Archie Reed • Stephen G. Bennett

Silver Clouds Dark Linings

A Concise Guide to Cloud Computing

Archie Reed
Stephen G. Bennett

Silver Clouds, Dark Linings

A CONCISE GUIDE TO CLOUD COMPUTING

PRENTICE HALL

Upper Saddle River, NJ • Boston • Indianapolis •
San Francisco • New York • Toronto • Montreal • London
Munich • Paris • Madrid • Cape Town • Sydney • Tokyo •
Singapore • Mexico City

"Geek and Poke" cartoons used with permission from Oliver Widder, via CC2.0.

Visit us on the Web: informit.com/ph

Library of Congress Cataloging-in-Publication Data is on file

ISBN-13: 978-0-131-38869-7
ISBN-10: 0-131-38869-X

Associate Publisher
Greg Wiegand

Senior Acquisitions Editor
Katherine Bull

Development Editor
Susan Zahn

Managing Editor
Kristy Hart

Project Editor
Lori Lyons

Copy Editor
Keith Cline

Proofreader
Apostrophe Editing Services

Technical Reviewer
Jonathan Penn

Publishing Coordinator
Cindy Teeters

Cover Designer
Alan Clements

Cover Illustration
RapidEye/
iStockphoto.com

We would both like to thank our good friends and industry colleagues for providing us with time, insight and coffee throughout the development of this book. Special thanks to Jonathan Penn who worked as technical editor, and to Senior Acquisitions Editor Katherine Bull and the team at Pearson Education who helped us work through this project.

ARCHIE:
For my family, Katrina, Danielle and Jade.

STEVE:
To my supportive wife Marie and my children Sam and Sophia for allowing me to take some family time to work on this book.

Acknowledgments

Special thanks to Jonathan Penn who worked as technical editor, and to our Editors Katherine Bull and Lori Lyons along with the whole team at Pearson Education, who helped us work through this project.

A special thank you to Oliver Widder for allowing us to use a select set of Geek and Poke cartoons.

We would both like to express our personal thanks to our good friends and industry colleagues for providing us with time, insight, and coffee throughout the development of this book. With this in mind we would like to expressly thank Ada Lovelace, Alan Turing, Andre Durand, Charles Babbage, Chris Whitener, Christian Verstraete, Chris Hoff, Clive Sinclair, Cutter Brown, Darian Shimy, Darren Kruse, Dave Champine, Doug Cavat, Eran Fiegenbaum, Eric Norlin, Gary Terrell, Ian Huyn, Jason Witty, Jim Reavis, John Howie, Keith Cosgrove, Laura Yecies, Marc Benioff, Marco Casassa Mont, Martin Duursma, Martin Kuppinger, Nils Puhlmann, Peter Coffee, Peter Mell, Phil Windley, Rebecca Lawson, Rob Zuber, Russ Daniels, Sharon Weinbar, Tim Berners-Lee, Trisha Gross, Werner Vogels, Ross Jimenez, and Winston Bumpus.

About the Authors

Archie Reed is HP Chief Technologist for Cloud Security. He is a 20-year experienced manager and technologist, with a wide range of leadership, architecture, product, R&D and implementation experience gained in high profile environments. Archie has worked to deliver both commercial and internal business solutions, as well as managed both engineering and corporate development for multi-million user, multi-tenant service provider. Archie has been an involved in multiple standards efforts including OASIS to the Cloud Security Alliance. Archie is a regular speaker at executive events, conferences and analyst meetings on topics including Security, Privacy, Cloud Computing, Identity Management and Business Technology Optimization. Archie's previous books include "The Definitive Guide to Identity Management" (Realtime Publishers 2003), "Migrating to Windows 2000 and Exchange 2000" (Realtime Publishers, 2001), and "Implementing Directory Services" McGraw Hill (2000), alongside many white-papers and magazine articles. Before becoming an IT professional Archie held prestigious roles as a croupier, a barman and a roustabout.

Stephen G. Bennett is a Senior Enterprise Architect at Oracle, and previously with BEA where he was the Americas SOA Practice Lead within BEA's consulting division. Stephen is a 25-year experienced technologist focused on providing thought leadership, best practices, and architecture guidance around SOA and cloud computing. Stephen has spoken at several conferences and customer events. Stephen has co-chaired a number of working groups within the Open Group organization around SOA Governance and TOGAF/SOA. Stephen has delivered technology management and enterprise architecture consulting to many enterprise customers. Before becoming a consultant, Stephen was involved in delivering enterprise-wide mission critical systems within the investment banking industry.

We look forward to your feedback:

Twitter:	@concisecloud
Email:	concisecloud@gmail.com
Web:	http://www.concisecloud.com/

Table of Contents

Preface

Shift Happens!

This book is written for business and technology leaders (CEOs, CIOs, CISOs, Enterprise Architects) wanting to take advantage of the opportunities "cloud services" create, without being taken advantage of.

Your job is to prioritize your business investment choices based on strategic value and the financial and growth implications. Your role is to clearly see shifts before they occur and prepare to make the best moves for your business. In this role you will have been hearing many things regarding cloud—cloud computing, cloud services, cloud everything. You may have heard that this is the greatest solution to your technology ills and corporate cost management needs. You may have heard that cloud is the riskiest approach to business or technology values since the latest financial crisis. Truth always lies in the middle of the extremes. This thing called cloud has gained more than just a marketing ring to it, and it's time to clearly understand what it is and how to use it.

The goal, therefore, is to provide a concise overview of where cloud services best sit with respect to common business and technology models. To guide your approach, we offer a concise discussion on the pros and cons, say the silver clouds and dark linings that exist, and guidelines for using cloud services.

Cloud computing got its name from the symbol used to represent the Internet in flowcharts and diagrams—a cloud. Not entirely useful as a basis for an IT evolution. Less concise today is creating a simple definition of cloud services. The search for a good definition has seen huge increase in searches on Google. Google Trends, as shown in Figure 1, shows the relative meteoric rise of "cloud computing" searches since 2007, and a slow decline in often compared solutions such as "grid computing" and "distributed computing" and the minimal focus on "utility computing."

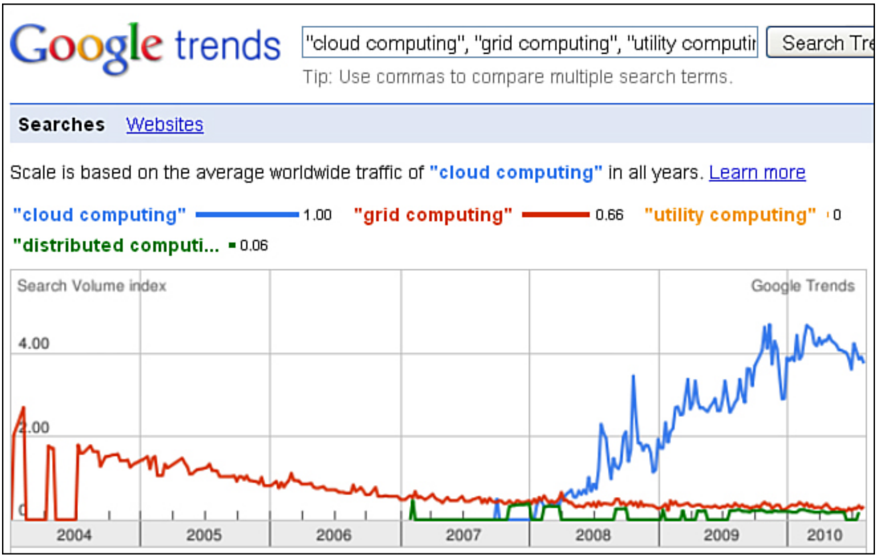

Google trends
"cloud computing", "grid computing", "utility computi
Search Tre
Tip: Use commas to compare multiple search terms.
Searches
Websites

Scale is based on the average worldwide traffic of "cloud computing" in all years. Learn more

Figure I.1
"Google Trends" for cloud computing versus related terms

We'll forgo the blind men trying to understand an elephant by touching various parts and coming up with completely different descriptions. Instead, consider a cloud-related analogy—say, a group of skydivers—all preparing for their dive hoping for clear skies but facing patchy cloud cover. They pack their chutes, maybe relying on an expert to do so—mitigating one level of security concerns. They get aboard the aircraft and begin to ascend. They have a choice—dive or not. Those who dive exit the aircraft at different times, each descending separately, but looking at the same clouds. The skydiver above sees a wonderful collection of fuzzy white shapes, soft and welcoming. The skydiver entering the cloud is buffeted and suddenly loses visibility, blinded by the cloud itself. Passing through the cloud, the skydiver underneath sees dark and ominous coverage, and the sun is completely blocked.

Dramatic? Maybe, but the situation remains that everyone has an opinion on clouds these days along with a definition, rather a description of what clouds are. We plan to offer you guidance to manage that challenge, and take advantage of the situation.

This book does not focus on any of the low-level technical details of cloud solutions, such as the application programming interfaces or the choices of particular technologies per se for several reasons:

Cloud service choices and approaches should be strategic—and while core technologies require attention, the arena is changing fast in terms of players, the interfaces, and the possible opportunities, such that it would be both foolish and inconsequential to spend time on their specifics.

The cloud computing model(s) your organization chooses has a fundamental impact on your business strategy and visa-versa. As such, these require a business level decision and a clear understanding of the shifts in business architecture, and more critically, the business models that result.

Why is this important?

The push for business and technology teams to look for every opportunity to better manage and optimize cost control, business agility, quality, delivery times, and more is immense and pervasive. For enterprises of all sizes, cloud services offer a way to deliver on these goals and minimize capital expenditure, minimize real estate requirements, focus on business priorities, and deliver better results more quickly. A CIO may choose to dismiss web-based services like Facebook and Twitter as toys, or Salesforce.com as an unsafe place to store critical business data. The issue is that can blind you to understanding the change in architecture that has let these companies deliver their solutions in a more efficient, expedient, and exceptional way. These approaches also show how data can be separated from the application showing true service orientation and platform delivery models.

Startups must be aware that in the venture capital and angel communities, a key watchword is cloud… Are you a cloud solution? What is your cloud strategy? If you plan to build a solution and are looking for funds to build out a data center, your solution had better be beyond compelling!

For enterprises, it is entirely possible to begin using cloud services for some technology and back office requirements—email, project management, customer relationship management, sales force automation, and so forth. For startups the situation exists that unless they are a technology focused business, it is arguably critical that they utilize cloud services for technology and back office requirements.

However, cloud is not all about technology. Cloud computing is a disruptive technology that will drive us toward a much more service-based model of delivering business services, and it is driving a huge change to business models alongside social, economic, and political landscapes.

A final point: There are numerous popular and well-known providers of public cloud services today, including Google, Salesforce.com, Amazon, Rackspace, Microsoft, and so on. While we will refer to these and others in examples, this book is not intended to provide you with in-depth reviews of how and where each vendor fits into the cloud ecosystems. The goal we have is to provide you with enough knowledge to understand what vendors are talking about, what is important, and what is not.

Let’s begin to enter the clouds and understand what makes sense in your situation….

PART I

Shift Happens!

IN THIS PART

CHAPTER 1

Introduction to Cloud Computing

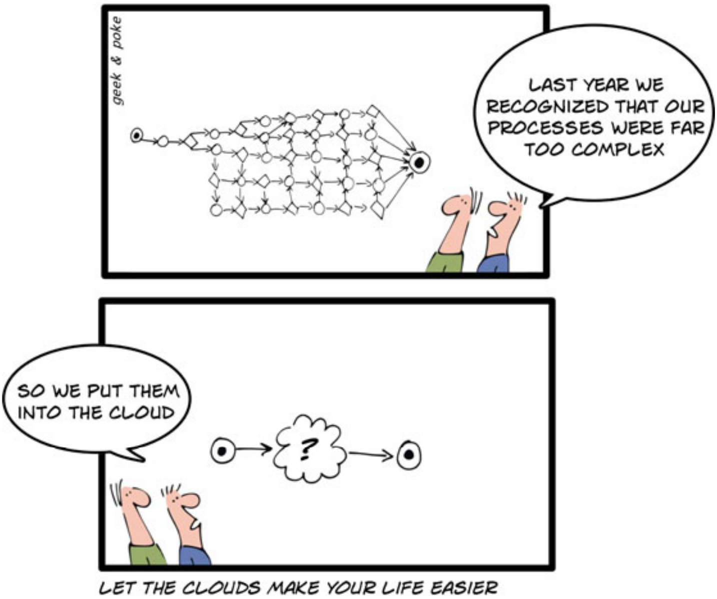
geek & poke
LAST YEAR WE RECOGNIZED THAT OUR PROCESSES WERE FAR TOO COMPLEX
SO WE PUT THEM INTO THE CLOUD
?
LET THE CLOUDS MAKE YOUR LIFE EASIER

Put everything in the cloud—Geek and Poke

Cloud services are arguably the most rapidly growing and evolving approach to delivering applications and services from anywhere to any customer, on any device. A shift is happening with cloud computing that spans the realms of technology and business; a shift that will dramatically change business and how it uses technology to deliver on its requirements. Are you ready?

The Cloud Services Market

Cloud is a logical but fundamental shift in how individuals, enterprises, governments, and more conduct business, interact, and use technology. The ability to have specialized tasks undertaken by third parties is the way in which business has evolved for decades. Think of FedEx for logistics, supply chain, and transport services; ADP for payroll, HR, and benefits administration; the Big Four accounting firms for tax and audit capabilities; or one of the many production facilities located around the world. This ability to hand off critical tasks that can be done more efficiently by a third party, whether they are core or noncore to your business, is a common business model and is how cloud services can benefit you, too. There are several dimensions to cloud computing. Commonly, you will experience sales pitches in terms of public and private cloud solutions—public clouds being solutions offered by third parties, and private clouds being cloudlike solutions you implement within your own data centers. Regardless of where the cloud service is housed however, the benefits are found in being able to pick and choose the most appropriate service when needed, and your business becomes focused on optimizing your own unique IP, business methodologies, and capabilities, while linking in the nonessential services from the best source. It is about delivering quickly and supporting your operational agility.

And it is critical that you understand how you can take advantage of the best opportunities for your organization, too, because your partners and competitors are likely already doing so. A study in August 2009, by F5 Computing of more than 200 mid- to large organizations found that 80% were in trial stages for public and private cloud services deployments for their businesses. Organizations are adopting cloud services aggressively, as detailed in Figure 1.1, with 50% reporting that they have already deployed a public cloud services implementation. Consequently, cloud services are also meriting budgetary consideration, with 66% of respondents indicating that they have a dedicated budget for cloud services initiatives.

Most types of organizations can benefit from cloud services. Large enterprises can often find private clouds compelling because they deal with the maintenance or replacement of legacy systems, cost management, the requirements to launch new services faster, and similar broader competitive issues. Small companies and start-ups can find it easier to make use of the newer business solutions and offer new services to compete with established or much larger competitors. Almost all organizations experience business pressures that can be alleviated through the right application of cloud services.

