



Cloud Computing

Implementation,
Management,
and Security

John W. Rittinghouse
James F. Ransome



CRC Press
Taylor & Francis Group

Cloud Computing

Cloud Computing

Implementation,
Management,
and Security

John W. Rittinghouse
James F. Ransome



CRC Press

Taylor & Francis Group

Boca Raton London New York

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

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

© 2009 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Version Date: 20131121

International Standard Book Number-13: 978-1-4398-0681-4 (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 CRC Press Web site at
<http://www.crcpress.com>

Contents

Foreword	xiii
Preface	xix
Introduction	xxv
What Is the Cloud?	xxvi
The Emergence of Cloud Computing	xxvi
The Global Nature of the Cloud	xxvii
Cloud-Based Service Offerings	xxviii
Grid Computing or Cloud Computing?	xxxi
Is the Cloud Model Reliable?	xxxi
Benefits of Using a Cloud Model	xxxii
What About Legal Issues When Using Cloud Models?	xxxii
What Are the Key Characteristics of Cloud Computing?	xxxiv
Challenges for the Cloud	xxxvi
Chapter 1 The Evolution of Cloud Computing	1
1.1 Chapter Overview	1
1.2 Hardware Evolution	2
1.2.1 First-Generation Computers	3
1.2.2 Second-Generation Computers	4
1.2.3 Third-Generation Computers	5
1.2.4 Fourth-Generation Computers	6
1.3 Internet Software Evolution	7
1.3.1 Establishing a Common Protocol for the Internet	12
1.3.2 Evolution of Ipv6	13

- 1.3.3 Finding a Common Method to Communicate Using the Internet Protocol 13
- 1.3.4 Building a Common Interface to the Internet 15
- 1.3.5 The Appearance of Cloud Formations—From One Computer to a Grid of Many 20
- 1.4 Server Virtualization 24
 - 1.4.1 Parallel Processing 25
 - 1.4.2 Vector Processing 26
 - 1.4.3 Symmetric Multiprocessing Systems 26
 - 1.4.4 Massively Parallel Processing Systems 27
- 1.5 Chapter Summary 28

Chapter 2 Web Services Delivered from the Cloud 29

- 2.1 Chapter Overview 29
- 2.2 Communication-as-a-Service (CaaS) 30
 - 2.2.1 Advantages of CaaS 31
 - 2.2.2 Fully Integrated, Enterprise-Class Unified Communications 32
- 2.3 Infrastructure-as-a-Service (IaaS) 34
 - 2.3.1 Modern On-Demand Computing 36
 - 2.3.2 Amazon’s Elastic Cloud 37
 - 2.3.3 Amazon EC2 Service Characteristics 38
 - 2.3.4 Mosso (Rackspace) 42
- 2.4 Monitoring-as-a-Service (MaaS) 44
 - 2.4.1 Protection Against Internal and External Threats 45
 - 2.4.2 Delivering Business Value 47
 - 2.4.3 Real-Time Log Monitoring Enables Compliance 48
- 2.5 Platform-as-a-Service (PaaS) 48
 - 2.5.1 The Traditional On-Premises Model 49
 - 2.5.2 The New Cloud Model 49
 - 2.5.3 Key Characteristics of PaaS 49
- 2.6 Software-as-a-Service (SaaS) 50
 - 2.6.1 SaaS Implementation Issues 51
 - 2.6.2 Key Characteristics of SaaS 52
 - 2.6.3 Benefits of the SaaS Model 53
- 2.7 Chapter Summary 54

Chapter 3	Building Cloud Networks	57
3.1	Chapter Overview	57
3.2	The Evolution from the MSP Model to Cloud Computing and Software-as-a-Service	59
3.2.1	From Single-Purpose Architectures to Multipurpose Architectures	60
3.2.2	Data Center Virtualization	61
3.3	The Cloud Data Center	62
3.4	Collaboration	62
3.4.1	Why Collaboration?	65
3.5	Service-Oriented Architectures as a Step Toward Cloud Computing	70
3.6	Basic Approach to a Data Center-Based SOA	72
3.6.1	Planning for Capacity	73
3.6.2	Planning for Availability	73
3.6.3	Planning for SOA Security	74
3.7	The Role of Open Source Software in Data Centers	75
3.8	Where Open Source Software Is Used	77
3.8.1	Web Presence	78
3.8.2	Database Tier	81
3.8.3	Application Tier	83
3.8.4	Systems and Network Management Tier	87
3.9	Chapter Summary	101
Chapter 4	Virtualization Practicum	103
4.1	Chapter Overview	103
4.2	Downloading Sun xVM VirtualBox	104
4.3	Installing Sun xVM VirtualBox	106
4.4	Adding a Guest Operating System to VirtualBox	112
4.5	Downloading FreeDOS as a Guest OS	112
4.6	Downloading the 7-Zip Archive Tool	114
4.7	Adding a Guest OS to Sun xVM VirtualBox	115
4.8	Chapter Summary	127
Chapter 5	Federation, Presence, Identity, and Privacy in the Cloud	129
5.1	Chapter Overview	129
5.2	Federation in the Cloud	129
5.2.1	Four Levels of Federation	132

