | 13.4.6 AT&T | 271 |
| 13.5 | Specialized Cloud Software Providers | 271 |
| | 13.5.1 Appistry | 271 |
| | 13.5.2 BMS Software | 271 |
| | 13.5.3 Nasuni | 273 |
| | Summary | 274 |

Chapter 14 Practice Fusion Case Study 275

| | | |
|------|--------------------------------|-----|
| 14.1 | Overview | 275 |
| 14.2 | Practice Fusion | 275 |
| 14.3 | Non-Trivial, Maybe Life-Saving | 276 |
| 14.4 | Typical User | 278 |
| 14.5 | Practice Fusion Resources | 280 |
| 14.6 | Summary | 280 |

| | | |
|--|--|------------|
| Chapter 15 | Support and Reference Materials | 281 |
| Overview | | 281 |
| The NIST Definition of Cloud Computing | | 281 |
| 15.1 | Characteristics of Cloud Computing | 284 |
| 15.2 | Commonly Cited Benefits of Cloud Computing | 284 |
| 15.3 | Most Cited Risks of Cloud Computing | 285 |
| 15.4 | Coping Strategies for Perceived Risks Associated with Cloud Computing | 286 |
| 15.5 | Threats to Security in the Cloud | 287 |
| 15.6 | Reasons for Capacity Planning | 288 |
| 15.7 | Step-by-Step Work Plan for Capacity Planning with Amazon EC2 | 288 |
| 15.8 | Cloud Capacity Planning and Classical Approach Compared | 293 |
| 15.9 | SLA Failures and Potential Solutions | 293 |
| 15.10 | Coping Strategies for Security Threats | 294 |
| 15.11 | General Questions to Ask When Migrating to the Cloud | 295 |
| Questions for Avoiding Lock-In | | 295 |
| 15.12 | Vendor Questions About Security for Cloud Providers (CP) | 296 |
| 15.12.1 | Data Security (At Rest) | 296 |
| 15.12.2 | Data Security (In Transit) | 296 |
| 15.12.3 | Authentication | 297 |
| 15.12.4 | Separation Between the Customers | 297 |
| 15.12.5 | Cloud Legal and Regulatory Issues | 297 |
| 15.12.6 | Incident Response | 297 |
| 15.12.7 | Data Export Restrictions | 298 |
| 15.12.8 | Questions to Ask Potential Cloud Providers About Costs | 298 |
| Index | | 301 |

About the Author

David E. Y. Sarna

David E. Y. Sarna is a technologist, serial entrepreneur, and author of the popular blogs EyeOnTheCloud.com and GoogleGazer.com. Mr. Sarna is a Certified Systems Professional, a Certified Computer Programmer and Certified Data Processing Auditor. He is the co-author, with George Febish, of PC Magazine Windows Rapid Application Development (published by Ziff-Davis Press) which went into three printings and was translated into several languages; he has also written five other books and more than 120 articles published in professional magazines. His longtime column “Paradigm Shift” was the most popular feature in Datamation for many years. Mr. Sarna holds several patents in the fields of bar code and kiosk technologies. He has been honored by the Computer Measurement Group, Inc., by IBM, and by Microsoft Corporation, where he was a founding Regional Director of the Microsoft Developers Network. He has lectured widely and has appeared on television many times, including multiple national appearances on the Fox Network, CNN, and MSNBC.

Mr. Sarna is the founder and managing director of Hendon, Stamford Hill & Co., Inc. (HSH), strategy consulting (www.hshco.com). He has more than 35 years of experience as a merchant banker, management consultant and as an executive of high-technology companies. Prior to founding HSH, Mr. Sarna served for many years on the Advisory Board of Hudson Venture Partners, a well-known New York venture capitalist.

He has served as a board member, director and executive officer of the Ramaz School, and on the Board of Yavneh Academy, both prestigious not-for-profit schools.

Mr. Sarna was founder, chairman, chief executive officer, and a director of ObjectSoft Corporation, a publicly traded company which he founded in 1990. In 1988, Mr. Sarna founded Image Business Systems Corporation (IBS), a software company specializing in document image processing; the

company was founded as a spin-off of International Systems Services Corp. (ISS), which Mr. Sarna co-founded in 1981. IBS developed ImageSystem, the first large-scale client-server software for document image processing; it was marketed by IBM. Warburg Pincus and IBM were major investors in IBS, which went public and was listed on the NASDAQ.

At ISS, he architected ISS Three, a computer capacity planning and expert systems tool which ISS successfully marketed and ultimately sold successfully to UCCEL Corp., now part of Computer Associates. ISS itself was successfully sold to a public company.