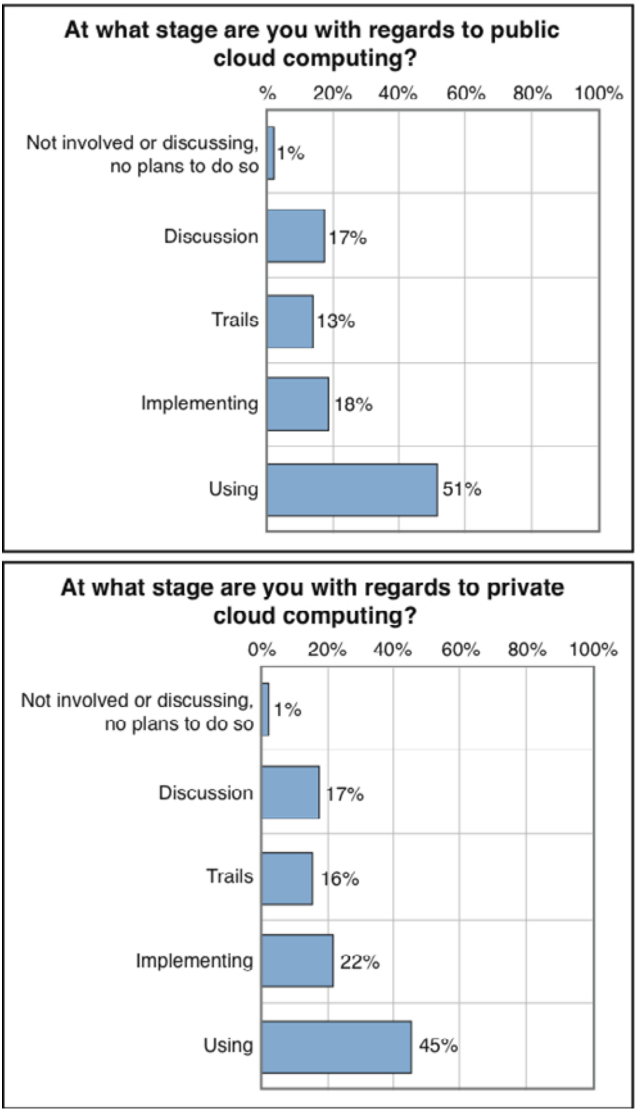

FIGURE 1.1
Stages of use for a public cloud and a private cloud

Legacy solutions must provide a baseline of capabilities, from supporting existing data to providing appropriate new or improved functionality. In addition, until cloud services are seen as being a dominant model for IT delivery, the use of cloud services may be politically sensitive in some organizations, for either valid regulatory, governance, or security reasons, or alternatively, from a job-security perspective. (In Chapter 12, “Creating a Successful Cloud Roadmap,” we discuss the chasms that need to be crossed for different types of organizations, and this is one area in which organizations require best practices before moving forward.) These organizations are definitely more likely to use or be aligned with private clouds, because IT departments try to leverage internally cloud architecture benefits to optimize their data centers. This is often portrayed as, and sometimes parlayed into, an entrée to more-public cloud options.

New solutions have an advantage of generally being able to be architected to use new technologies. This is certainly true for most start-ups benefiting from public cloud offerings. In particular, infrastructure offerings of storage, compute, and networking enable a start-up to create its solution without significant investment in such hardware and its related installation and management requirements. The same potential is there for larger organizations that need immediate capacity without a hit on capital expenditure.

Governments are seeing similar reasons to chase cloud solutions. On September 15, 2009, Vivek Kundra, chief information officer (CIO) within the U.S. Office of Management and Budget, gave a talk at NASA Ames Research Center on the administration's long-term cloud computing policy. In that talk, Kundra noted that of a $77-billion federal IT budget, the U.S. government spent $19 billion on infrastructure alone. The key goals were noted as cutting costs and reducing the environmental impact of the government's computer systems. Citing examples, such as in doubling of federal energy consumption between 2000 and 2006 and duplication of efforts and associated costs across agencies, Kundra saw cloud computing as an incredibly strategic force to mitigate these challenges.

October 2009, IDC released its "IT Cloud Services Forecast: 2009-2013,"[1] and estimated that of the $400 billion customers would spend on IT, $17.4 billion (5% of spend) will be consumed as cloud services. By 2013, customer spending on IT cloud services will grow almost threefold, to $44 billion (10% of spend). While acknowledging there are risks, the expectation is that few mission-critical systems will be moved to the cloud, but significant benefits can be gained elsewhere through nonessential and controlled approaches.

January 27, 2010, the U.K. government announced its strategy to create a private governmental "cloud computing" solution. As reported in the *Guardian,*[2] this is "part of a radical plan that it claims could save up to £3.2bn a year from an annual bill of at least £16bn." In one example of expected benefits, they note that "cloud-based infrastructure could cut costs of government computing significantly and also satisfy its drive for a 'green' agenda by reducing power usage. The Inland Revenue, for example, is presently seeing a huge demand for its online tax return system—but that peaks every tax season and then drops substantially."

Obviously, the goal is to support the peaks and valley's as needed, and share the resources among other departments throughout the rest of the year.

Fundamentally, the market for cloud services is nascent but growing explosively due to a combination of unbridled exuberance, and even more important (as we examine in this book), a compelling set of business drivers. Guy Rosen's State of the Cloud for May 2010[3] shows that of the top 500K sites worldwide, more than 3,000 sites were hosted by cloud infrastructure service providers as of April 2009 and more than

5,000 sites by April 2010. That’s around 40% growth year over year. But, as noted in the same report for March 2010, cloud solutions constitute just 1.01% of the sample. Looking at the higher-level cloud services is more challenging as a whole, but almost every version of the metrics shows significant growth, too. A poster child for cloud services is Salesforce.com. Their 2010 annual financial report showed year-over-year worldwide growth of 17,000 corporate customers to a total of 72,000, just shy of 31%.Fiscal year revenues for 2010 were $1.3 billion, a 21% year-over-year increase.

The growth in cloud services based on vendor metrics so far is remarkable, and the room to grow is immense. The business drivers for cloud services include intense economic pressures and harsh realities being experienced globally, time-to-market concerns, competitive pressures, criminal threats, and more.

Cloud services have achieved a level of awareness faster and greater than many previous technology solutions. Cloud services are having a global impact in so many aspects of the business world right now, from individuals to global corporations, from small businesses to the largest of governments, from the richest nations to the poorest. Even senior executives are asking their CIOs what the “cloud strategy” is.

However, confusion surrounds what cloud services are, and how to best capitalize on all the options available. With these factors in mind, it is also the case that those interested and ultimately influencing a cloud strategy range well beyond technology professionals. The use of cloud services will increase significantly as a result. So, there are lots of silver clouds offering huge cost savings, speed of delivery, and more. However, there are some dark linings in those clouds, and risks are being ignored as the allure of cheaper solutions becomes a focus. The goal of this book is to help you develop the best approach for your organization to get the most from cloud services solutions.

Cloud is not a panacea!

It is neither possible nor sensible to wholesale move your entire enterprise to using cloud services and thus prosper. Established business will have existing, often purpose-built infrastructure that they depend on. Larger organizations will be especially aware of their existing systems on which they depend and cannot change in an instant. Concerns about performance, reliability, availability, and security are often mentioned as barriers to adopting cloud services, and the subsequent requirements must be understood to successfully manage any migration and the associated risks. This generally requires long-term planning and project management. Smaller organizations, start-ups especially, are looking for and are able to gain immediate benefits from cloud services. They can adapt and manage the risks because the cost benefits have significantly greater weighting in such evaluations.

It is sensible to look for a combination of tactical and strategic moves to take advantage of the opportunity cloud services offer. It makes sense to focus on key initiatives and requirements that can be met by cloud services. It makes sense to piece together the right parts of cloud services to improve your business processes, speed up system and product delivery, or even create a completely new product or business. Consider how small concepts and capabilities joined together can create something incredible!

Robert Kearns invented the intermittent windshield wiper in 1963, and filed his first patent around the technology in 1964. After showing his invention to several car companies, Kearns saw the concept stolen and patents infringed when major car companies started to roll out their own. The road to common use of the intermittent wiper and the subsequent decades of lawsuits against Ford, Chrysler, General Motors, and Mercedes for patent infringement forms the basis of the 2008 film *Flash of Genius.* The courtroom scene was compellingly watchable, as Kearns argued against the Ford lawyer's charge that the patent was invalid because it was an obvious use of existing parts. Not so obvious is a core requirement for a valid patent. This argument was countered by Kearns, who showed that although it may have been made of common components, the resultant solution was far from common, but rather gestalt.

Whereas Ford asked a scientist to testify that the invention was a simple set of existing circuits, Kearns pointed out that when Charles Dickens wrote the classic *Tale of Two Cities,* it was not the use of common words that made it great or original, it was the arrangement of them into something new.

In many respects, cloud computing can be seen as a set of simple components, technologies, and processes, itself built upon a legacy of more common ones. Yet with a flash of genius, the cloud can deliver new, unique, and incredibly valuable solutions. The cloud offers an immense wealth of choice components and services for enterprises of all sizes to build new things in new ways.

Kearns, however, did not completely change a historical business model. For that, we can look at a much more strategic and game-changing example.

In 2001, Apple Computer introduced the iPod. Some considered the iPod a simple MP3 player that would need to compete with a multitude of existing products, ultimately appealing primarily to the Apple zealots. Its unique selling point was a new method to control the device called the scroll wheel. Less compelling to many at the time was the binding of the iPod to a simple media management tool called iTunes. However, Apple's combination of hardware, software, and services created a new market model that monumentally changed the music industry.

Regardless of anyone's specific proclivities or polarized position about Apple and its walled-garden approach to the business, it is clear that Apple has won a majority of the online music market for now.

Up until 1999, record companies worldwide largely owned their domains, managing the majority of music production and distribution from the artist through to the consumer. By 2005, the market had massively shifted, and those same dominant companies were beginning to reel from continuous hits to that dominant position. Customers could download music from the Internet. Furthermore, in another part of the industry, artists if they so choose could begin to create, promote, and distribute their own product on the Internet, organizing their own licensing for products such as T-shirts and tchotchkes. The heavy lifting of organizing tours and bookings was now possible through direct online access to venues, unless those venues were locked into contracts, of course.

The market changed and control was being eroded. Music sales were declining, and piracy was an easy target to blame. The reality was much more complex, however.

The once-dominant companies had failed to observe and respond effectively to a multitude of societal, technological, economic, and business models

Quite simply, Apple created an environment that people wanted to be part of…an experience. The choices remain for anyone to avoid or vacate the Apple environs, its unique approaches, and rights management mélange. Microsoft offers an arguably worthy competitive vertical solution in Zune, comprising hardware, software, and services that match Apple's. There are many individual component solutions that when combined offer the same sorts of capabilities, yet the Apple solution is still compelling enough for most to remain loyal. Moreover, it is not just a technical or business model that makes it so. It is a mystique and market presence that allows Apple to maintain such a following, never mind the accessory and secondary markets that have been created in their wake.

Meanwhile, back in the broader music industry, the industry group, primarily in the United States, responded to the market threats with measures that have been decried as everything from draconian and misdirected (suing individuals, lobbying political support to create or update laws, and more). As a result of not responding to a shift in the market, they lost value and their position of power.

Hey, shift happens! Be ready for it.

The movie industry has faced a similar problem. Because of the move to digital distribution of their content that can be easily moved over the Internet, users can more easily access it through this new means. Unless a business model is found that embraces and extends this reality, the movie industry risks suffering a similar experience, arguably in a much shorter time frame. They are experiencing unprecedented shift.

Worldwide, many are excited by the prospect of owning Amazon's Kindle. Although there is much to admire in the hardware in the Kindle device, Amazon is attempting to create a vertical market juggernaut like Apple's. However, competitors are coming in thick and fast to try to own this new vertical market, using a similar model, and therein lies another key point to be made. The benefits of cloud services can be capitalized on often without requiring significantly new intellectual property. However, business models will be broken as a result.

For now, let's be clear: This is not the story of patent litigation or legal issues. It would be a mistake to say that cloud will devastate the landscape of business. However, it is not a stretch to say that the landscape will shift radically in the next ten years because of how cloud services will be used to revolutionize markets. Here, we exhort you to critically consider how your business will be impacted by the cloud approach to using resources, and prepare for some of the most significant changes to business processes and opportunities for change in decades.

Cloud Services Benefits

Common attributes of cloud offerings include massive scalability, near-immediate availability and provisioning, increased cost management controls, and more. However, while we consider the benefits here, a number of dark linings lurk around our silver clouds, and throughout this book we examine them in relation to the usage models of cloud services. Each organization will determine different sets of benefits and risks however. So, your mileage *will* vary.

The definitions of *cloud services* and related cloud computing architectures—and there are many—span a huge range of opportunities and architectures. You may choose to source cloud services internally or externally to your organization, and it is our position throughout this book that most organizations will end up with a hybrid mix of options. Therefore, not all the benefits, or risks, will apply to your situation. Part of the goal in this book is to try to isolate those areas that vary and provide tools to determine your best path forward.

Therefore, as with all things, the benefits and risks attributed to cloud should be considered relative to your current circumstances and measured against your capabilities in relation to any strategic constraints and opportunities that exist. Let's consider building cloud up before we break it down. So, before we detail how cloud solutions are defined, consider the potential benefits from the business and technology viewpoints.

Benefits of Cloud Services

The benefits ascribed to cloud span both business and technology spheres. For business leaders, your ability to maintain or gain agility and your management options are greatly enhanced. For an IT department trying to deliver services to support the business, cloud services offer a new way to architect and source solutions. By finding cloud services that match noncore delivery of IT services, the IT department can concentrate on finding and delivering the best services for the business more effectively.

At a high level the benefits of cloud services can be categorized as:

- Agility
- Business focus
- Cost and budget control
- Scalability and capacity management
- Governance and compliance
- Security
- Optimized infrastructure
- Isolation
- Mobility
- Refactorization

Agility

From a business perspective, there is much more to consider today beyond your ability to manage your core business and deliver great and timely products and services. Today, competitive pressures, marketing challenges, budget issues, and more are considerable requirements. Your ability to manage situations quickly and efficiently is the key.

The biggest benefit of cloud computing to business today can be framed in terms of agility. Cloud services can offer huge savings in terms of time (for example, when IT capabilities must be delivered quickly). Scaling up or down with cloud services does not usually require additional hardware or software. Cloud services offer minimal setup time, minimal time to scale, and less cash outlay. This is because as a business model, cloud service providers generally host massively scaled systems' capacity that can be switched on upon request.