5.2.2	How Encrypted Federation Differs from Trusted Federation	134
5.2.3	Federated Services and Applications	134
5.2.4	Protecting and Controlling Federated Communication	135
5.2.5	The Future of Federation	136
5.3	Presence in the Cloud	136
5.3.1	Presence Protocols	137
5.3.2	Leveraging Presence	138
5.3.3	Presence Enabled	139
5.3.4	The Future of Presence	139
5.3.5	The Interrelation of Identity, Presence, and Location in the Cloud	140
5.3.6	Federated Identity Management	140
5.3.7	Cloud and SaaS Identity Management	141
5.3.8	Federating Identity	143
5.3.9	Claims-Based Solutions	144
5.3.10	Identity-as-a-Service (IaaS)	144
5.3.11	Compliance-as-a-Service (CaaS)	145
5.3.12	The Future of Identity in the Cloud	146
5.4	Privacy and Its Relation to Cloud-Based Information Systems	147
5.4.1	Privacy Risks and the Cloud	149
5.4.2	Protecting Privacy Information	150
5.4.3	The Future of Privacy in the Cloud	151
5.5	Chapter Summary	152

Chapter 6 Security in the Cloud 153

6.1	Chapter Overview	153
6.2	Cloud Security Challenges	158
6.3	Software-as-a-Service Security	162
6.3.1	Security Management (People)	164
6.3.2	Security Governance	165
6.3.3	Risk Management	165
6.3.4	Risk Assessment	165
6.3.5	Security Portfolio Management	166
6.3.6	Security Awareness	166
6.3.7	Education and Training	167
6.3.8	Policies, Standards, and Guidelines	167
6.3.9	Secure Software Development Life Cycle (SecSDLC)	168

6.3.10	Security Monitoring and Incident Response	169
6.3.11	Third-Party Risk Management	169
6.3.12	Requests for Information and Sales Support	169
6.3.13	Business Continuity Plan	170
6.3.14	Forensics	170
6.3.15	Security Architecture Design	172
6.3.16	Vulnerability Assessment	173
6.3.17	Password Assurance Testing	173
6.3.18	Logging for Compliance and Security Investigations	173
6.3.19	Security Images	173
6.3.20	Data Privacy	174
6.3.21	Data Governance	175
6.3.22	Data Security	175
6.3.23	Application Security	176
6.3.24	Virtual Machine Security	177
6.3.25	Identity Access Management (IAM)	177
6.3.26	Change Management	178
6.3.27	Physical Security	178
6.3.28	Business Continuity and Disaster Recovery	179
6.3.29	The Business Continuity Plan	180
6.4	Is Security-as-a-Service the New MSSP?	181
6.5	Chapter Summary	182

Chapter 7 Common Standards in Cloud Computing 183

7.1	Chapter Overview	183
7.2	The Open Cloud Consortium	183
7.3	The Distributed Management Task Force	185
7.3.1	Open Virtualization Format	186
7.4	Standards for Application Developers	187
7.4.1	Browsers (Ajax)	188
7.4.2	Data (XML, JSON)	189
7.4.3	Solution Stacks (LAMP and LAPP)	192
7.5	Standards for Messaging	193
7.5.1	Simple Message Transfer Protocol (SMTP)	193
7.5.2	Post Office Protocol (POP)	194

7.5.3	Internet Messaging Access Protocol (IMAP)	194
7.5.4	Syndication (Atom, Atom Publishing Protocol, and RSS)	194
7.5.5	Communications (HTTP, SIMPLE, and XMPP)	202
7.6	Standards for Security	205
7.6.1	Security (SAML OAuth, OpenID, SSL/TLS)	205
7.7	Chapter Summary	212

Chapter 8 End-User Access to Cloud Computing 213

8.1	Chapter Overview	213
8.2	YouTube	214
8.3	YouTube API Overview	215
8.3.1	Widgets	216
8.3.2	YouTube Player APIs	217
8.3.3	The YouTube Custom Player	218
8.3.4	YouTube Data API	218
8.4	Zimbra	219
8.4.1	Zimbra Collaboration Suite (ZCS)	221
8.5	Facebook	221
8.5.1	Facebook Development	223
8.6	Zoho	223
8.6.1	Zoho CloudSQL	225
8.7	DimDim Collaboration	226
8.8	Chapter Summary	233

Chapter 9 Mobile Internet Devices and the Cloud 235

9.1	Chapter Overview	235
9.2	What Is a Smartphone?	236
9.3	Mobile Operating Systems for Smartphones	237
9.3.1	iPhone	237
9.3.2	Google (Android)	237
9.3.3	Blackberry	241
9.3.4	Windows Mobile	241
9.3.5	Ubuntu Mobile Internet Device (MID)	243
9.4	Mobile Platform Virtualization	251
9.4.1	KVM	253
9.4.2	VMWare	254
9.5	Collaboration Applications for Mobile Platforms	256

9.6	Future Trends	257
9.7	Chapter Summary	258
9.8	Closing Comments	258
Appendix A Virtualization Practicum (Linux)		261
A.1	Overview	261
A.2	Adding a Linux-Based Guest Operating System to VirtualBox	262
A.3	Downloading OpenSolaris as a Guest Operating System	263
A.4	Using the 7-Zip Archive Tool	264
A.5	Adding the OpenSolaris Guest OS to Sun xVM VirtualBox	265
A.6	Summary	281
Appendix B Executive Scenario for Cloud Migration		285
Index		297