From 1976 to 1981, Mr. Sarna was employed at Price Waterhouse & Co. as a management consultant, beginning as a senior consultant and rising to the position of senior manager. At the start of his career, Mr. Sarna worked for Honeywell, Inc. and a hardware engineer from 1969 to 1970, and for IBM Corp. from 1970 to 1976 in the large systems division of IBM World Trade Corp. in engineering and sales capacities.

Mr. Sarna holds a B.A. degree cum laude with honors from Brandeis University and did his graduate work in Computer Science at the Technion-Israel Institute of Technology.

Ivan Gelb collaborated with Mr. Sarna on matters related to cloud economics and capacity planning. He is past president and a director of Computer Measurement Group. He is also President of Gelb Information Systems Corporation (GIS), a consulting firm that provides management and technical consulting services in the United States and internationally. His extensive information technology (IT) background includes determination of optimum hardware and software requirements for mainframe and client-server systems; effectiveness evaluation of computer systems and related organizations; data communications systems design and implementation; computer systems end-to-end availability management, performance management and capacity planning; development of software packages; and proprietary measurement data analysis techniques.

During his more than 30 years of experience, Mr. Gelb performed technical and management services for more than 100 organizations such as JP Morgan, Merrill Lynch, PepsiCo, the FBI, the State of California, the New Jersey State Office of Information Technology, and the New York City Board of Education. He is a speaker at various technical conferences, writes articles and serves as editor for a number of trade publications.

Preface

I was first exposed to what would become the Internet way back in 1969, while I was still an undergraduate at Brandeis University working on ARPANet, the forerunner to the present Internet, which operated at the then blazing speed of 2,400 bps, ultimately increased to 50 Kbps (see http://www.computerhistory.org/internet_history/). I have been privileged to enjoy a front-row seat watching the technology speed up, evolve and mature over the past 45 years.

Without hesitation, I make this bold statement: Cloud computing will have a greater effect on our lives than the PC revolution and the dot-com revolution combined.

This book details how to go beyond the theory and build “industrial-strength” robust and highly scalable cloud computing applications for the enterprise. We discuss

- Whose platforms are available today
- What tools facilitate development
- How to fit the different pieces together
- How much it costs

We look at actual case studies, and examine costs, technologies, and problems that were overcome.

In this book, I’m assuming that I’m singing with the choir and that the choir knows the tune, if not quite all the lyrics. It’s not my objective to convince you to develop cloud-based applications. You’re already convinced. Another book in this series, *Cloud Computing: Implementation, Management, and Security* (paperback) by J W Rittinghouse, Hypersecurity LLC, provides a somewhat more technical (e.g., nuts and bolts) understanding of what cloud computing really means.)

I aim to help you select the best tools, to follow the best practices, and to avoid pitfalls so you can build effective and appropriate cloud applications.

Cloud Computing and Web 3.0 are disruptive technologies, and the technology is changing, developing, and improving with breathtaking speed. My blog eyeonthecloud.com keeps up with daily news and developments and is a useful supplement to this book.

Author's Acknowledgements

I have many people to thank.

My agent, Carole Jelen McClendon of Waterside Agency, who believed in me, and introduced me to my editor, Theron Shreve. Theron has been a pleasure to know and a privilege to work with, and I learned much from him. He made many, many contributions to this book. Darice L. Moore carefully and meticulously copy-edited the manuscript. My long-time friend Ivan Gelb contributed greatly to the chapter on cloud economics and capacity planning and made many helpful contributions throughout. My friend Eli Schor and Yishai Pinchover, both of Cav Systems read the manuscript and made many helpful suggestions.

A book like this could not be written without the assistance and cooperation of many vendors, from A (Amazon) to Z (Zend). Thanks are due to all the vendors for fact-checking, and for their courtesy in allowing the reuse of copyrighted material from their respective Web sites, including illustration credits, as noted.

In the vendor community, special thanks are due to Darcy Hansen and Craig Sowell, Director, Cloud Marketing and Communications, IBM; Helen Phung and Matt C. Douglass, Practice Fusion; Gordon Evans and Ariel Kelman, Salesforce.com; Jerry Huang, Gladinet; Vicki Morris, Candice Heidebrecht, and Pete Malcolm, Abiquo; Jim Kerr and Brian Daly, Unisys; Phil Libin, Evernote; Betsy Zikakis, Rightscale; Lars Leckie, Hummer Winblad; Junaid Saiyed, Elastra; Bob Gordon, and Andy Kicklighter, Senior Technical Marketing Manager – Cloud at Nimsoft, CA Technologies; and Paul Martiz, Tod Nielsen, and Aaron T. Feigin, VMware.