Suppose, for instance, that you need to scale rapidly for a new project or a seasonal rush. Companies can model these situations using internal resources, but likely at some point they will need to expand beyond that capacity. A decision is made whether to use an external provider to fill the gap; in the world of cloud services, this is called *cloud bursting.*

Cloud services, as a concept, are available over Internet technologies and enable us to interact or consume them from almost anywhere on any device. While issues of form factor and communication speed create some limitations today, the business benefit of being able to bring key resources to bear on a critical or time-sensitive problem is recognized as a huge benefit to the agility of any business. Having mobile and remote capabilities allows organizations to recruit employees/contractors who can deliver but who cannot or will not travel to their physical locations. Popularized in the 1990s, offshoring was a first example of this business transition: Business services could be offered from anywhere. However, the advent of cloud services means that more capabilities are available to you and to those providers (and from them, too).

Cultural issues aside, web conferencing services, such as the pioneering WebEx, show how product demonstrations no longer require someone to be physically in the room to represent the company. There are now a multitude of meeting options ranging from Citrix's GoTo services for remote access and support, to HP's Halo room for the "in the same room" meeting experience, and to Skype for making video and audio calls worldwide for much less than traditional telecommunication carrier costs. All of these examples illustrate the opportunity to use technology to deliver business results faster and at highly cost-effective price points.

Servicing your customers at scale is possible only through improvements to the scale and functionality of your service and support capabilities. Support over the Internet is one way to do this. In this scenario, you either expose your support model through a web application or you use a provider who will manage that support through a web application on your behalf (à la cloud). Outsourcers such as EDS/HP Enterprise Services, Centerbeam, and more have been offering these types of support services for more than a decade, but the ability to focus support into web-based solutions decreases the number of staff required to answer phones and deal with people directly.

Business Focus

By using the best service from a cloud service provider, a business can potentially focus more energy and talent on optimizing existing revenue streams and aggressively pursuing the development of new ones.

For example, cloud services can enable businesses to gather information, ideas, feedback and so forth from a much wider set of sources (such as customers, partners) than was ever possible by traditional means. This approach is known as *crowdsourcing.*

Popular crowdsourcing approaches have primarily evolved from the world of Web 2.0 solutions. One business that relies on crowdsourcing is Wikipedia, an online encyclopedia. Wikipedia employs a small organization of less than 50 employees, while utilizing several thousand key volunteers and tens of thousands of other contributors from around the world. While some entries are questionable in terms of veracity, substance or even legality, the overall effort resulted in a much more dynamic and comprehensive set of data than traditional printed encyclopedias could ever match.

Yelp offers the ability to source a set of opinions on a wide range of vendors, from restaurants, to retailers, and more. From this, others can view ratings and comments about those vendors and decide whether to use them. Yelp also shows that these types of solutions can be manipulated (for example, when they gain notoriety, or, when not enough people provide opinions).

Consider a company in crisis. On April 20, 2010 British Petroleum's oil drilling platform, the Deepwater Horizon in the Gulf of Mexico suffered a series of catastrophic failures and collapsed into the water with devastating results to life and nature. The amount of oil escaping was immense. Estimates ranged wildly from 5 to 200 thousand barrels of oil a day, flowing non-stop for over three months. The point here is that BP used crowdsourcing as one approach to deal with the cleanup efforts by creating a "Deepwater Horizon suggestions" page.[4] As of July over 20,000 suggestions had been submitted, and at least 10 had been tested for use. The US federal government also set up a site with information on volunteering to help with the clean-up effort.[5]

This is not to say that crowdsourcing is all perfect. Using services in the cloud like this opens up the potential for anyone in the world with Internet access to "join in" with the crowd. The majority of participants are likely to offer positive input, whereas other individuals or groups are less valuable; some are trolls seeking to make noise, and some are vandals seeking to abuse the system. Mitigation against these and similar issues is centered on access controls, the ability to curate the input, verification processes, and so forth.

Vendors such as Ning, Big Tent, SocialGo, and many others enable for community-based social networking solutions, even crowdsourcing, in a more controlled environment with stricter access controls to the various parts of their services in the cloud. These additional controls can make crowd-based efforts more compelling to community-based or vertical-focused organizations.

Cost and Budget Control

Although the initial costs of using cloud services may appear less, a better expectation should be that cloud services offer more control over costs or better budget management capabilities. Most cloud services enable you to pay on a monthly, weekly, or per-use basis. Choosing a cost-effective cloud service provider can result in significant savings, but more important is finding a cloud service provider than can accurately report usage patterns to you so that you can confirm the accuracy of your billing based on use.

One noted advantage of using public cloud services is the use of operational expenditure (OpEx) over capital expenditure (CapEx). However, understanding the implications of CapEx and OpEx is critical to effectively managing budgets. The difference between buying a house and renting one is the amount of cash that (usually) goes out the door at one time, and the same concept applies here. IT data centers are generally CapEx-intensive, because they require initial outlays of cash upfront to build out. Cloud services are generally booked as OpEx because they are consumed through a services agreement over time. Although CapEx can be depreciated over time, essentially allowing costs to be defrayed against profits, the initial drain on cash at hand is often seen as detrimental and to be avoided in the business world. Having flexibility of where to spend OpEx versus CapEx enables an organization to better manage expenditure such that the costs occur over time rather than in one big lump sum, which enables for better cash-flow management.

The concept of Green IT illustrates the significant costs borne by many organizations. The operational costs associated with building and maintaining data centers and with scaling power, cooling, and even basic hardware requirements are significant. Over time business workloads fluctuate and data centers are often left fully powered on, o whether the infrastructure is used. For example, some organizations run financial systems all year, but they hit peak load only once a quarter or once a year during financial reporting periods. There are many ways that organizations can look to benefit from cloud deployments as they try to optimize and approach a Green IT model. For their own private clouds, the goal is to optimize their usage requirements and cycling systems when needed. Organizations use community or public cloud services when needed, essentially pushing their workloads to the cloud when their internal capacity is reached, or cloud-bursting. Although many cloud service providers do not provide utilization data, there is an assumption that cloud service providers have optimized their infrastructure and management tools to maximize usage.

In many cases, organizations no longer have to hire people to manage system updates and backups and therefore can save on staffing costs and on storage management.

Scalability and Capacity Management

The capability to scale quickly to extreme capacity is not a common attribute or goal of traditional data centers and IT departments. As a result, many businesses experience times when they are hampered by a process that takes weeks or months, sometimes even years, to get new compute resources and applications online. This is despite that many large enterprises have already spent millions on computer hardware and have thousands of servers under management (at least on the books). Cloud computing promises the capability to scale massively in terms of systems, connections, bandwidth, storage, and more on an almost immediate basis. The converse is also important, where the service will shrink back down just as quickly if the need is no longer there. These benefits support cost management goals and enable you to grow in line with your requirements in a linear rather than a reactive mode.

The appearance of unlimited capacity is quite appealing. Cloud services need to be tuned to carefully manage actual capacity against expected requirements, yet deal with unexpected demands when necessary, too. Predictive analysis is critical in this respect, and offering those delivery and management capabilities to customers is something that both IT and cloud service providers must do to deliver this benefit to businesses.

Governance and Compliance

Governance is the process used to ensure that regulations, rules, and mandates are followed within an organization. *Compliance* refers to the facility to monitor and validate that the organization is meeting the governance requirements. Although cloud services introduce potential challenges in matching these requirements, the opportunities to improve an organization's governance and compliance stance can be greatly improved.

Taking a service-based approach to delivering both business and IT functionality that incorporates the governance processes allows for closer alignment to compliance reporting. Sourcing specific, metered services from cloud service providers that specifically match governance models allows for better reporting, more accurate alignment between service usage and fluctuations in business workflow, and ultimately, faster time to market. This also brings up the need for a service catalog and related management tools to ensure usage matches expectations. We will talk about this in Chapters 7 through 10 as this closely relates to overall operations management and security as well.

In some cases, organizations have found that by using a third party to manage their IT services they are better able to deliver services within their governance models and compliance mandates. For small businesses in particular, the cost of hiring to manage these requirements can be significant itself, much like security. Using pooled resources and services allows business to capitalize on the best of that common capability.

It is important to note that although cloud services provide potential improvements to your governance and compliance position, you cannot abdicate responsibility for compliance. While different models of cloud computing architecture actually vary the amount of "control" you and the service provider have over the service, none of them changes your responsibilities to protect data, privacy, or service levels.

This is a critical point made that is reiterated throughout the text: Although you can look to the cloud for specific capabilities and functions to support governance, compliance, and even security, you cannot abdicate these obligations. For example, if a breach of your customer or employee data occurs, the responsibility to report and rectify remains yours. The cloud service provider may help deal with the issue. The cloud service provider may suffer some of the same financial repercussions. However, the cloud service provider is not the owner of the problem. Therefore, any reliance on the cloud's controls and safeguards, and on the governance and compliance practices of a cloud service provider, must be carefully evaluated in line with your governance, and compliance needs to ensure that the solution enhances your posture, instead of increasing either direct or ancillary risks.

Security

Whereas many potential users of cloud services have a visceral or gut reaction that suggest cloud services are not secure, that is certainly not always the case. We discuss security opportunities and challenges in much more detail later in the book, but for now consider many organizations find significant improvements in their security posture by using cloud services. Often, because of their size or financial limitations, small and even medium-sized businesses cannot hire security specialists or respond quickly to security incidents. Using specific best-of-breed capabilities in the cloud to prevent or identify fast-moving or specific threats is a benefit to all organizations, perhaps best illustrated by the vendors who provide antivirus and antispam solutions outside the traditional IT environment.

Optimal Infrastructure

The ability to host multiple capabilities in the same hardware, software, or service allows for the follow-on benefit of optimized use of the infrastructure. Multitenancy as a model for using compute resources has existed since the 1960s, when IBM challenged traditional time-sharing models by adding virtualization with its VM/370 series.

Multitenancy can occur at any or all levels of the architecture, as follows:

- **Virtual layer:** Virtualization provides the ability to create specific environments for each process, application, or operating system. This model isolates everything above the virtual layer itself but allows for the use of pooled resources below that layer, most commonly hardware such as networking, processor, memory, input/output, and storage resources.
- **Application layer:** The user interfaces of most web applications allow for specific fixed graphical and behavioral elements alongside customized elements associated with a specific organization, individual, or function. If one component or functional element fails in the delivery of the interface, it is easier to replace because those elements derived from other parts of the application and can be easily reconstructed.
- **Database layer:** Data for multiple applications, through to multiple customers, can be stored in the same database, and thus allow for the focus on optimal data structures rather than entire infrastructures to support each individual requirement.

To gain the best advantage of cloud requires that each of these architectural layers be considered carefully to determine the best layer or layers to optimize for multitenancy.

Isolation

Almost a corollary to multitenancy is both the ability and requirement for isolation at each of the layers. While an application vendor such as Salesforce.com may use a common database layer for managing customer data, it is essential that the security mechanisms around the application layer maintain isolation between the various customers. Fundamental for any type of cloud provider is the requirement to offer isolation at each level of service exposure to their cloud users. This may be the infrastructure, platform, or software applications. In addition, APIs and management tools must also ensure that isolation exists in terms of identity management and access models, key management and encryption, and user interfaces. This is a specific set of technical requirements that need to be carefully evaluated when using a third-party solution in a private cloud and in any public cloud.

Mobility

The idea of web-based services has been around for many years, as have application hosting and outsourcing. The ability to get to the services from anywhere from any device has been a goal that is finally being broadly realized, and it threatens business models and IT departments who must contend with the security issues associated with data being available on devices either temporarily or long term.

In August 2009, ABI Research[6] released a report that said mobile cloud computing subscribers would total nearly one billion by 2014. The ABI report contained the following reasoning:

> There are two primary reasons why ABI believes cloud computing will become a disruptive force in the mobile world. The first is simply the number of users the technology has the power to reach: far more than the number of smartphone users alone. The second reason has to do with how applications are distributed today. Currently, mobile applications are tied to a carrier. If you want an iPhone app, for example, you have to first have a relationship with the mobile operator who carries the iPhone. If you want a Blackberry app, the same rule applies. But with mobile clouding computing applications, as long as you have access to the web, you have access to the mobile application.

Although there are many arguments against this position, the general direction in cloud services is to support more open standards and therefore the dependence on specific carriers is certainly less than in years previous. This means that offering employees, customers, and partners better access and links to your organization via an exploding mobile world is absolutely possible.

Refactorization

Some like to think of cloud computing as an opportunity to do away with all their existing infrastructure challenges and costs—a “burn it to the ground” or rebuild scenario. This is truly not the real option, especially for any enterprise with more than a year under its belt with existing IT, and especially not for medium to large enterprises with many current or legacy systems in place. The reality is that cloud services allows an IT department to refactor some or all of their existing systems over time and usually take advantage of cost-effective new ways to deliver IT services to the business as a result.

The same concept can also be applied to data center design. Historically, IT has designed data centers using the model of high availability, focusing on repair as a core requirement instead of considering the best mode for recovery to availability. This includes using certified hardware with comprehensive support and maintenance contracts. This includes concepts such as fail-fast, highly integrated systems, and deploying on stable QA tested solutions on a prepared basis with massive change control to ensure the ability to roll back failures.

Large cloud service providers are approaching infrastructure design with a qualitatively different approach. They focus on low-cost commodity hardware where possible. Faults should be simply routed around until some standard recovery can be achieved on a scheduled or even ad hoc basis. This is much more cost efficient in many modes, yet requires a change in thinking that may benefit IT. Modeling can be important here, but if you are not prepared to consider alternative approaches, you will never get to model it. Also note that it is often difficult to model third-party cloud services well, given the potentially vast array of failures that can occur. Traditional IT approaches do not immediately match well to these new architectures and will require some refactoring. We all know that hardware will fail, networks will fail, and an entire data center may fail. Truly, any part of the stack, including the human components (from operators to users), may fail in some way. So what does that mean? We need to monitor different things. It is common to model small failures in a larger system and monitor for those failures (e.g., disk out of space, router offline). Cloud services are modeled as services, and therefore if you refactor how to manage them in this context, you begin to see dependencies across the whole system rather than in terms of just the data center itself, which in turn allows you to focus on core services rather than all services. This change can truly impact IT's ability to deliver and maintain high-availability services.

Summary of Benefits

Using cloud services to supplement or replace IT functions should allow an IT department to deliver more innovative capabilities to the business by focusing more on the service delivery and less on the hardware and software updates. Most benefits derive from the pooled nature of cloud services being offered through multitenancy architectures (or more simply, economies of scale). Costs, risks, controls, and more are aggregated across thousands of customers rather than one individual organization's data center.

Many definitions of cloud computing identify self-service or self-provisioning as a benefit. Although self-service can help make cloud services easier to use, the reality is that self-service requires a mode of operations that includes service catalogs, automated provisioning and de-provisioning, and more to be effective. Therefore, we consider self-service as something that although beneficial in small environment is in reality a result of delivering good service-oriented architectures that may or may not be cloud based in nature.