Foreword

While there is no arguing about the staying power of the cloud model and the benefits it can bring to any organization or government, mainstream adoption depends on several key variables falling into alignment that will provide users the reliability, desired outcomes, and levels of trust necessary to truly usher in a “cloud revolution.” Until recently, early adopters of cloud computing in the public and private sectors were the catalyst for helping drive technological innovation and increased adoption of cloud-based strategies, moving us closer to this inevitable reality. Today, driven in large part by the financial crisis gripping the global economy, more and more organizations are turning toward cloud computing as a low-cost means of delivering quick-time-to-market solutions for mission-critical operations and services. The benefits of cloud computing are hard to dispute:

1. Reduced implementation and maintenance costs
2. Increased mobility for a global workforce
3. Flexible and scalable infrastructures
4. Quick time to market
5. IT department transformation (focus on innovation vs. maintenance and implementation)
6. “Greening” of the data center
7. Increased availability of high-performance applications to small/medium-sized businesses

Gartner, in a February 2, 2009, press release, posed the question of why, when “the cloud computing market is in a period of excitement, growth and high potential. . . [we] will still require several years and many

changes in the market before cloud computing is a mainstream IT effort”¹ In talking with government and industry leaders about this, it became clear that the individual concerns and variables that were negatively impacting business leaders’ thought processes regarding cloud computing (and therefore preventing what could be even more growth in this market) could be boiled down to one addressable need: a lack of understanding. Let’s take this case in point: GTRA research showed that the most common concern about implementing cloud programs was security and privacy, a finding supported by an IDC study of 244 CIOs on cloud computing, in which 75% of respondents listed security as their number-one concern.² It is true that moving from architectures that were built for on-premises services and secured by firewalls and threat-detection systems to mobile environments with SaaS applications makes previous architectures unsuitable to secure data effectively. In addition, at a March 2009 FTC meeting discussing cloud computing security and related privacy issues, it was agreed that data management services might experience failure similar to the current financial meltdown if further regulation was not implemented. In short, some executives are simply too scared to move forward with cloud initiatives.

However, this concern, while valid, is not insurmountable. Already there are countless examples of successful cloud computing implementations, from small organizations up to large enterprises that have low risk tolerance, such as the U.S. Department of the Navy. The security community is also coming together through various initiatives aimed at education and guidance creation. The National Institute of Standards and Technologies (NIST) is releasing its first guidelines for agencies that want to use cloud computing in the second half of 2009, and groups such as the Jericho forum are bringing security executives together to collaborate and deliver solutions. As with any emerging technology, there exists a learning curve with regard to security in a cloud environment, but there is no doubt that resources and case studies exist today to help any organization overcome this.

The same types of pros and cons listed above can be applied to other concerns facing executives, such as data ownership rights, performance, and availability. While these are all valid concerns, solutions do exist and are being fine-tuned every day; the challenge is in bringing executives out of a state of unknown and fear and giving them the understanding and

1. “Cloud Application Infrastructure Technologies Need Seven Years to Mature,” Gartner, Inc., December 2008.
2. “IT Cloud Services User Study,” IDC, Inc., October 2008.

knowledge necessary to make informed, educated decisions regarding their cloud initiatives.

In this book, Drs. Rittinghouse and Ransome do a tremendous job of educating, dispelling myths, and giving detailed examples and steps which will provide the reader with a proper understand of cloud computing, its risks, and how to implement and manage an effective cloud strategy. This is all done in a manner that is reader-friendly but with enough detailed technical language to be complete, and not so much that a nontechnical leader will be lost.

In the Introduction and Chapter 1, Drs. Rittinghouse and Ransome lay the foundation for the reader's proper understanding of cloud computing, detailing its history and evolution and discussing how new technologies such as virtualization played a huge role in the growth and acceptance of cloud computing. Chapter 2 then educates us on the different types of services which can be delivered from the cloud, providing detail on Software-as-a-Service (SaaS), Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), Monitoring-as-a-Service (MaaS), and Communication-as-a-Service (CaaS).

Chapter 3 dives into the heart of what it means to build a cloud network, including a look at the roles that service-oriented architecture (SOA) and open source software play in the process. Following this, Chapter 4 is dedicated entirely to the topic of virtualization, a critical component of any cloud network and one of the technologies which is a foundation of cloud concepts.

Security and privacy, one of the largest areas of concern for anyone building a cloud network, are covered in Chapters 5 and 6. These chapters look at how federation in the cloud and federated services and applications can be used to increase security, build trust, and mitigate risk. Dr. Ron Ross, a senior computer scientist at NIST, recently said, "You're never going to have complete trust. We don't live in a risk-free environment—we have to manage risk, not avoid it." These chapters give the reader a wealth of guidance, practical applications, and process, which can be used to keep risk at an acceptable level in any cloud network.

Chapter 7 shifts focus to look at common standards in cloud computing, including standards for application development, messaging, and security. Social networking and collaboration is the focus of Chapter 8, in which the authors discuss end-user access to cloud computing (You Tube, Facebook, etc.). Chapter 9, the book's final chapter, discusses in detail how

mobile Internet devices react with cloud networks—a topic which is critical now and will only increase in importance as users expect more and more applications to be delivered to their smartphones and other mobile devices.