My children, Ariel, Leora, and Shira, helped in many ways, and are a source of never-ending joy. Last and most dear is my wife of 35 years, the love of my life, Dr. Rachel C. Sarna, without whose constant support, encouragement, and competence very little would get accomplished in my house.

Executive Summary

Cloud Computing is a True Paradigm Shift

Cloud computing is a major paradigm shift. Most of us are already using cloud computing in our daily lives for personal use, and now enterprises are rapidly moving key applications to the clouds for agility (speed of implementation and speed of deployment), improved customer experience, scalability, and cost control.

From Do It Yourself to Public Cloud—A Continuum

Cloud computing is a natural development to meet needs that have been evident for more than forty years. *Virtualization* is the key technology that enables cloud computing. Remote hosting has developed from simply renting infrastructure to providing and maintaining standardized virtual servers that can be scaled up and down as demand fluctuates. Private (often on-premise) clouds can provide increased utilization compared with deploying, installing, and maintaining traditional farms of servers deployed on a task-per-server basis. Public clouds offer increased efficiency, but are perceived as being less secure. Newer hybrid solutions, such as IBM Cloudburst, Eucalyptus, and Windows AppFabric provide a consistent development experience for easy migration, redundancy, and scalability.

Cloud Computing: Is It Old Mainframe Bess in a New Dress?

Cloud computing is very much in vogue. Many vendors are simply relabeling their offerings to pretend they are cloud applications, but they are not the real McCoy.

The vision of a computer utility goes back to the 1960s, to John Marthey, Fernando Corbató, and Fred Guenberger. The rapid development of

the personal computer pushed aside interest in its development, which in any event required ubiquitous, high-speed Internet access to become a reality.

Many vendors offer managed platforms as a service. Universal standards are emerging, but there is not yet a universal standard as vendors fight for competitive advantage.

Commercial software developers and well as enterprise developers are building robust, multitenant software-as-a-service applications to run efficiently on these platforms, and usage is anticipated to explode over the next few years.

Moving Into and Around the Clouds and Efforts at Standardization

Most enterprise IT organizations have either implemented or are studying cloud projects. The two most commonly expressed fears are:

- How do we keep our data safe?
- How do we prevent being locked in to a single vendor?
- How do we move legacy applications to the cloud?

Portability of data and applications is crucial. Several versions of Linux have been optimized for the clouds. Linux, Apache, and the programming languages C++, Python, and Java, as well as PHP, have been widely adopted and are supported by many vendors. Leading tools like Eucalyptus and RightScale have also been adopted by many vendors, enhance portability, and prevent lock-in, as does the use of “wrappers” like Zend. VPN-Cubed IPsec supports hybrid clouds across multiple vendors.

The Simple Cloud API project empowers developers to use one interface to interact with a variety of cloud application services, enabling them to more easily access new technologies from cloud vendors.

Tools like Abiquo and 3Tera’s AppLogic facilitate cloud management across vendors, hypervisor independence, and the support of thousands of virtual machines. Elastra’s CloudServer facilitates provisioning and enforcement of policy rules.

Serious efforts are underway in the vendor community to promote portability within the cloud as well as emerging standards for high-speed interclouding and Open PaaS. The Distributed Management Task Force, Inc., (DMTF) is working to develop universal cloud service portability through the work of the Open Cloud Standards Incubator (OCSI).

Cloud Economics and Capacity Management

The goal of capacity planning is to ensure that you always have sufficient but not excessive resources to meet customers' needs in a timely fashion. In this chapter, we look at the economics of cloud computing, and the tools of capacity management needed to ensure adequate performance without overpaying.

Queueing theory explains how overutilization of resources causes poor response times and erratic performance. Key Volume Indicators (KVI) are a technique for relating computer metrics to units of forecastable work.

We discuss tools for evidence-based decision making, including measuring workloads, forecasting changes, modeling the expected workloads on different physical and virtual configurations to meet business needs at optimal cost, and validating the models for accuracy and robustness.

We discuss questions to ask cloud vendors about elasticity (scalability), and how to work through make versus buy decisions.

Demystifying the Cloud: A Case Study Using Amazon's Cloud Services (AWS)

Amazon began providing Amazon Web Services in 2005, and the early adopters did their best to treat it as a black art, known only to the cognoscenti.