In summary, cloud computing provides significant opportunities. Thinking these through in relation to your own business challenges is important. Perhaps even more critical, you want to remember that your competitors are also considering how to capitalize on these opportunities.

However, cloud computing has challenges. Your organization may already be ready to adopt and adapt to new technologies or ways of achieving your business goals, but finding success with cloud services may require a mindset that allows you to change how you manage risk and control. To achieve success with cloud, regardless of internal or external options, IT needs to be the service broker and aggregator for the business, providing guidance, cost management, and governance in this new model.

What Is Cloud?

This is the story of how cloud provides us all with the opportunity to truly rethink how we do things, by rearranging things in a new way. The reality of cloud is that it is a culmination of many parts: changing business models, changing Internet functionally (increased speed of transmission, increased reliability), baseline standards, significantly lower costs of hardware and networking.

We consider cloud at the highest level to provide a means by which adequately secured, global, highly scalable, and flexible services can be delivered and consumed using Internet standards through an as-needed, pay-per-use business model.

Cloud services exist as the current incarnation of our evolving technology and business models. Implicit in the term *cloud services* are a number of evolved capabilities in terms of cloud computing, cloud operations, and cloud standards:

- Cloud computing is the infrastructure, including the data centers, networking, and communication standards.
- Cloud operations are the management tools, the APIs, and the many disciplines associated with managing the cloud environments.
- Cloud standards are rapidly evolving from the existing Internet and Web 2.0 standards that have led us to this stage. The goal of cloud standards efforts is to align cloud computing and cloud operations capabilities. In early incarnations, cloud standards have primarily been de facto in nature, but a significant number of organizations are working to evolve them (or to create de jure standards for the benefit of all).

The concepts of large-scale usage of compute resources such as utility and grid computing are derived from traditional utility providers such as power or water. The adoption of utility computing is seen as a good parallel for cloud services, because they have initially been used for noncritical processing.

Let’s start off with our chosen principal definition to cloud and then move to focus on what’s really core to cloud and how you can take advantage of it. While you will find many definitions available, we see the most useful, comprehensive, and popular definition as being from the National Institute of Standards (NIST). The NIST Definition of Cloud Computing Version 15 offers the following definition of cloud computing as defined by Peter Mell and Tim Grance:

> Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of

five essential characteristics, three service models, and four deployment models.

The next three sections present the essential characteristics, service models, and deployment models of cloud computing as defined by NIST.

The NIST Essential Characteristics of Cloud Computing

The essential characteristics of cloud computing as defined by NIST are shown in Figure 1.2.

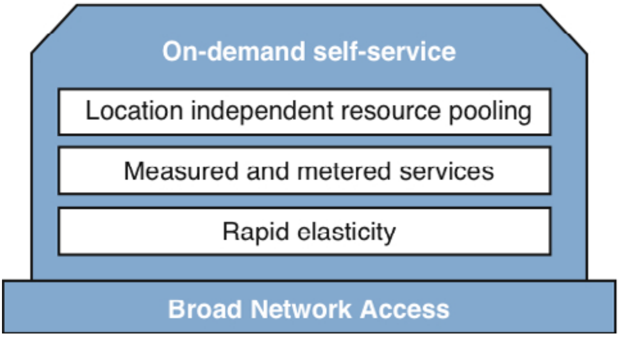
On-demand self-service
Location independent resource pooling
Measured and metered services
Rapid elasticity
Broad Network Access

FIGURE 1.2
Cloud essential characteristics based on the NIST definition

Key to these characteristics introduced here is the concept of multitenancy—the idea that many different applications, users, and even businesses may take advantage of the resources being used, as follows:

- **On-demand self-service:** A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.
- **Broad network access:** Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).
- **Resource pooling:** The provider's computing resources are pooled to serve multiple consumers using a multitenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

- **Rapid elasticity:** Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.
- **Measured service:** Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

The NIST Service Models of Cloud Computing

The service models of cloud computing as defined by NIST are shown in Figure 1.3.

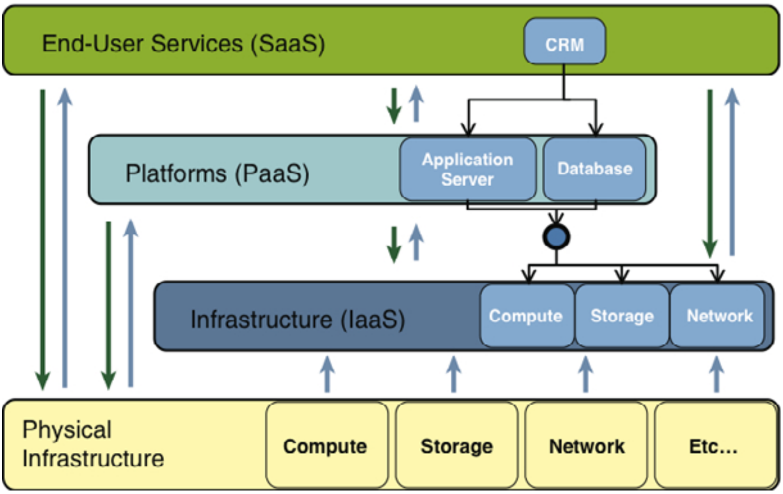

FIGURE 1.3
Cloud service models based on the NIST definition

Each of the service models offers different levels of capabilities and responsibilities to the provider and the consumer of the service. In addition, these service models may utilize a custom architecture, or the physical infrastructure, to exist. These services may rely on a service offering from a lower service model. In this way, cloud services in the higher levels, those being platform and software, may be entirely built atop other cloud services. The most compelling way to think about these models is that if everything has a set of APIs from the lowest to the highest levels, the service-orientated nature of cloud services becomes very clear, as follows:

- **Software as a service (SaaS):** The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

- **Platform as a service (PaaS):** The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.
- **Infrastructure as a service (IaaS):** The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

To elaborate the potential use of lower services by higher level services, Figure 1.3 includes an example SaaS customer relationship management (CRM) application that resides entirely atop the various cloud services deployment options, using an application server and database in a PaaS scenario, which in turn rely on an IaaS option for compute, network, and storage capabilities.

Critically, in Figure 1.3, we also introduce the concept of XaaS. XaaS represents anything as a service. That said, we will try to avoid confusion or pollution of the “aaS” nomenclature by focusing on SaaS, PaaS, and IaaS as the core models for cloud services, often combined into its own acronym of *SPI.*

These models can also be deployed in isolation, utilizing their own compute, network, storage, and related infrastructure. In these cases, the infrastructure is architected to deliver the best performance for the service delivery requirements and not generic capabilities as defined by each service layer.

The NIST Deployment Models of Cloud Computing

The deployment models of cloud computing as defined by NIST are shown in Figure 1.4 and are described in the following list.

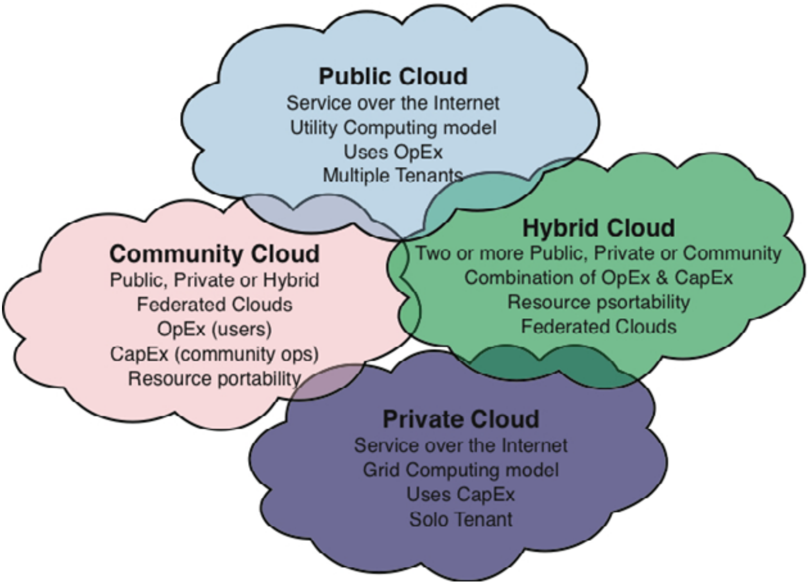

FIGURE 1.4
Cloud deployment models based on the NIST definition

- **Private cloud:** The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premises or off premises.
- **Community cloud:** The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premises or off premises.
- **Public cloud:** The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
- **Hybrid cloud:** The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

What Can Cloud Do for Me?

Cloud computing is about moving services, computation, and data—for cost and business advantage—to an internal or external, location-transparent service or services. By making data, services, capacity, and more available in a service-based model, they can be much more easily incorporated into business and IT processes that can be ubiquitously accessed. This is often at much lower cost, increasing the value to the business by enabling opportunities for enhanced collaboration, integration, and analysis on a shared common platform.

From just these few definitions, we see both commonality and disparity, and those same issues exist across all the many and varied approaches out there. As a result, there is a multitude of overlapping definitions of *cloud.* The analysts have them, the marketers have them, users have them, and standards groups have them, and so on. We have one, too, but we won't get wrapped up in debates about it. Instead, we focus on the capabilities that the cloud model provides.

Perhaps most aggravating when investigating cloud services options is the marketing approach to relabel or brand anything and everything "as a service." This is called *cloud washing.* This just becomes confusing initially and an exercise in futility in the long run as terms become overloaded. We see "storage as a service" and "security as a service" competing acronym-wise with "software as a service." So, one of the advances is to say that each of these options fits into one of the three delivery models, SPI, and be done with it.

To understand the cloud then, it is better to focus on the attributes possible rather than ascribe a specific definition to it. Not only do the definitions change depending on which part you focus on, but also so do the benefits and risks.

A breakdown of these solutions with examples of popular solutions helps to understand the nuances of each layer in more detail (see Figure 1.5).

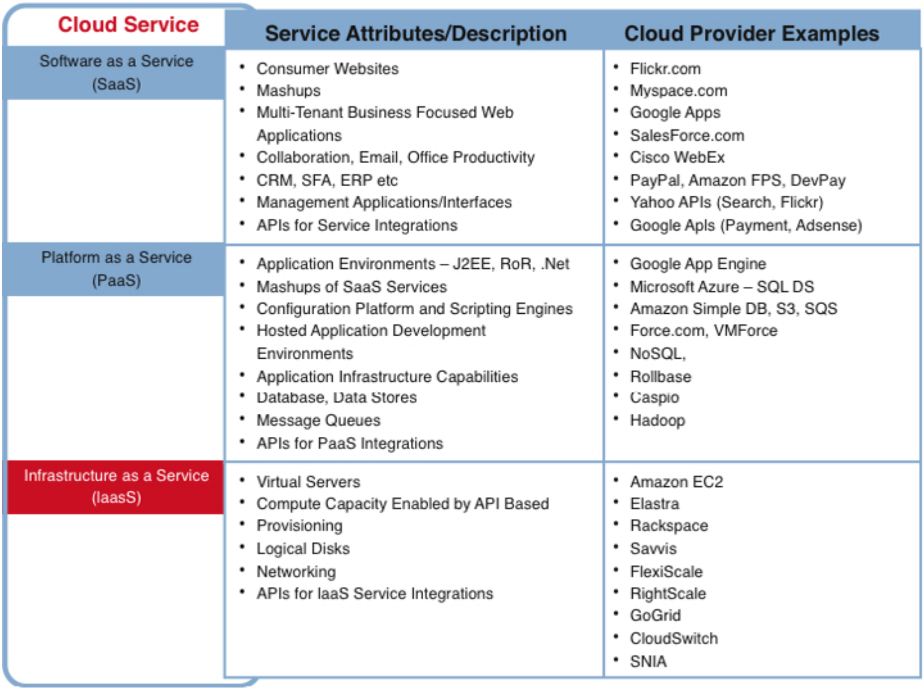

Cloud Service	Service Attributes/Description	Cloud Provider Examples
Software as a Service (SaaS)	• Consumer Websites • Mashups • Multi-Tenant Business Focused Web Applications • Collaboration, Email, Office Productivity • CRM, SFA, ERP etc • Management Applications/Interfaces • APIs for Service Integrations	• Flickr.com • Myspace.com • Google Apps • SalesForce.com • Cisco WebEx • PayPal, Amazon FPS, DevPay • Yahoo APIs (Search, Flickr) • Google Apls (Payment, Adsense)
Platform as a Service (PaaS)	• Application Environments – J2EE, RoR, .Net • Mashups of SaaS Services • Configuration Platform and Scripting Engines • Hosted Application Development Environments • Application Infrastructure Capabilities • Database, Data Stores • Message Queues • APIs for PaaS Integrations	• Google App Engine • Microsoft Azure – SQL DS • Amazon Simple DB, S3, SQS • Force.com, VMForce • NoSQL, • Rollbase • Caspio • Hadoop
Infrastructure as a Service (IaaS)	• Virtual Servers • Compute Capacity Enabled by API Based • Provisioning • Logical Disks • Networking • APIs for IaaS Service Integrations	• Amazon EC2 • Elastra • Rackspace • Savvis • FlexiScale • RightScale • GoGrid • CloudSwitch • SNIA

FIGURE 1.5
Cloud architecture sample capabilities and industry examples

Most of the examples are public cloud options, and the opportunity to fashion your own solution such as these as private cloud options does negate some of the opportunities we have discussed so far.

The addition of cloud-driven processes and web-based services to the SPI framework is intended to illustrate the higher-level business models and processes that can be considered. Furthermore, here we illustrate the term *mashup,* which comes from the Web 2.0 environments. A mashup is a composite web-based service created by mixing together other web-based services, or more appropriately, higher-level cloud services. This creates useful and varied solutions for users, but also introduces management challenges as delivery is affected via distributed capabilities. This essentially becomes a chained service model, and as with all integrations, or should we say "mashups," things can go wrong across the service components.

For a business example more grounded (ahem) in supply chain management, consider the history of the Boeing 787 Dreamliner. Planned to be the most advanced technical and component-based airplane in large-scale commercial aviation, Boeing approached the effort by taking a nontraditional approach to collaborating, sourcing, and integrating myriad components from hundreds of suppliers and subcontractors. The project was beset by delays as a result of too many issues. As a result, the aircraft finally lifted off for its maiden test flight on December 15, 2009 (more than two years after its original schedule of August 2007, and well over budget and suffering from a decreased set of capabilities).

History shows that there were numerous problems for Boeing in terms of supply chain management: integration due to standardization issues, component shortages, quality-control issues, and more. Similar issues come to the forefront when an organization chooses to use cloud services without clear strategic goals and management tools being in place.

To succeed with cloud services, business and IT leaders must recognize and deal with the fact that the role of IT is changing to include much more comprehensive supply chain management and vendor management. If cloud services are to be used, the traditional IT team makeup is incomplete and so must be enhanced with more legal, contractual, and business expertise. Those familiar with outsourcing are in a much better place to manage this new approach to delivering business value.