We feel that completing this book, readers will have a thorough, well-rounded understanding of cloud computing, the knowledge necessary to overcome fears, and will be armed with the guidance necessary to make smart, strategic decisions regarding their cloud initiatives. Ultimately, this book will play a part in ushering in the “cloud revolution” and will help overcome the lack of understanding currently preventing even faster adoption of cloud computing.

Kelly Yocum

Parham Eftekhari

Co-Founders, Government Technology Research Alliance

Kelly Yocum and Parham Eftekhari are the co-founders of the Government Technology Research Alliance (GTRA), an organization that provides government CXO leaders a forum in which to collaborate, strategize, and create innovative solutions for today’s most pressing IT needs. Kelly is GTRA’s executive director and is responsible for strategic direction, business development, and work with solution and technology providers for the GTRA Government Council. She also serves as the CEO for GOVTek, a collaborative online information resource for government technology executives and industry experts. Kelly was formerly CEO of ConVurge, a business intelligence conference company, where she founded several councils for government technology including SecureGOV, ArchitectureGOV, MobileGOV, and HrGOV, which are currently managed by GTRA. She invented a unique government-to-industry collaboration model, called GTRA Roundtable Meetings, which foster an innovative discussion forum for government and industry experts.

Parham Eftekhari serves as director of research and curriculum development for GTRA, where he is responsible for overseeing all research conducted with senior government technology executives and industry leaders on technology and leadership issues. Parham’s areas of expertise include transparency/open government, enterprise architecture, security, virtualization, information sharing, social networking/Web 2.0, knowledge management, green IT, records management, mobility, and cloud computing.

Parham is also responsible for growing GTRA's councils with key government leaders and assisting in the government-to-industry collaboration model. Parham is also vice president of GOVTek, where his primary focus is to oversee the content, research, and resources shared on the site. Parham formerly served as director of technology research for Proactive Worldwide, managing the full life cycle of competitive intelligence, strategic, and market assessment research studies. Together, Parham and Kelly run the semiannual GTRA Council Meeting Symposia, which bring together executive-level decision makers from both the public and private sectors to collaborate, share ideas, and discuss solutions to current challenges. This forum is a unique model for government and technology collaboration in which the concepts of cloud computing and the cloud's value to the next generation of consumers and practitioners in both government and commercial sectors are presented.

Preface

There are lots of books on cloud computing in the market today. This one is not intended for “supergeeks” looking for the next revelation in “geek know-how.” In fact, it attempts to present cloud computing in a way that anyone can understand. We do include technical material, but we do so in a way that allows managers and technical people alike to understand what exactly cloud computing is and what it is not. We try to clear up the confusion about current buzzwords such as PaaS, SaaS, etc., and let the reader see how and why the technology has evolved to become “the cloud” as we know and use it today.

In the Introduction we explain what cloud computing is, its characteristics, and the challenges it will face in the future. The biggest challenges that companies will face as they move into the cloud are secure data storage, high-speed access to the Internet, and standardization. Storing large amounts of data in centralized locations while preserving user privacy, security, identity, and their application-specific preferences raises many concerns about data protection. These concerns, in turn, lead to questions about the legal framework that should be implemented for a cloud-oriented environment.

In Chapter 1 we discuss the evolution of cloud computing, including hardware, software, and server virtualization. In order to discuss some of the issues involved in the cloud concept, it is important to place the development of computational technology in a historical context. Looking at the cloud’s evolutionary development, and the problems encountered along the way, provides some key reference points to help us understand the challenges that had to be overcome by those who were responsible for the development of the Internet and the World Wide Web. These challenges fell into three primary categories: hardware, software, and virtualization. We discuss how the rules computers use to communicate came about, and how the

development of networking and communications protocols helped drive the technology growth we have seen in the last two decades or so. This, in turn, has driven even more changes in protocols and forced the creation of new technologies to mitigate concerns and improve the methods used to communicate over the Internet. The rise of web browsers led to huge growth in use of the Internet and a migration away from the traditional data center toward cloud computing.

In Chapter 2 we discuss the advent of web-based services delivered from the cloud, including Communication-as-a-Service (CaaS), Infrastructure-as-a-Service (IaaS), Monitoring-as-a-Service (MaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). As technology has migrated from the traditional on-premises model to the new cloud model, service offerings have evolved almost daily. We provide some basic exposure to where the technology is today, and we give you a feel for where it will likely be in the not too distant future.

In Chapter 3 we discuss what is required from service providers to make the services described in Chapter 2 available. We describe the basic approach to service-oriented architecture (SOA) as it applies to data center design, how companies can build highly automated private cloud networks that can be managed from a single point, and how server and storage virtualization is used across distributed computing resources. We discuss what it takes to build a cloud network, the evolution from the managed service provider model to cloud computing and SaaS and from single-purpose architectures to multipurpose architectures, the concept and design of data center virtualization, the role and importance of collaboration, SOA as an intermediate step and the basic approach to data center-based SOA, and lastly, the role of open source software in data centers and where and how it is used in the cloud architecture.

In Chapter 4 we provide a virtualization practicum that guides you through a step-by-step process for building a virtualized computing infrastructure using open source software. The beauty of virtualization solutions is that you can run multiple operating systems simultaneously on a single computer. So that you could really understand how powerful that capability is, we show you how to do it for yourself. We show you how to download and install the Sun VirtualBox, how to install and configure it, and how to add a virtual operating environment on top of your existing operating system. In learning the basics of using the Sun xVM VirtualBox, you will also gain knowledge about what virtualization is and how it can be used.