Amazon's Web Services is the oldest and most mature of the public cloud service providers. An easy way to get started with AWS is to use Gladinet to create a Z disk that appears local but connects to the cloud. S3Fox Organizer is a free tool for moving static content from our own computers to the cloud. A custom instance of a virtual server on EC2 can be easily created and configured just the way we want it. Amazon's CloudWatch is useful for monitoring EC2 instances and Elastic Load Balancers in real time or by using Auto Scaling to dynamically add or remove Amazon EC2 instances based on Amazon CloudWatch metrics. Nimsoft's Nimsoft Monitoring Solution (NMS) for AWS is one tool for monitoring and controlling Amazon-hosted cloud solutions.

Virtualization: Open Source and VMware

Virtualization is the main ingredient of cloud computing. While it's an old idea, it is modern and fast, and low-cost, mass-produced hardware has made virtualization cost-effective. Many powerful hypervisors, including Xen, KVM, and QEMU, are open source. VMware is the commercial leader, but

is based on open source. Citrix is a form of virtual desktop, but today it often rides on VMware. Amazon uses a modified version of Xen. Monitoring is essential to managing the performance of virtual systems. Microsoft has its own patented approach in Microsoft Azure. EMC's VPLEX is an important new technology for moving blocks of storage across the cloud. Interesting partnerships have been announced among VMware, Google, Salesforce.com, Eucalyptus, and Amazon that will help grow the entire industry and prevent lock-in to a single vendor.

Securing the Cloud: Reliability, Availability, and Security

Reliability, Availability, and Security (RAS) are the three greatest concerns about migrating to the cloud. *Reliability* is often covered by a service level agreement (SLA). *Availability* addresses not only whether sufficient resources are available but also how long provisioning of new resources can take and how quickly they can be deprovisioned as needs scale back down. The main goal of *Security* is to limit access only to those approved, to let those with approved see and/or modify only the data they are entitled to see and no other data, and to ensure that no one can requisition resources beyond their budget.

However, many commercial service providers have better tools and facilities for ensuring RAS than do their clients. ISO 27001 and SAS 70 are two recognized standards designed for independently ensuring that third parties handling data have sufficient controls in place. These standards have been adapted for cloud security. The Cloud Security Alliance has been developing cloud-specific standards that will further improve on such standards. CloudAudit is developing an open, extensible, and secure interface that allows cloud computing providers to expose Audit, Assertion, Assessment, and Assurance (A4) information for cloud infrastructure (IaaS), platform (PaaS), and application (SaaS) services to authorized clients.

Scale and Reuse: Standing on the Shoulders of Giants

There are two principle attributes of cloud computing: scalability and code reuse. Service-Oriented Architecture (SOA) is a flexible set of design principles used during the phases of systems development and integration. SOA separates functions into distinct units, or services, comprised of unassociated, loosely coupled units of functionality that have no calls to each other embedded in them. Developers make them accessible over a network in

order to allow users to combine and reuse them in the production of applications. SOA, Web 2.0, and SOA 2.0 promote code reuse in a cloud environment. Calls to cloud-provided services, such as Google's AJAX APIs, also let you implement rich, dynamic Web sites entirely in JavaScript and HTML.

Windows Azure

Microsoft Azure represents a major evolution both of operating systems and of Microsoft's overall strategy. While written entirely from the ground up, it benefits from a long, mostly distinguished, and expensive pedigree. It seems to be the first-to-market component of Midori, the descendant of Cairo, Microsoft's two-decades-ago planned, never released object-oriented distributed operating system. Midori's strong emphasis on concurrency issues, a willingness to break compatibility, and the idea of using a hypervisor "as a kind of Meta-OS" fits Microsoft's long-term strategy.

Azure is a great place to develop and host .Net applications, an adequate place to build and host LAMP applications, and a very good place for hosting applications developed in a mixed environment.

Google in the Cloud

Google is believed to manage one of the two or three largest server farms in the world. Recently, it has begun making its infrastructure available to others for a fee. Its widely used offerings, in addition to search, include Google Apps for Business, Google Maps, Google Finance, and Google Voice. More recently, it has introduced Google App Engine, and in its own unique way, it is now a general cloud services provider.

Google is aiming to be an enterprise cloud vendor. Its approach to development in the cloud may be summarized as:

- Stand on our tall shoulders (Use Google's extensive code base)
- Develop your applications in Java, PHP, or Python
- Use the GQL datastore as an alternative to SQL
- Let Google worry about resource allocation, load balancing, and scalability

GWT, Google App Engine and Google Apps Script offer clear evidence of Google's big push to bring enterprise development to the cloud. Google

App Engine should prove to be a worthy competitor to Amazon Web Services, one of App Engine's major competitors for hosting environments.

Enterprise Cloud Vendors

Traditional enterprise vendors all have cloud strategies.