Cloud services can enable the business to gain much greater control over its IT-dependent decisions, as long as it has the correct management processes and tools in place. This is discussed extensively in Part III, "Life in the Cloud—Planning and Managing the Cloud," of this book, where we examine the approaches and requirements for managing cloud services solutions.

Summary

Ultimately, cloud services offer opportunities to source complete or partial services for all IT and business processes. This move to a multisourced, multicapability, or hybrid market model has risks just like those faced by Boeing (when different suppliers held different responsibilities and different roles, which resulted in several time and budget failures). That is not to say that using cloud services is bad, but rather, a reminder that adequate management of their use and integration is critical.

Key points to consider from this chapter are

- The cloud services market is developing at a rate faster than most other technology and business models.
- Cloud services are changing business models and industries and are creating new opportunities for all.
- There are many definitions of cloud services, and the NIST definition is the most widely accepted definition used today.
- The benefits of cloud services span both business and technology spheres.
- Cloud is not a panacea.

The discussion so far has focused primarily on defining our topic and examining the benefits associated with cloud services. Chapter 2, "Evolution or Revolution?," examines the risks that must be considered alongside those benefits.

Endnotes

[1] http://blogs.idc.com/ie/?p=543

[2] www.guardian.co.uk/technology/2010/jan/27/cloud-computing-government-uk

[3] http://www.jackofallclouds.com/2010/05/state-of-the-cloud-may-2010/

4 http://www.deepwaterhorizonresponse.com/go/doc/2931/546759/

5 http://www.serve.gov/oilspill.asp

[6] http://www.abiresearch.com/research/1003385-Mobile+Cloud+Computing

CHAPTER 2

Evolution or Revolution?

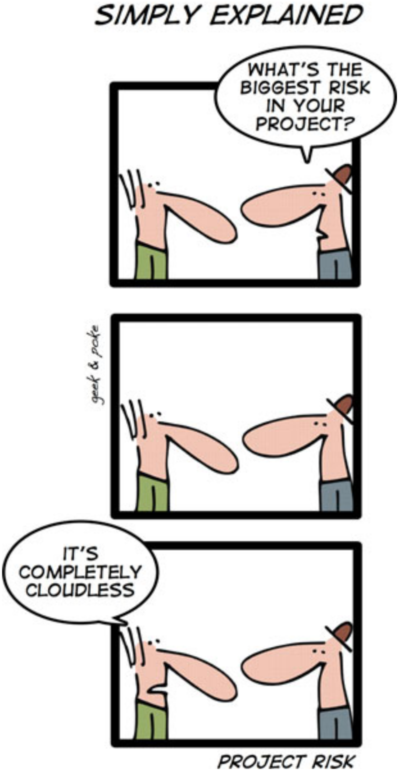
SIMPLY EXPLAINED
WHAT'S THE BIGGEST RISK IN YOUR PROJECT?
geek & poke
IT'S COMPLETELY CLOUDLESS
PROJECT RISK

The biggest risk to any new project —Geek and Poke

Most agree that cloud is an evolution of existing technologies and business models. Those who don't are generally marketing something.

In this chapter, we look at the new and somewhat unique group of capabilities ascribed to cloud services. In so doing, we contrast our discussion that focused on the benefits of cloud services in Chapter 1, "Introduction to Cloud Computing." The approach is to examine the dark linings, those critical concerns that need to be addressed in a cloud initiative.

As noted in Chapter 1, there is huge potential to create revolutionary businesses by assimilating and expanding on the approaches and technologies offered by cloud computing. However, we should not consider this entirely new. Figure 2.1 shows the approaches and technologies that have supported the evolution to cloud services.

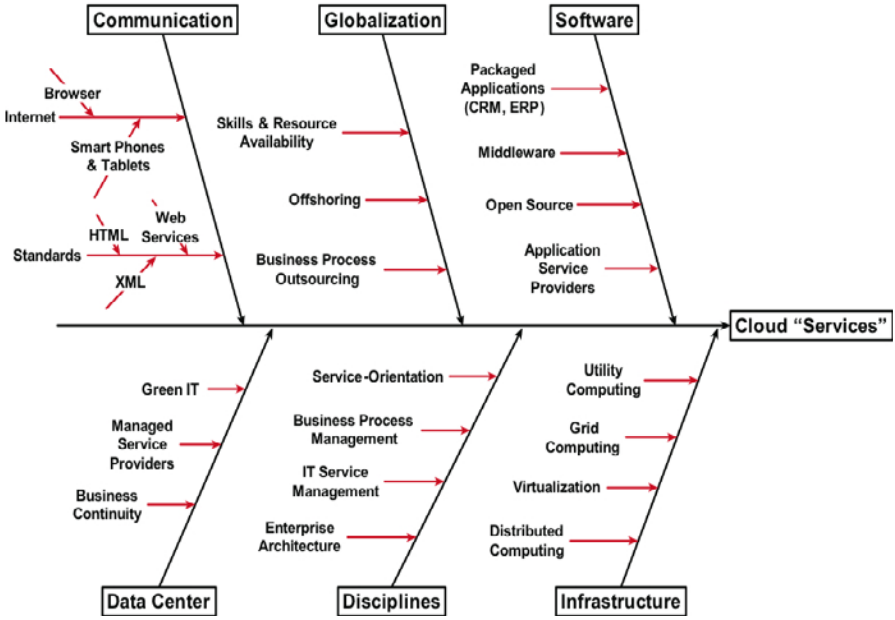

FIGURE 2.1
Sources of the evolution of cloud services

In this chapter, we cover the capabilities that make up cloud computing. We examine these core capabilities and then take a look at the benefits and risks of cloud computing.

Cloud Capabilities

Capabilities that are assumed essential for cloud computing are most often derived from or direct adaptations of previous approaches, most of them detailed in Figure 2.1. For example, virtualization has been around for decades in various forms. For another example, in 2005 Tim O'Reilly detailed several principles of Web 2.0. The first was "Web as the platform." Then O'Reilly noted that the "core competencies" identified included "services, not packaged software," "cost effective scalability," "remixable data source and data transformation," and "software above the level of a single device." These sound similar to how many talk about cloud services. Further, self-service portals, Web 2.0 UIs available on multiple devices have been evolving over the past 10 to 20 years, service-oriented architectures have been a mainstay of large-scale enterprise architectures for the past decade.

So the real question is this: What's new?

The conjoining and combining of these capabilities are included as part of the NIST essential characteristics of cloud computing discussed in Chapter 2. More discussion is warranted to fully understand how to deliver on these characteristics.

- **On-demand self-service:** Cloud services are expected to be easy to request, such that a customer can acquire access to resources or services without having to go through a complex or manually intensive ordering process. This usually means using a web-based management tool, using some simple API calls, or capitalizing on some level of integrated authentication and provisioning capabilities (commonly known as federation).

- **Resource pooling, multitenancy, and shared resources:** To achieve scale and remain cost-effective, cloud service providers must capitalize on the actual capital outlay and ongoing expenses by offering common access to the core infrastructure capabilities (most commonly, compute capacity, storage, and network infrastructure). In private or internal cloud environments, this is less of an issue. However, be aware that providers of public or external clouds are not usually inclined to provide customized environments—yet.

- **Broad network access:** The concept of being able to access cloud services anytime, anywhere, and from any device is supported by broad network access. Without standards to integrate atop common carrier solutions that have been borne from the Internet, this capability will be challenging. At this level, the contention between net neutrality and the costs to build out the broad network capabilities is one of the most difficult equations to balance. Next, this characteristic is the most difficult to support through the private cloud deployment model, as by definition private clouds are generally limited to organizational network boundaries. The contention here is that by offering secure remote access to private clouds, such as VPNs over public networks, mitigates that potential limitation. Device support, however, can be the most challenging. For example, the popularity of the Apple ecosystem is vast, yet supporting enterprise applications require the approval and ongoing support of the Apple app store approval process, something that many have complained about as being arbitrary, slow, and even anticompetitive. The Android ecosystem, however, is promoted as being open and without any approval process. This then creates increased security risks, as there is no standard approach to the assessment of the virtue of any applications. In this characteristic, it is critical that you understand the opportunities and limitations of supporting a specific vendor or set of vendors.

- **Elastic or rapid scaling up and down:** The combination of these capabilities is an expectation that cloud services should provide unlimited expansion in terms of capacity, size, or speed. Although a guarantee of infinite capacity is impossible, it is the role of the provider to manage the capacity of cloud services appropriately to create that potential.

- **Measured services, pay-per-use, consumption-based pricing:** Capacity in cloud services is expected to be easily bought and sold in a pay-per-use model.

An additional expectation is that those services should be immediately accessible when required and immediately released when no longer needed. After cloud services have been provisioned, it is expected that there will be a simple pricing model based primarily on actual usage, billed in simple increments allowing for easier budget control management.

Each of these capabilities is a key tenet of cloud service solutions. The next step is to discuss the overall benefits and risks associated with this new business model and to decide whether the capabilities need to be adopted as a whole to succeed with a cloud services initiative.

Benefits Versus Risks

Most organizations that adopt cloud services realize huge benefits. In September 2009, IDC took a poll at its annual Enterprise Panel to rate the top benefits that enterprises attribute to cloud services, the public cloud in particular (see Figure 2.2).

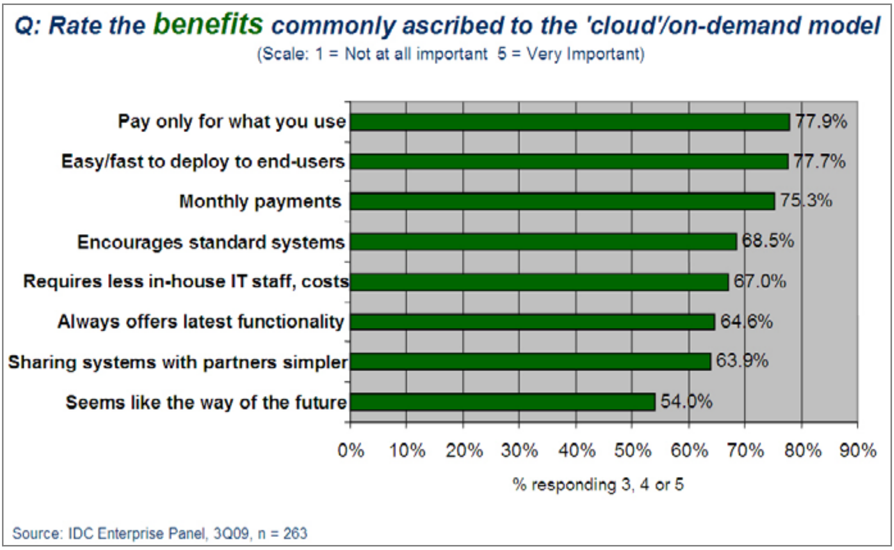

FIGURE 2.2
IDC survey of enterprise customers' benefits ascribed to cloud computing models

The highest rankings related to cost-control aspects such as pay per use, monthly payments, and less staffing and lower costs. Interestingly, there was little reference to business flexibility, increasing the ability to enter new markets, and similar growth opportunities.

These rankings certainly align with the goals of many controllers, CFOs, and COOs. Given the current utilization rates of many IT environments, there is a belief that such services should be an order of magnitude cheaper to run than running the current environments, and cloud services offer that promise.

It is perceived that popular cloud solutions such as Amazon's EC2 (Elastic Compute Cloud), Microsoft's Azure, Google's App Server, and Salesforce.com are proving the point. However, the decrease in capital expenses often sees a similar decrease in other critical areas of concern (e.g., security, privacy, manageability, and compliance). The possible savings may not always be as significant, however, if a company is already using internal cloud solutions, or if the company has virtualized its data center completely to achieve optimal use. Fundamentally, without adequate management and strategic architectures, the benefits will not be delivered or significant.

IDC also surveyed the respondents on the top challenges or risks associated with the cloud model. Nearly 90% of CIO-level respondents identified *security* as the number one concern related to cloud computing. You will see similar reports where the concerns are in different orders. Regardless of the actual order that results, the relative numbers for each show that they are all high-level and common concerns and thus must be considered almost as equals. As the cloud architectures evolve, expect to see these concerns move around, but for now they serve as a good list on which to base our next discussion.

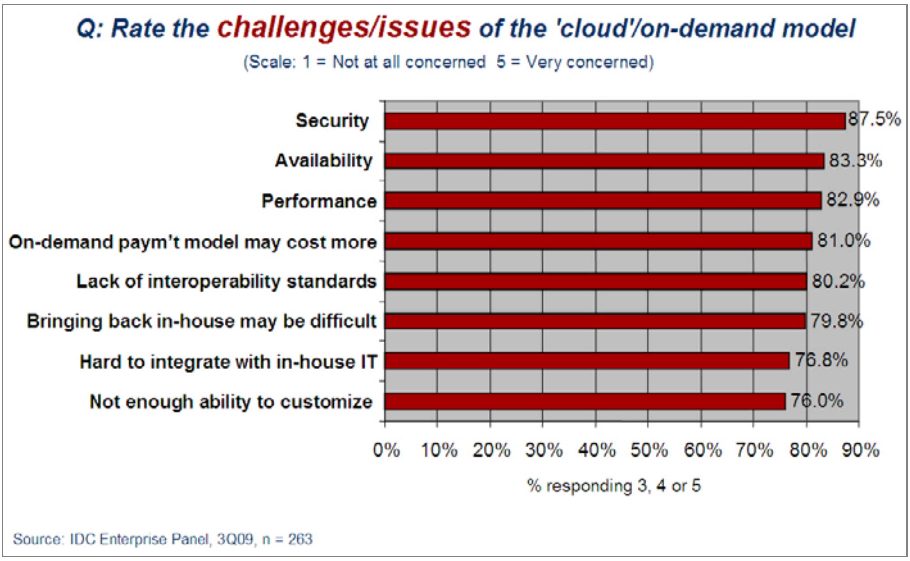

FIGURE 2.3

IDC survey of enterprise customers' challenges and issues ascribed to cloud computing models

The trouble here is that security is a potentially vast and polarizing topic. The high ranking of security could well be the result of a visceral reaction rather than a clearly definable concern. For example, it is commonly assumed that cloud services are less secure than those housed in IT data centers. However, it's not always true. Consider, for example, that many small and medium enterprises do not have security staff. By using cloud services, those enterprises can significantly increase their security posture. The same may be true for many larger organizations if their business units make use of vertical solutions and services or go so far as to avoid IT when acquiring capabilities (what is known as shadow or rogue IT).

Security as a term without context is almost as nebulous as *cloud.* To achieve "security," we first need to determine what data is involved, the value of the service being provided, the deployment model being used from the SPI stack, and ultimately (using these inputs for analysis), the risk and cost-benefit analysis. Together, these will help determine the potential savings, and if there is increased risk, whether it is acceptable risk to the business. For example, savings may allow the business to create new products or services, or enter new markets that decrease business risk.