Chapter 5 discusses the importance and relevance of federation, presence, identity, and privacy in cloud computing and the latest challenges, solutions, and potential future for each in the cloud. Building a seamless federated communications capability in a cloud environment, one that is capable of supporting people, devices, information feeds, documents, application interfaces, and other entities, depends on the architecture that is implemented. The solution chosen must be able to find such entities, determine their purpose, and request presence data so that others can interact with them in real time. This process is known as discovery.

The extension of virtualization and virtual machines into the cloud is affecting enterprise security because the traditional enterprise network perimeter is evaporating. In Chapter 6 we identify security as the greatest challenge in cloud computing, particularly with regard to the SaaS environment. Although there is a significant benefit to leveraging cloud computing, security concerns have led some organizations to hesitate to move critical resources to the cloud.

Corporations and individuals are concerned about how security and compliance integrity can be maintained in this new environment. Even more concerning, though, is the corporations that are jumping to cloud computing while being oblivious to the implications of putting critical applications and data in the cloud. Chapter 6 addresses the security concerns of the former and educates the latter. Moving critical applications and sensitive data to a public and shared cloud environment is a major concern for corporations that are moving beyond their data center's network perimeter defense. To alleviate these concerns, a cloud solution provider must ensure that customers can continue to have the same security and privacy controls over their applications and services, provide evidence to these customers that their organization and customers are secure and they can meet their service-level agreements, and show how can they prove compliance to their auditors.

Regardless of how the cloud evolves, it needs some form of standardization so that the market can evolve and thrive. Standards also allow clouds to interoperate and communicate with each other. In Chapter 7 we introduce some of the more common standards in cloud computing. Although we do not analyze each standard in depth, you should gain a feel for how and why each standard is used and, more important, a better understanding of why they evolved. Most current standards evolved from necessity, as individuals took a chance on new innovation. As these innovative techniques became

acceptable to users and implementers, more support for the technique ensued. At some point, the innovation began to be considered a “standard,” and groups formalized protocols or rules for using it. We discuss the Open Cloud Consortium and the Distributed Management Task Force as examples of cloud-related working groups.

Innovation leading to success in cloud services depends ultimately on acceptance of the application by the user community. In Chapter 8 we present some of the applications that are gaining acceptance among end users. We look at some of the most popular SaaS offerings for consumers and provide an overview of their benefits and why, in our opinion, they are helping to evolve our common understanding of what collaboration and mobility will ultimately mean in our daily lives. We examine five particularly successful SaaS offerings, YouTube, Zimbra, Facebook, Zoho, and DimDim, looking at them from both the user perspective and the developer/implementer perspective. This dual perspective should give you a clear understanding of how such offerings are transforming our concept of computing by making much traditional desktop-type software available from the cloud.

In Chapter 9 we detail the transition from fixed devices connected to the Internet to the new mobile device–empowered Internet. While it is essentially the same Internet, it has become tremendously more accessible, and advances in telephony, coupled with the use of the Internet, have led to some very compelling, powerful offerings. In this chapter we provide an overview of the more common offerings and how their widespread use will affect the cloud computing world. When more than 90% of your user base depends on mobile devices for common applications such as email, contacts, and media streaming or sharing, you cannot take the same approach as you used with statically connected Internet devices such as laptops and desktop PCs. It is a brave, new cloud-based world we are entering.

We hope that what you take away from reading this book is knowledge that separates hype from reality in talking about cloud computing. It seems that everyone you ask has a different answer. Most of the time, each answer you hear is based on one person’s experience with the cloud or with his or her desire to capitalize on the cloud for profit. Our intent is to present the cloud as an evolving, changing entity that does so out of demand from the Internet community itself. The technologies that are used in the cloud often give rise to new uses. For example, 10 years ago, you needed custom applications to watch video, the right codec had to be used for the right software,

etc. It was more trouble than watching the video was worth. Today, there is a *de facto* standard. Look at how YouTube has come about as a result of such innovation. After you read this book, you will know about the cloud, but not from the perspective of any one source; you will know from the perspective of how technological innovation has actually made it what it is.

Introduction

The purpose of this book is to clear up some of the mystery surrounding the topic of cloud computing. In order to understand how computing has evolved, one must understand the evolution of computing from a historical perspective, focusing primarily on those advances that led to the development of cloud computing, such as the transition from mainframes to desktops, laptops, mobile devices, and on to the cloud. We will also need to discuss in some detail the key components that are critical to make the cloud computing paradigm feasible with the technology available today. We will cover some of the standards that are used or are proposed for use in the cloud computing model, since standardization is crucial to achieving widespread acceptance of cloud computing. We will also discuss the means used to manage effectively the infrastructure for cloud computing. Significant legal considerations in properly protecting user data and mitigating corporate liability will also be covered. Finally, we will discuss what some of the more successful cloud vendors have done and how their achievements have helped the cloud model evolve.

Over the last five decades, businesses that use computing resources have learned to contend with a vast array of buzzwords. Much of this *geek-speak* or marketing vapor, over time, has been guilty of making promises that often are never kept. Some promises, to be sure, have been delivered, although others have drifted into oblivion. When it comes to offering technology in a *pay-as-you-use* services model, most information technology (IT) professionals have heard it all—from allocated resource management to grid computing, to on-demand computing and software-as-a-service (SaaS), to utility computing. A new buzzword, *cloud computing*, is presently in vogue in the marketplace, and it is generating all sorts of confusion about what it actually represents.