IBM was a very early proponent of both virtualization and cloud computing. IBM Smart Business cloud solutions support clouds built behind the enterprise firewall, or the IBM cloud. IBM's public cloud offering is still new, while its private cloud offerings are, for the cloud, very mature.

IBM has partnered with Red Hat, SOASTA, RightScale, and others. HP has partnered with Microsoft; Oracle bought Sun outright and partners with Amazon AWS.

Hewlett Packard is primarily selling cloud solutions to the enterprise, and the enterprise is typically hosting the solutions in private clouds. Its SaaS offerings are still specialized and limited. HP's BSM 9.0. addresses hybrid delivery models and management of the "consumerization of IT," i.e., people who use non-company-owned devices on a company network. HP, Intel, and cloud software maker Enomaly have partnered to offer a full end-to-end IaaS platform for cloud service providers. HP and Microsoft are investing \$250 million to significantly simplify cloud technology environments for businesses of all sizes.

Oracle has reclassified its clusters as private clouds while slowly moving into the "enterprise private cloud." Oracle customers can now use their existing Oracle licenses or acquire new licenses to deploy Oracle software on Amazon's EC2. Oracle has also announced its intention to license others as well. The newly introduced Oracle Secure Backup Cloud module makes it possible to move database backups to the Amazon Simple Storage Service (S3) for offsite storage. Its VeriScale architecture (part of Oracle's acquisition of Sun) optimizes load balancing by implementing the networking logic locally in the service instance's containers and treating the networking logic as part of the application.

CA Technologies (formerly Computer Associates) acquired 3Tera. Its AppLogic offers an innovative solution for building cloud services and deploying complex enterprise-class applications to public and private clouds using an intuitive graphical user interface (GUI). CA has extended its partnership with NetApp, integrating CA's virtualization, automation, and service assurance offerings with NetApp's storage management solutions. In addition, CA is planning SaaS offerings for most of its IT management solutions. It also partners with a Carnegie Mellon and a host of resellers.

Unisys places strong emphasis on security; this has carried over to its cloud offerings. Unisys Stealth security solution, an innovative, patent-pending data protection technology initially designed for government applications, is now available to commercial clients. Unisys has also partnered with VMware on the software side and with its parent, EMC, on the hardware side, among others.

Cloud Service Providers

Large cloud service providers include Rackspace, GoGrid, and Joyent (a newer entry), as well as robust cloud offerings from AT&T. EngineYard is a specialized provider of cloud services for Ruby on Rails developers. Other interesting SaaS vendors include NetSuite, Intuit, and Intacct, as well as cross-platform vendors like 3Tera, Appistry, Elasta, RightScale, BMS, and Nasuni.

Practice Fusion Case Study

Practice Fusion has demonstrated that cloud computing enables it to offer sophisticated applications to a wide audience at extremely low cost, while respecting HIPAA privacy and security mandates. Physicians are armed with good and complete data at the point of care; this is a significant paradigm shift from traditional paper-centric processes. A cloud-based environment prepares providers by focusing on the condition rather than by asking repeated questions around past medical history of the patient because they couldn't find it in a traditional paper chart or noninteroperable environment.

Support and Reference Materials

Charts and tables review the basic definitions of cloud computing, its characteristics, delivery models, and deployment models. Commonly cited benefits are listed, and the main concerns articulated. Pathways are identified for mitigating the risks. We also specifically articulate security concerns and pathways for mitigating security risks. Questionnaires are provided to ask internally and to vendors regarding:

- When to migrate to the cloud
- How to avoid lock-in
- What security is available
- What migrating to the cloud will cost

Chapter 1

Cloud Computing is a True Paradigm Shift



Figure 1.1 Cumulus clouds; photograph taken at Swifts Creek, in the Great Alps of East Gippsland, Victoria, Australia. Image by Fir0002/Flagstaffotos. Licensed under Gnu Free Documentation License (GFDL).

Chapter Overview

In this introductory chapter, we look at what cloud computing really is, why it's generating such excitement, and who are the major players. We make the point that most of us are already using cloud computing in our daily lives for personal use and show how enterprises are moving key applications to the clouds for improved customer experience, scalability, and cost control.

1.1 Introduction

A **cloud**, of course, is a visible mass of droplets or frozen crystals floating in the atmosphere above the surface of the Earth or another planetary body. A cloud is also a visible mass attracted by gravity. Lately, cloud computing has been exerting a strong gravitational pull all of its own—one that has been attracting a mass of money.

The big players in cloud computing are Google, Amazon, and, of late, Microsoft and IBM. Maybe Oracle/Sun, maybe HP will join them. Rack-space, GoGrid, and AT&T want in too.