Availability and performance are closely related requirements, and only specific analysis of uptime and recovery plans can truly determine which solution is more likely to achieve an organization's required levels. Although there have been widely publicized outages for Google's Gmail, Amazon Web Services, and others, the reality is likely that many data center-based applications and services have suffered similar if not more extensive downtime. To analyze whether a cloud service offers better availability than one from IT, both potential providers must be clear on what metrics are critical to you and thus what metrics need to be compared. Uptime of the messaging system, core business applications, public website, intranet, and other key applications are the levels to examine when considering software as a service (SaaS) offerings. Underlying technology is of concern, but the overall statistics impact how business is done. In addition, the accessibility from other devices, locations, and improvements in security posture or performance can be added to the evaluation.

The challenge underlying these top three concerns of security, availability, and performance is a combination of trust, management, and service levels. For public cloud service providers that are new to the market, a lack of established credentials and history of uptime means that trust is the main challenge they must overcome. To do so, the approach is to attract customers with offers of extreme discounts or free services. Realize that these approaches will not work long term.

Cost control requires careful consideration, and the requirement for cost-benefit analysis is almost as critical as a risk analysis to ensure success in using cloud services. The cost benefits of cloud services seem to be obvious initially (e.g. less hardware and personnel costs associated with delivering cloud services). Often forgotten costs when considering the overall return on investment (ROI) of cloud services relate to the migration into and out of the cloud service. Ongoing networking costs are likely to rise as usage increases, but as with all ROI calculations, this is offset by the benefits.

Another risk that could impact your ROI is vendor lock-in. The rapid evolution of offerings is creating significant churn in vendor APIs, UIs, and management tools. One cost will be maintaining alignment with the vendors' offerings as they are managed centrally and updates are pushed to you on the vendor's life cycle, not yours. Until standards evolve, another cost will be anything associated with migrating from one set of APIs to another, especially in the platform as a service (PaaS) and SaaS layers, but information as a service (IaaS) is also susceptible because few standards at this layer truly support cloud computing architectures. In general, the higher in the cloud service stack you are, the more likely lock-in will occur because the services become less generic.

In terms of vendor lock-in, use of public cloud services ties your business and technology life cycle to that of the vendors. As an example, Salesforce.com has offered a proprietary programming environment along with custom APIs in its Force.com platform for many years. Migration from that environment to another would entail a considerable effort and cost (until competitors develop migration tools). Salesforce.com has created a tool to support customer migration from Microsoft CRM to the Force.com platform. That said, the reverse is also possible through a number of solutions. Microsoft, in its common approach, provides toolsets, whereas commercial vendors such as NettMore in the United Kingdom offers a tool called CRM Migrate to facilitate the migration of data from Salesforce to Microsoft Dynamics CRM 4. The understanding, however, is that both these environments support large numbers of customers and warrant development of such tools. That, indeed, is the issue for most of these solutions: Until they become big enough to warrant competitor attention, they will not have standard or commercially available methods to migrate.

Vendor viability is often a concern for smaller providers as well as for start-ups. The possibility of having to migrate in the event of vendor failure is one concern, but it is the risk of using an early-stage cloud service provider. To ascertain a vendor's position, its financial position may be a critical data point. The vendor's ability and willingness to provide may indicate the provider's ability to deal with issues in the future, such as operational issues, e-discovery requests, or even security breaches.

The evaluation of these risks in a comprehensive risk-benefits analysis will not only define your approach, but also feed your requirements for your cloud services operations management approach. To that end, Part III, "Life in the Cloud—Planning and Managing the Cloud," deals with key aspects of governance, compliance, legal concerns, operations, and security in more detail.

By efficiently preparing to use cloud services, real and detailed concerns can be identified and managed through process, planning, risk analysis, and governance. Managing these requirements can be considered in the mode of operations management. In this case, focus on key operational life cycles to clearly understand what changes when cloud services are introduced. In addition, education of key architects, managers, and operational staff is needed.

The cloud services approach suggests that simplifying the frontend is of considerable benefit. However the reality is that if architects and (perhaps more so) developers do not understand distributed systems and service-oriented models, you will fundamentally fail to deliver benefits through the cloud and (often more critically) fail to secure your environment.

As previously noted, consider each cloud service provider's options in terms of cost, features, speed to market, and scalability. For now, consider how each of these applies to specific implementation models available (see Figure 2.4).

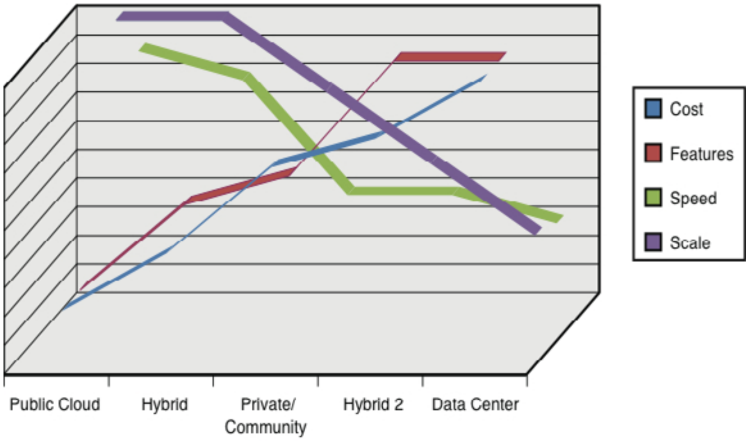

FIGURE 2.4
Comparison of benefits versus deployment options for IT capabilities

A simple alignment illustrates that cloud services provide the greatest speed and scale for the least cost. However, when we look at the features metric, it is also quite low for cloud services. Features often include complex requirements in terms of security, management capabilities, compliance controls, migration facilities, and more. To achieve these requires a cloud service provider to improve its service or for you, the user, to add additional cloud services that offer the required functionality, alongside the necessary integration (much like traditional IT shops do in their data centers, and with a high likelihood of increased costs and complexity).

Fully expect cloud services to increase their features over time, and although costs might not increase significantly, or at all, as features are added, the choice will remain aligned with these metrics. Cloud service providers' primary path to profit is through cost containment for the majority of expenditures that you must deal with when deploying your own data center, as follows:

- Hardware
- Power and cooling
- Provisioning
- Networking
- Security
- Billing
- Operations
- Personnel

These attributes are what can be considered as common to any cloud service. They underlie any implementation and are the focus in any drive to lower costs. This means that a cloud service provider will be pushing to minimize the amount of customization in their environment. While this does apply to the platform and software layers, wherein customization creates increased costs in all the management areas previously listed, the cost implications are acute in the infrastructure where standardization is most cost-effective. This applies to any type of cloud and must be considered as a limiting factor for higher-level usage. For example, some database solutions require direct access to the hardware to optimize their internal workings. With a virtual machine manager between the database and the hardware, such optimizing is devalued.

With much less dependence on these infrastructure capabilities, there will be an increased focus and opportunity in the PaaS and SaaS models. This will result in the following:

- Significant growth in PaaS and SaaS offerings.
- More focus on the attack vectors, or the dark linings of the silver clouds of PaaS and SaaS.
- IT must become a service broker because the traditional IT processes of "develop, deploy, and manage" are changing to include and possibly be subsumed by "source, compose, offer, and manage."
- IT must adapt to using agile approaches for application and service development, as well as for asset and personnel allocation.

The key needs are a combination of business and technology requirements. You will find there is overlap, and the goal is to ensure that the understanding is clear: Success in using cloud services requires multiple disciplines to work together for the benefit of your business.

Cloud offers many of the same opportunities and presents many of the same challenges service-oriented architecture practitioners have had for many years, but on a much more significant scale. Cloud reasserts the need for comprehensive enterprise architecture to be aligned with business architecture; otherwise, the result will be (again) many silos of cloud services.

Summary

The point to reiterate is that cloud services enable you to revolutionize how you deliver and consume business processes (at a minimum) and change business models (on a broad scale).

Can you afford to ignore this trend? The answer is: not for long!

In Chapter 1, we discussed how Apple's strategy in terms of mobile devices, client applications, the App Store, and content delivery service from the Internet resulted in an ecosystem that dominates the consumer space.

Consider how the smartphone market is rapidly changing, and the resultant change in fortunes of vendors in that space. In the context of business email, Research in Motion, with its BlackBerry devices and BlackBerry Enterprise Server (BES), used to own the smartphone market. Some would argue that it still does. However, 2010 is a significantly different market. IT departments, continuously under cost management constraints, are increasingly changing their approach to employee phones and are no longer forcing business users to take a company-provided phone, allowing for much more choice by users. As consumers, the appeal of an iPhone or Android-based smartphone is high. The rank and file are not purely driving this adoption, but more important, usage is being driven by executive management who often embrace and advocate new technologies ahead of IT and the broader business. As a result, Apple and Google are encroaching on RIM's traditional market from below and the side, and illustrate another example where changing business models impact incumbents in negative ways.

One aspect to consider when weighing up the business benefits versus the risks is how your users, customers, employees, constituents, or partners are experiencing the procurement of other services. App stores, slick interfaces, and support services that offer immediate fulfillment are driving their expectations. Cloud services in the enterprise will increasingly be measured against this model, and IT must offer instant provisioning from a portal but remain secure and well managed. As we have discussed, however, the risks exist across the board. Worse, your existing architectural design and operational processes are not entirely adapted to deliver cloud services. Finally, to move forward, you must recognize that the overall threat to landscape is now no longer just technical but a competitive business issue, as well.

Key points to consider from this chapter are

- Cloud services have evolved from multiple technologies creating new businesses outright and new business opportunities for your organization.
- While public cloud service providers may offer solutions that meet your business needs, their business priorities do not always entirely align with their customers.
- IT must adapt to using agile approaches for application and service development, as well as for asset and personnel allocation.
- IT's role will increasingly include service broker requirements.
- Manage risks (dark linings) to allow your business to realize meaningful benefits from cloud services (silver clouds)!

In Chapter 3, "Reflections on the Shift," we review our introduction to cloud computing and prepare to plan for cloud services.

CHAPTER 3

Reflections on the Shift

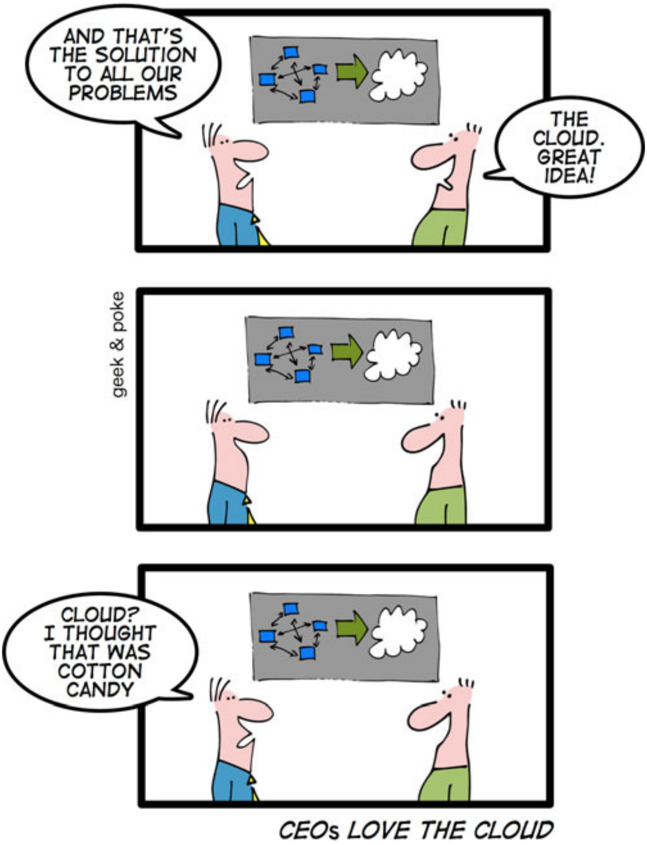
AND THAT'S THE SOLUTION TO ALL OUR PROBLEMS
THE CLOUD. GREAT IDEA!
geek & poke
CLOUD? I THOUGHT THAT WAS COTTON CANDY
CEOs LOVE THE CLOUD

The solution to all your problems—Geek and Poke

Think about your current information technology infrastructure. Think about how you source that technology. What about your business processes? How do these things relate to your financial and business models?

Are you a start-up looking for every opportunity to save time and money? Because start-ups typically are looking for these opportunities, start-ups are generally more accepting of the cloud model and the public cloud offerings specifically. This is the approach that founders, angel investors, and the venture capital community expect of start-ups: to effectively control their burn rate, while ensuring quick entry into the market and managing constantly varying adoption and usage patterns. Therefore, building data centers and managing their operations at the infrastructure layer is generally frowned upon, unless of course their target business is as an infrastructure provider. Start-ups can concentrate on designing their solution, instead of designing, funding, and building scalable infrastructure. Services such as Amazon, Rackspace, Google, Yahoo! and others, where many of the global scaling and security issues are managed, allow this instant scaling capability without the front-loaded costs. Cloud services, especially low- and fixed-cost IaaS and PaaS, provide instant and automated scaling for SaaS and business process-based start-ups and thus lower the barriers to entry.

Many start-ups today would not even exist, and many others would not have seen anything nearing profitability, without cloud services. Zynga, a social gaming company rode to incredible success atop the Facebook platform and Amazon Web Services, reported income of more than $600 million in 2009 according to *Business Insider*.[1]

And it is important that you do consider companies like Facebook as pursuing a PaaS and SaaS business model, as much as Amazon Web Services offers IaaS. This is the approach many new companies take to ensure others build from, an on top of, their solutions, ensuring revenue opportunities for the services that they provide.

For large organizations, an IT department usually evaluates technology decisions on behalf of the business users. Processes will arguably have been developed over many years, creating substantial documentation, cultural knowledge, architecture diagrams (for both facilities and technologies), and so forth. Now consider that your technology requirements have the potential to exist in and be delivered through a combination of private and public cloud services. Critically, large organizations can look to the start-up space as a model for experimenting with new business or service offerings. The fast start and ability to clearly manage costs is compelling, but the security concerns, primarily around data, limit the scope of what is possible in the public cloud model. Therefore, the private and community cloud models, as discussed in Chapter 1, "Introduction to Cloud Computing," are much more accepted. Data concerns revolve primarily around intellectual property (IP) and personally identifiable information (PII) of employees, customers, and so forth.

The New York Times serves as a good example of this. It initiated a project for its TimesMachine (http://timesmachine.nytimes.com/browser) website to improve access to its articles. TimesMachine is a collection of full-page image scans of the newspaper from 1851 to 1922 (i.e., the public domain archives).