What Is the Cloud?

The term *cloud* has been used historically as a metaphor for the Internet. This usage was originally derived from its common depiction in network diagrams as an outline of a cloud, used to represent the transport of data across carrier backbones (which owned the cloud) to an endpoint location on the other side of the cloud. This concept dates back as early as 1961, when Professor John McCarthy suggested that computer time-sharing technology might lead to a future where computing power and even specific applications might be sold through a utility-type business model.¹ This idea became very popular in the late 1960s, but by the mid-1970s the idea faded away when it became clear that the IT-related technologies of the day were unable to sustain such a futuristic computing model. However, since the turn of the millennium, the concept has been revitalized. It was during this time of revitalization that the term *cloud computing* began to emerge in technology circles.

The Emergence of Cloud Computing

Utility computing can be defined as the provision of computational and storage resources as a metered service, similar to those provided by a traditional public utility company. This, of course, is not a new idea. This form of computing is growing in popularity, however, as companies have begun to extend the model to a cloud computing paradigm providing virtual servers that IT departments and users can access on demand. Early enterprise adopters used utility computing mainly for non-mission-critical needs, but that is quickly changing as trust and reliability issues are resolved.

Some people think cloud computing is the next big thing in the world of IT. Others believe it is just another variation of the utility computing model that has been repackaged in this decade as something new and cool. However, it is not just the buzzword “cloud computing” that is causing confusion among the masses. Currently, with so few cloud computing vendors actually practicing this form of technology and also almost every analyst from every research organization in the country defining the term differently, the meaning of the term has become very nebulous. Even among those who think they understand it, definitions vary, and most of those definitions are hazy at best. To clear the haze and make some sense of the new

1. <http://computinginthecloud.wordpress.com/2008/09/25/utility-cloud-computingflashback-to-1961-prof-john-mccarthy>, retrieved 5 Jan 2009.

concept, this book will attempt to help you understand just what cloud computing really means, how disruptive to your business it may become in the future, and what its advantages and disadvantages are.

As we said previously, the term *the cloud* is often used as a metaphor for the Internet and has become a familiar cliché. However, when “the cloud” is combined with “computing,” it causes a lot of confusion. Market research analysts and technology vendors alike tend to define cloud computing very narrowly, as a new type of utility computing that basically uses virtual servers that have been made available to third parties via the Internet. Others tend to define the term using a very broad, all-encompassing application of the virtual computing platform. They contend that anything beyond the firewall perimeter is in the cloud. A more tempered view of cloud computing considers it the delivery of computational resources from a location other than the one from which you are computing.

The Global Nature of the Cloud

The cloud sees no borders and thus has made the world a much smaller place. The Internet is global in scope but respects only established communication paths. People from everywhere now have access to other people from anywhere else. Globalization of computing assets may be the biggest contribution the cloud has made to date. For this reason, the cloud is the subject of many complex geopolitical issues. Cloud vendors must satisfy myriad regulatory concerns in order to deliver cloud services to a global market. When the Internet was in its infancy, many people believed cyberspace was a distinct environment that needed laws specific to itself. University computing centers and the ARPANET were, for a time, the encapsulated environments where the Internet existed. It took a while to get business to warm up to the idea.

Cloud computing is still in its infancy. There is a hodge-podge of providers, both large and small, delivering a wide variety of cloud-based services. For example, there are full-blown applications, support services, mail-filtering services, storage services, etc. IT practitioners have learned to contend with some of the many cloud-based services out of necessity as business needs dictated. However, cloud computing aggregators and integrators are already emerging, offering packages of products and services as a single entry point into the cloud.

The concept of cloud computing becomes much more understandable when one begins to think about what modern IT environments always

require—the means to increase capacity or add capabilities to their infrastructure dynamically, without investing money in the purchase of new infrastructure, all the while without needing to conduct training for new personnel and without the need for licensing new software. Given a solution to the aforementioned needs, cloud computing models that encompass a subscription-based or pay-per-use paradigm provide a service that can be used over the Internet and extends an IT shop's existing capabilities. Many users have found that this approach provides a return on investment that IT managers are more than willing to accept.

Cloud-Based Service Offerings

Cloud computing may be viewed as a resource available as a service for virtual data centers, but cloud computing and virtual data centers are not the same. For example, consider Amazon's S3 Storage Service. This is a data storage service designed for use across the Internet (i.e., the cloud). It is designed to make web-scale computing easier for developers. According to Amazon:

Amazon S3 provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web. It gives any developer access to the same highly scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites. The service aims to maximize benefits of scale and to pass those benefits on to developers.²

Amazon.com has played a vital role in the development of cloud computing. In modernizing its data centers after the dot-com bubble burst in 2001, it discovered that the new cloud architecture it had implemented resulted in some very significant internal efficiency improvements. By providing access to its systems for third-party users on a utility computing basis, via Amazon Web Services, introduced in 2002, a revolution of sorts began. Amazon Web Services began implementing its model by renting computing cycles as a service outside a given user's domain, wherever on the planet that domain might be located. This approach modernized a style of computing whereby IT-related capabilities could be provided “as a service”

2. <http://aws.amazon.com/s3>, retrieved 5 Jan 2009.

to users. By allowing their users to access technology-enabled services “in the cloud,” without any need for knowledge of, expertise with, or control over how the technology infrastructure that supports those services worked, Amazon shifted the approach to computing radically. This approach transformed cloud computing into a paradigm whereby data is permanently stored in remote servers accessible via the Internet and cached temporarily on client devices that may include desktops, tablet computers, notebooks, hand-held devices, mobile phones, etc., and is often called *Software as a Service* (SaaS).

SaaS is a type of cloud computing that delivers applications through a browser to thousands of customers using a multiuser architecture. The focus for SaaS is on the end user as opposed to managed services (described below). For the customer, there are no up-front investment costs in servers or software licensing. For the service provider, with just one product to maintain, costs are relatively low compared to the costs incurred with a conventional hosting model. Salesforce.com³ is by far the best-known example of SaaS computing among enterprise applications. Salesforce.com was founded in 1999 by former Oracle executive Marc Benioff, who pioneered the concept of delivering enterprise applications via a simple web site. Nowadays, SaaS is also commonly used for enterprise resource planning and human resource applications. Another example is Google Apps, which provides online access via a web browser to the most common office and business applications used today, all the while keeping the software and user data stored on Google servers. A decade ago, no one could have predicted the sudden rise of SaaS applications such as these.

Managed service providers (MSPs) offer one of the oldest forms of cloud computing. Basically, a managed service is an application that is accessible to an organization’s IT infrastructure rather than to end users. Services include virus scanning for email, antispam services such as Postini,⁴ desktop management services such as those offered by CenterBeam⁵ or Everdream,⁶ and

3. <http://www.salesforce.com>, retrieved 5 Jan 2009.

4. In September 2007, Google acquired Postini, recognized as a global leader in on-demand communications security and compliance solutions. This is further evidence of the aggregation of cloud service providers.

5. CenterBeam delivers services over the Internet using a SaaS model.

6. In November 2007, Dell signed an agreement to acquire Everdream, a leading provider of SaaS solutions for remote service management. The planned acquisition was a key component in Dell’s strategy of enabling customers to simplify IT. Everdream’s capabilities complement those provided by the recently acquired SilverBack Technologies, further enabling end-to-end remote management of customers’ IT environments.

application performance monitoring. Managed security services that are delivered by third-party providers also fall into this category.

Platform-as-a-Service (PaaS) is yet another variation of SaaS. Sometimes referred to simply as web services in the cloud, PaaS is closely related to SaaS but delivers a platform from which to work rather than an application to work with. These service providers offer application programming interfaces (APIs) that enable developers to exploit functionality over the Internet, rather than delivering full-blown applications. This variation of cloud computing delivers development environments to programmers, analysts, and software engineers as a service. A general model is implemented under which developers build applications designed to run on the provider's infrastructure and which are delivered to users in via an Internet browser. The main drawback to this approach is that these services are limited by the vendor's design and capabilities. This means a compromise between freedom to develop code that does something other than what the provider can provide and application predictability, performance, and integration.

An example of this model is the Google App Engine. According to Google, "Google App Engine makes it easy to build an application that runs reliably, even under heavy load and with large amounts of data."⁷ The Google App Engine environment includes the following features

- Dynamic web serving, with full support for common web technologies
- Persistent storage with queries, sorting, and transactions
- Automatic scaling and load balancing
- APIs for authenticating users and sending email using Google Accounts
- A fully featured local development environment that simulates Google App Engine on your computer

Currently, Google App Engine applications are implemented using the Python programming language. The runtime environment includes the full Python language and most of the Python standard library. For extremely lightweight development, cloud-based mashup platforms (Ajax modules that are assembled in code) abound, such as Yahoo Pipes or Dapper.net.

7. <http://code.google.com/appengine/docs/whatisgoogleappengine.html>, retrieved 5 Jan 2009.

Grid Computing or Cloud Computing?

Grid computing is often confused with cloud computing. Grid computing is a form of distributed computing that implements a *virtual supercomputer* made up of a cluster of networked or Internetworked computers acting in unison to perform very large tasks. Many cloud computing deployments today are powered by grid computing implementations and are billed like utilities, but cloud computing can and should be seen as an evolved next step away from the grid utility model. There is an ever-growing list of providers that have successfully used cloud architectures with little or no centralized infrastructure or billing systems, such as the peer-to-peer network BitTorrent and the volunteer computing initiative SETI@home.⁸

Service commerce platforms are yet another variation of SaaS and MSPs. This type of cloud computing service provides a centralized service hub that users interact with. Currently, the most often used application of this platform is found in financial trading environments or systems that allow users to order things such as travel or personal services from a common platform (e.g., Expedia.com or Hotels.com), which then coordinates pricing and service delivery within the specifications set by the user.

Is the Cloud Model Reliable?

The majority of today's cloud computing infrastructure consists of time-tested and highly reliable services built on servers with varying levels of virtualized technologies, which are delivered via large data centers operating under service-level agreements that require 99.99% or better uptime. Commercial offerings have evolved to meet the quality-of-service requirements of customers and typically offer such service-level agreements to their customers. From users' perspective, the cloud appears as a single point of access for all their computing needs. These cloud-based services are accessible anywhere in the world, as long as an Internet connection is available. Open standards and open-source software have also been significant factors in the growth of cloud computing, topics we will discuss in more depth later.