In 2007, its first project was to move from a set of individual images of papers from 1851 to 1922 to the popular Adobe PDF format. Whereas it had previously done this conversion on-the-fly, concerns over traffic increases impacting the overall site performance led it to look at performing the conversions in bulk, ahead of time. As reported on The New York Times Open blog, by Derek Gottfrid,[2] he initially copied around 4TB data across to Amazon's Simple Storage Service (S3), along with a system image and some scripts:

> I then began some rough calculations and determined that if I used only four machines, it could take some time to generate all 11 million article PDFs. But thanks to the swell people at Amazon, I got access to a few more machines and churned through all 11 million articles in just under 24 hours using 100 EC2 instances, and generated another 1.5TB of data to store in S3. (In fact, it worked so well that we ran it twice, since after we were done we noticed an error in the PDFs.)

This clear success met both very restrictive cost and time constraints.

For cloud service providers, the goal of optimizing their service delivery will also drive them to consider using other cloud service providers for noncore or nonessential service delivery. The need to scale is key and is based on the ability to achieve what most IT professionals would like to see in their own data centers—optimal use of all infrastructure, not having idle systems, and not hitting a ceiling while delivering.

The opportunity to use cloud services is immense in terms of scope and effort. For small to medium-size enterprises, the task is considered daunting, and for large and global enterprises, even more so. This is why it is unlikely any existing organization will switch immediately to the cloud. Even when considering all the artifacts that exist to describe an organization's orthodoxy, its canons, its ideologies, and its dogmas, a wholesale switch to move those things into a cloud service provider environment to save money or increase speed of delivery is improbable.

So, cloud services adoption will be a set of tactical and strategic moves, staged to meet the business goals and priorities. Given the potential scope of cloud services opportunities, we cannot divine this for you in a book. However, in the following chapters, we will help you determine which order to follow, and on which areas to focus to gain the most benefit in scenarios related to common business models and problems.

For example, consider that cloud services offer the promise of having many operational issues managed by the provider. This does not obviate the requirement for technical architecture, governance, and management, but it does decrease the need for many manual and repetitive task requirements being placed on your staff.

We will, and you must, focus on balancing critical business needs against risk. Both requirements range across cost management, security, governance, compliance, customer satisfaction, delivery, and similar axes. In all cases, the two core data points you need to have are your data classifications and a clear sense of service value.

You may choose to put mission-critical processes into the cloud. Despite it being mission critical, the process or capability might not be a core capability, and as such could be handled by a service provider much more effectively. Consider payroll as a key example. Many companies worldwide use and trust ADP, Intuit, or others for this critical business process. The dominance of FedEx and DHL and others to ship critical business documents is rarely questioned, because it is rarely a core competency of most organizations. This would then direct your approach and evaluation in terms of service guarantees, service-level agreements, disaster recovery, and critically, your data handling requirements.

If you are considering nonmission-critical processes, it is arguably less important to focus on things such as disaster recovery per se, but it remains critical to understand what data is going to be used. This will direct your needs in terms of data protection and privacy management. Even nonmission-critical processes can involve data that needs to be protected.

All of your regulatory, governance, and compliance activities are dictated by understanding the value and use of your data. Data classification is core and critical to evaluating the value of cloud services and any associated risks!

The challenge is that many organizations have not done this, or the classifications are not maintained. Regardless the choice of cloud services you make, it is your responsibility to maintain privacy and confidentially in accordance with regulations. You cannot abdicate this responsibility! Cloud service providers may offer some level of protection, even greater than you could provide. Some may offer remediation in the event of a breach. Ultimately, however, the impact will be on your business, and that is the core issue to deal with.

You can outsource your IT, infrastructure, IT security, and even business processes, but you cannot outsource responsibility or associated risk. (Insurance helps, but ultimately risk lies with your business.)

Throughout the next few chapters, we delve into the next layer of cloud services options and architectures, and then in later chapters, we focus on the business of cloud services.

Endnotes

[1] www.businessinsider.com/chart-of-the-day-monthly-active-users-of-various-widgets-on-facebook-2010-4

[2] http://open.blogs.nytimes.com/2007/11/01/self-service-prorated-super-computing-fun/

PART II

The Cloud Service Alphabet Soup

IN THIS PART

CHAPTER 4

Introduction to Cloud Services

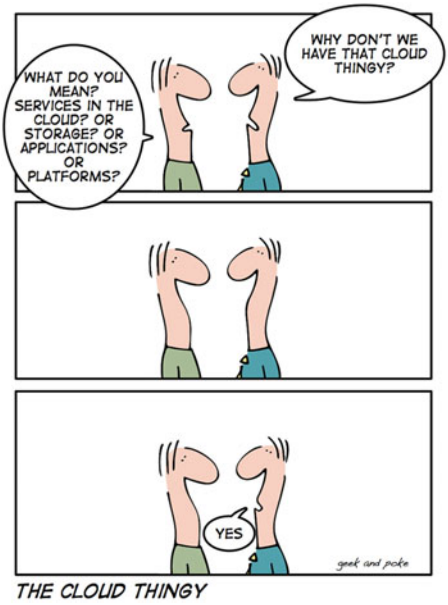
WHY DON'T WE HAVE THAT CLOUD THINGY?
WHAT DO YOU MEAN? SERVICES IN THE CLOUD? OR STORAGE? OR APPLICATIONS? OR PLATFORMS?
YES
geek and poke
THE CLOUD THINGY

Why don't we have a cloud?—Geek and Poke

In this chapter, we explore the lower layers of the common cloud architectures so that you have a clear understanding when we move on to planning for and managing cloud services. Our discussion here examines cloud services taxonomies that incorporate architectural, an understanding of which can enable you to better map from existing enterprise architectures into cloud-based architectures. A definition of usage models is important to understanding your role in the cloud services usage model.

Cloud Usage Models

To move further in the discussion about how to use cloud services requires that we introduce a set of players to the mix:

- **Service provider:** The organization providing the cloud services offers certain capabilities desired by the consumer.
- **Service consumer:** The person or organization using the cloud services with a specific set of requirements from a provider.
- **Service broker:** The person or organization obtaining the services on behalf of the consumer. Although service consumers may interact directly with service providers, in enterprises there is a need to think strategically and create a service broker role. More specifically, this role is associated primarily with the IT department for larger companies, yet external brokers may provide cloud services directly to any part of an organization.

Cloud service brokers usually appear to be cloud service providers themselves by offering consolidated sets of cloud services from one source. This is much like a traditional value-added reseller (VAR) or solution integrator (SI) who bundles third-party products and services together with additional benefits—perhaps single provider billing or support—and then sells them as a complete solution. This results in scenarios where cloud consumers may actually be sourcing cloud services via a network, or chain, of cloud providers. The critical requirement then is to understand where the core services are being provided from and whether that impacts any of your key performance indicators, and critically, your compliance obligations, security requirements, or management metrics. This includes gaining a comprehensive understanding of the geographical and related jurisdictional aspects of all components of any cloud service delivery chain.

Figure 4.1 provides a visual summary and some further detail.

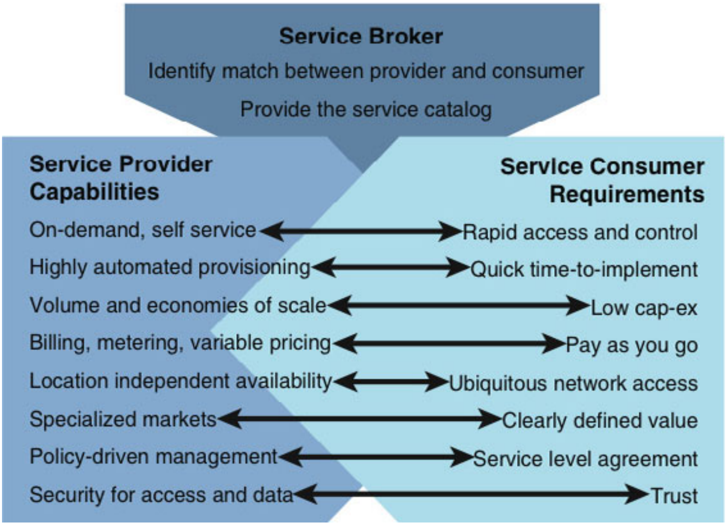
Service Broker
Identify match between provider and consumer
Provide the service catalog
Service Provider Capabilities
Service Consumer Requirements
On-demand, self service
Rapid access and control
Highly automated provisioning
Quick time-to-implement
Volume and economies of scale
Low cap-ex
Billing, metering, variable pricing
Pay as you go
Location independent availability
Ubiquitous network access
Specialized markets
Clearly defined value
Policy-driven management
Service level agreement
Security for access and data
Trust

FIGURE 4.1
Consumers, providers, and brokers of cloud services

The criticality of the service broker role in IT is illustrated through the impact of shadow IT (see Figure 4.2).

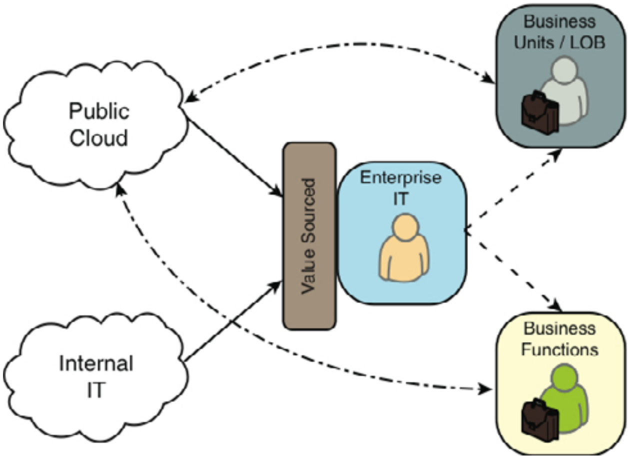

FIGURE 4.2
Shadow IT

In this case, business units or individuals directly procure products and services outside the normal management chain. Cloud service providers with pay-per-use models make it easy to use a credit card to acquire their offerings but incredibly hard for a business to track in terms of costs (and more critical, in terms of risks). This tactic is often a reaction to restrictive IT practices but also indicates that cases exist where IT is not adequately servicing the business needs. By taking the service broker role to a level of procuring business services, IT can gain, maintain, or increase its business relevance by ensuring that all cost management, contractual, legal, and regulatory requirements are being met by each service being procured—cloud or otherwise. This level of control also helps manage costs but will meet business needs only by offering a service catalog with cloud services attributes such as rapid provisioning and self-service.

This now supports a view of how cloud architecture is used by each player.

Cloud Architecture: High Level

In Chapter 2, “Evolution or Revolution?,” we introduced the high-level concepts associated with software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS). This is often referred to as the SPI model (see Figure 4.3).

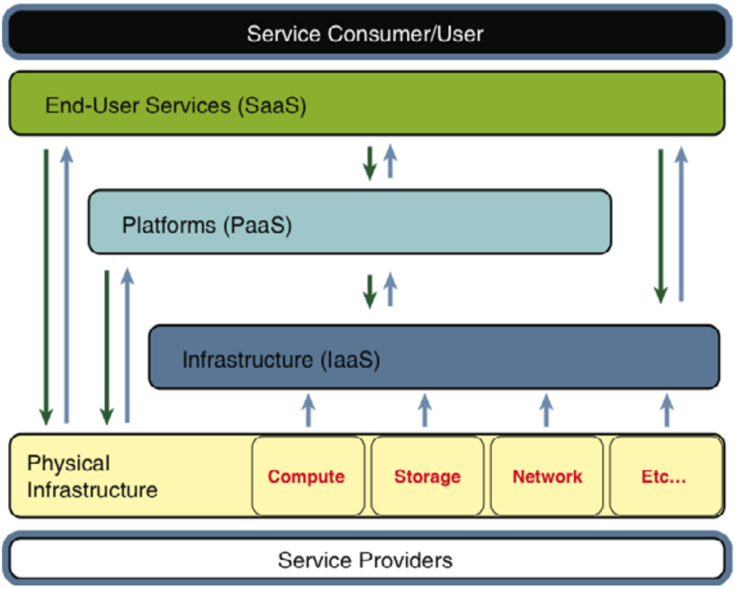

FIGURE 4.3
The SPI model (SaaS, PaaS, IaaS)

This serves to help discuss approaches to various cloud capabilities, but be aware that many folks will create an acronym for almost any cloud service: security as a service, storage as a service, communications as a service. Even though security might be one of the few broad categories for which it might make sense to create something specific, let's be honest and say these terms really do *not* help the discussion! The simple SPI model with appropriate capabilities associated with each, such as compute, network, and storage as capabilities of the infrastructure, is enough.

Let's break down that model by examining the key capabilities of each service:

- **Infrastructure as a service (IaaS):** IaaS refers to computing resources (for example, compute, storage, and network). The base of the stack provides raw capabilities assigned to each customer and minimizes the need to own unnecessary hardware or have costly resources to manage the underlying infrastructure.
- **Platform as a service (PaaS):** PaaS is a set of application and development services (such as Java, Python, or .NET) for building cloud-based applications and services (such as processing applications or SaaS solutions). PaaS solutions may utilize support from IaaS solutions but do not need to do so.
- **Software as a service (SaaS):** SaaS is the category for applications and business processing capabilities that are offered over the Internet or use Internet technologies. These may utilize IaaS or PaaS environments for all or part of their functionality and delivery, but do not require them, with many solutions using highly

proprietary environments for the bulk of their scalability and functional requirements.

To be clear, across all the SPI layers, physical infrastructure may be proprietary or inherited from a lower-layer service, whereas management tools are usually considered SaaS applications in their own right.

The model shows that any user can choose to employ cloud services from any level. For example, you might choose to use a SaaS solution such as Google Apps for email, a PaaS such as Microsoft Azure to run a custom application, while using a storage IaaS provider to support remote backup. None of these depends on the other.

Public Cloud Architectures

Figure 4.4 details a more specific breakdown of all the potential functional or architectural components that go into making up a public cloud solution. The technical aspects of these layers are shown in Figure 4.4.