8. SETI@home is a scientific experiment that uses Internet-connected computers in the Search for Extraterrestrial Intelligence (SETI). For more information, see <http://www.seti.org>.

Benefits of Using a Cloud Model

Because customers generally do not own the infrastructure used in cloud computing environments, they can forgo capital expenditure and consume resources as a service by just paying for what they use. Many cloud computing offerings have adopted the utility computing and billing model described above, while others bill on a subscription basis. By sharing computing power among multiple users, utilization rates are generally greatly improved, because cloud computing servers are not sitting dormant for lack of use. This factor alone can reduce infrastructure costs significantly and accelerate the speed of applications development.

A beneficial side effect of using this model is that computer capacity increases dramatically, since customers do not have to engineer their applications for peak times, when processing loads are greatest. Adoption of the cloud computing model has also been enabled because of the greater availability of increased high-speed bandwidth. With greater enablement, though, there are other issues one must consider, especially legal ones.

What About Legal Issues When Using Cloud Models?

Recently there have been some efforts to create and unify the legal environment specific to the cloud. For example, the United States–European Union Safe Harbor Act provides a seven-point framework of requirements for U.S. companies that may use data from other parts of the world, namely, the European Union. This framework sets forth how companies can participate and certify their compliance and is defined in detail on the U.S. Department of Commerce and Federal Trade Commission web sites. In summary, the agreement allows most U.S. corporations to certify that they have joined a self-regulatory organization that adheres to the following seven Safe Harbor Principles or has implemented its own privacy policies that conform with these principles:

1. Notify individuals about the purposes for which information is collected and used.
2. Give individuals the choice of whether their information can be disclosed to a third party.
3. Ensure that if it transfers personal information to a third party, that third party also provides the same level of privacy protection.

4. Allow individuals access to their personal information.
5. Take reasonable security precautions to protect collected data from loss, misuse, or disclosure.
6. Take reasonable steps to ensure the integrity of the data collected.;
7. Have in place an adequate enforcement mechanism.

Major service providers such as Amazon Web Services cater to a global marketplace, typically the United States, Japan, and the European Union, by deploying local infrastructure at those locales and allowing customers to select availability zones. However, there are still concerns about security and privacy at both the individual and governmental levels. Of major concern is the USA PATRIOT Act and the Electronic Communications Privacy Act's Stored Communications Act. The USA PATRIOT Act, more commonly known as the Patriot Act, is a controversial Act of Congress that U.S. President George W. Bush signed into law on October 26, 2001. The contrived acronym stands for "Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act of 2001" (Public Law P.L. 107-56). The Act expanded the definition of terrorism to include domestic terrorism, thus enlarging the number of activities to which the USA PATRIOT Act's law enforcement powers could be applied. It increased law enforcement agencies' ability to surveil telephone, email communications, medical, financial, and other records and increased the range of discretion for law enforcement and immigration authorities when detaining and deporting immigrants suspected of terrorism-related acts. It lessened the restrictions on foreign intelligence gathering within the United States. Furthermore, it expanded the Secretary of the Treasury's authority to regulate financial transactions involving foreign individuals and businesses.

The Electronic Communications Privacy Act's Stored Communications Act is defined in the U.S. Code, Title 18, Part I, Chapter 121, § 2701, Unlawful Access to Stored Communications. Offenses committed under this act include intentional access without authorization to a facility through which an electronic communication service is provided or intentionally exceeding an authorization to access that facility in order to obtain, alter, or prevent authorized access to a wire or electronic communication while it is in electronic storage in such a system. Persons convicted under

this Act can be punished if the offense is committed for purposes of commercial advantage, malicious destruction or damage, or private commercial gain, or in furtherance of any criminal or tortious act in violation of the Constitution or laws of the United States or any state by a fine or imprisonment or both for not more than five years in the case of a first offense. For a second or subsequent offense, the penalties stiffen to fine or imprisonment for not more than 10 years, or both.

What Are the Key Characteristics of Cloud Computing?

There are several key characteristics of a cloud computing environment. Service offerings are most often made available to specific consumers and small businesses that see the benefit of use because their capital expenditure is minimized. This serves to lower barriers to entry in the marketplace, since the infrastructure used to provide these offerings is owned by the cloud service provider and need not be purchased by the customer. Because users are not tied to a specific device (they need only the ability to access the Internet) and because the Internet allows for location independence, use of the cloud enables cloud computing service providers' customers to access cloud-enabled systems regardless of where they may be located or what device they choose to use.

Multitenancy⁹ enables sharing of resources and costs among a large pool of users. Chief benefits to a multitenancy approach include:

- Centralization of infrastructure and lower costs
- Increased peak-load capacity
- Efficiency improvements for systems that are often underutilized
- Dynamic allocation of CPU, storage, and network bandwidth
- Consistent performance that is monitored by the provider of the service

Reliability is often enhanced in cloud computing environments because service providers utilize multiple redundant sites. This is attractive to enter-

9. <http://en.wikipedia.org/wiki/Multitenancy>, retrieved 5 Jan 2009. Multitenancy refers to a principle in software architecture where a single instance of the software runs on a SaaS vendor's servers, serving multiple client organizations (tenants).