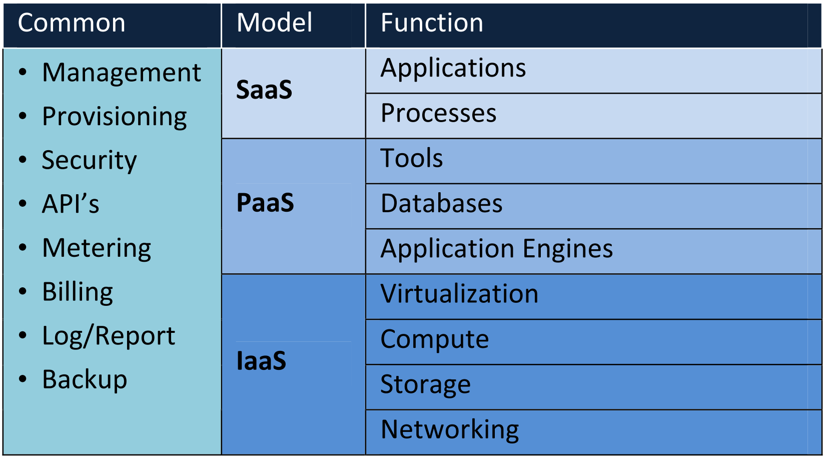

Common	Model	Function
• Management • Provisioning • Security • API's • Metering • Billing • Log/Report • Backup	**SaaS**	Applications
		Processes
	PaaS	Tools
		Databases
		Application Engines
	IaaS	Virtualization
		Compute
		Storage
		Networking

FIGURE 4.4
Common requirements and associated functions for the SPI model

All the components that go into making up each level of a service do not preclude the need for a broad set of general capabilities:

Private and Hybrid Cloud Architectures

Figure 4.5 details a more specific breakdown of all the potential functional or architectural components that go into making up a private cloud solution. To succeed with any cloud solution requires a level of support for each of the unique aspects we discussed in Chapter 2. Architectural elements such as automation, virtualization, and similar technologies and management approaches are important, and they can become a focus for IT departments to the detriment of business expectations. Specific issues that arise often center on user experiences with consumer services. Regardless of the regulation, security, and compliance issues faced by IT, the expectation is that a cloud services model will allow for instant provisioning through a simple service catalog available through a web portal. Traditional approaches to IT requests that involve long-term project planning, budget approvals, and change control need to be adjusted significantly to allow for this type of approach.

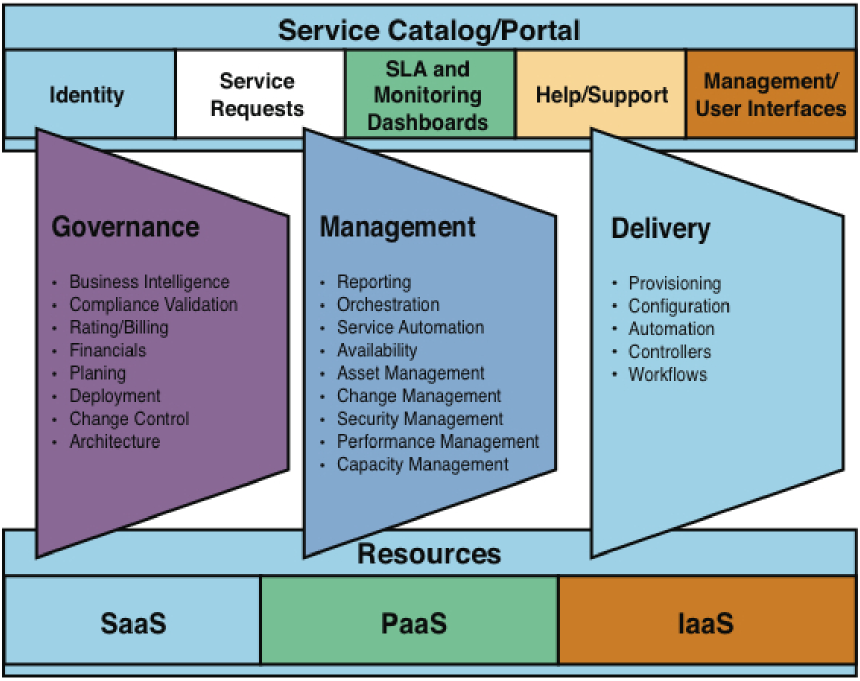

FIGURE 4.5
A full model for cloud services management

The critical difference between public or external clouds and private or internal clouds is the number of options, if not requirements, you have around service management and governance. In both cases, look for a service catalog that allows you to create your service model and request service delivery in the on-demand mode. Similarly, releasing your acquired or requisitioned services should be allowed for.

The line from traditional data centers through optimized IT in a private cloud model and ending in public cloud options introduces several pros and cons for each. As shown in Figure 4.6, the move to a private cloud minimally results in a service-oriented architecture.

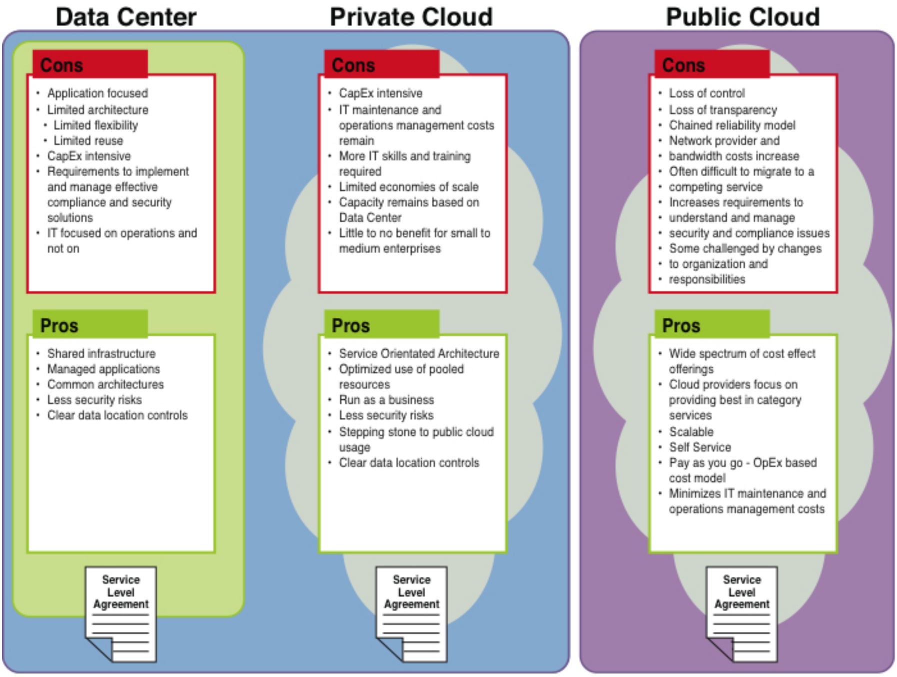
Data Center
Private Cloud
Public Cloud
Cons
Application focused
Limited architecture
Limited flexibility
Limited reuse
CapEx intensive
Requirements to implement and manage effective compliance and security solutions
IT focused on operations and not on
Pros
Shared infrastructure
Managed applications
Common architectures
Less security risks
Clear data location controls
Service Level Agreement
Cons
CapEx intensive
IT maintenance and operations management costs remain
More IT skills and training required
Limited economies of scale
Capacity remains based on Data Center
Little to no benefit for small to medium enterprises
Pros
Service Orientated Architecture
Optimized use of pooled resources
Run as a business
Less security risks
Stepping stone to public cloud usage
Clear data location controls
Service Level Agreement
Cons
Loss of control
Loss of transparency
Chained reliability model
Network provider and
bandwidth costs increase
Often difficult to migrate to a
competing service
Increases requirements to
understand and manage
security and compliance issues
Some challenged by changes
to organization and
responsibilities
Pros
Wide spectrum of cost effect offerings
Cloud providers focus on providing best in category services
Scalable
Self Service
Pay as you go - OpEx based cost model
Minimizes IT maintenance and operations management costs
Service Level Agreement

FIGURE 4.6

The pros and cons of traditional shared data centers, private clouds, and public clouds

A final note to this discussion focuses on service level agreements (SLAs). Each model of delivering services to your organization will introduce variation into how you monitor its success. Approaches to cloud computing often deal primarily with uptime in their SLAs, and their measure of availability may differ from your expectations; so ensure that you cater for the variations, as well as clearly evaluate the taxonomies and terminology in your agreements.

Cloud Architecture in Depth

In this section, we examine each service category in more detail and provide critical questions to be answered when considering using such cloud service providers.

Software as a Service

Although IaaS does get a significant amount of press as a result of the success Amazon has shown with AWS, SaaS represents a significant majority of cloud solutions. SaaS is very popular with small and medium-size organizations because it minimizes the amount of software and hardware to manage; and by being accessible through a browser or other common user interfaces (e.g., email clients), it allows for ubiquitous access at all times. In addition for such organizations, there are quantifiable advantages in terms of scalability, security, and reliability, allowing them to access capabilities that would traditionally require large enterprise funding and management capabilities to deliver. However, these benefits may not always create such significant differentials for larger organizations and call for clear cost and risk analysis.

SaaS in traditional architecture models is the business processes and applications, such as document editing and management, email, customer relationship management (CRM), sales force automation (SFA), human resources management (HRM), and so forth. Service providers of SaaS solutions deliver applications through web pages, rich Internet application (RIA) environments such as Adobe AIR and Microsoft Silverlight, or finally, support traditional desktop or server applications with software for integrated services such a Google Maps.

SaaS vendors are generally evolutions from application service providers (ASPs). The primary difference is in the licensing. ASPs were commonly aligned with traditional outsourcing-type models wherein customers would buy the license for the software being run, whereas the core of the SaaS approach is to support multitenancy, with licensing being on a per-customer or per-user model.

Offerings from SaaS providers cover most enterprise and productivity applications, including Adobe's eSignatures and Photoshop.com, Microsoft Exchange Online through Microsoft's Business Productivity Online Standard (BPOS) Suite, Google Apps, Oracle's Customer Relationship Management (CRM) Online, and so forth. Alongside these traditional applications, new and different communication and data management solutions that were not possible using traditional software solutions have been created, such as Facebook, Twitter, Ning, and LinkedIn.

Each type of organization will have different concerns that need to be considered with SaaS:

- **Large enterprises:** Integration with security management, identity management, backup and recovery, monitoring and reporting, forensics, and business failover
- **Small to medium-size enterprises:** True improvement in security, remote- and time-based support requirements, backup, recovery, and cost management

Platform as a Service

Platform as a service (PaaS) offers a managed application platform for building and operating applications and services, even SaaS.

The requirements of a comprehensive PaaS are capabilities to support developers in terms of suitable languages and APIs alongside controls for presentation logic, data management, and infrastructure access where needed. PaaS solutions eliminate the need for an organization to deploy costly infrastructure to support the related deployment of applications. The assumption is that cloud qualities of a PaaS would incorporate session management, transaction integrity, reliability, availability, and scalability. Most benefits attributed to PaaS solutions, including security, reliability, and elasticity, scale to support business capacity management.

Today, most PaaS solutions offer some level of migration support from existing architectures, incorporating developer tools. Some may also incorporate specific capabilities in terms of data services, including Structured Query Language (SQL), NoSQL, or other data query and management solutions. This can be important when considering how much direct access can affect processing speed.

Key areas to review include the following:

- Support for your preferred development environment:
    - *Java application platform vendors:* Google App Engine, Stax, and Morph Labs, for example.
    - *Microsoft:* Microsoft offers a set of software infrastructure services such as .NET Services and SQL Services running in an elastic operating environment called Windows Azure. Microsoft offers the Azure platform to third parties.
    - *Web application platforms and frameworks:* PHP, Ruby on Rails, Drupal, and similar.
    - *Proprietary platforms and frameworks:* Salesforce.com, LongJump, and Google App Engine.
    - *Proprietary niche vendors:* There are a number of highly specialized platforms, such as Facebook, WordPress, and Ning.
- Using PaaS does not negate the need for a solid and secure software development life cycle (SDLC).
- PaaS support for versioning, backup, recovery, and QA/test/release with live rollout/rollback.
- Does the PaaS provider offer any option beyond a contained or virtualized environment? Is the language support proprietary?

As a word of caution, consider customers who chose to jump into the cloud with a PaaS provider called Coghead. Founded in 2003 and formally launched in 2006, Coghead offered visual application building services in the cloud. Coghead even provided sample starter applications built with its system, such as CRM and issues tracking. Coghead competed with solutions from LongJump to Salesforce.com's Force.com PaaS. Unfortunately, despite a reported developer community of more than 25,000 and a shift to Amazon EC2 as its hosting environment in 2008 to control costs,[1] Coghead announced in early 2009 it was going out of business. After three years, users were given less than six weeks' notice to leave the service. Although many competitors immediately offered migration tools and services, it is clear that having the service actually close down as opposed to just a software environment not being supported any more creates a definitive end for all users involved. Very quickly, SAP announced it was buying Coghead, but SAP made it clear it was buying only the intellectual property, not the users, or any responsibilities to continue servicing the Coghead user community. A key difference between using a public or third-party-provided cloud service, as opposed to using an in-house solution, is that if a public cloud service goes out of business, there is an immediate shutdown of the service. Owning the platform and software associated with the service internally gives you much more time to deal with these sorts of events.

Infrastructure as a Service

Infrastructure as a service (IaaS) is a virtual or physical resource offered as a service. Although commonly associated with virtualized offerings such as Amazon's EC2, a vast array of infrastructure offerings incorporate direct physical access to hardware, more in tune with traditional managed service providers. Rackspace makes an effort to offer this kind of option to customers for example.

What IaaS providers strive to offer is highly efficient use of hardware while creating improved scale and availability for their customers. To achieve this, most cloud service providers try to virtualize all aspects of access to specific resources including compute, network, and storage. From this, it should be clear that virtualization is not specifically associated with the core compute capability. The IaaS approach to virtualization also gives us a new way to view deployments at scale. We are moving from chips with units of resistors and links to Moore's law in terms of speed increasing over time, to now considering compute and related capacity at warehouse levels, with complete units of compute capacity (memory, chip, storage, connectivity). This shift changes how we manage our deployments as well as drives a change in how we consider fault recovery at the hardware level.

Key areas to review include the following:

- Support for virtual environments and how will you best manage them:
    - VMware (vCloud)
    - Citrix Systems Xen-based solutions and Citrix Cloud Center (C3)
    - Microsoft HyperV
    - Red Hat's KVM
    - And more, proprietary implementations, or solutions derived from one of those listed above
- Are there potential performance issues that require direct access to hardware?
- Are there potential performance issues introduced between an application and the data store due to the impact of network latency?
- Can you migrate cleanly from one provider data center to another?
- Can you migrate cleanly from one provider to another?
- IaaS support for versioning, backup and recovery, and QA/test/release with live rollout/rollback.

For example, Amazon (EC2) uses a proprietary packaging of Xen's virtual environment called Amazon Machine Images and Rackspace. AT&T, T-Mobile, and CDW are examples of hosting providers planning to provide IaaS using VMware's vCloud product.

Cloud Architecture Conclusions

As cloud services emerged, the most common usage of IaaS was for development and testing of solutions. The ability to layer virtualized images of systems atop a common set of infrastructure, alongside the potential to create large-scale client environments (for example, desktops) to facilitate QA and load testing, was quite compelling. Having this proof point, along with improvements in reliability, led to many organizations defining IaaS as the preferred approach for deploying any new applications, using private cloud options in general. Today, however, the decision increasingly includes the option to move to public cloud services and even up the stack to PaaS and SaaS. SaaS options see increasing focus as both a migration option from in-place, or legacy systems, and as the first choice options for most new deployments, especially infrastructure and enterprise applications. The question then is: Can any application, platform, or service be cloud based?

The simple answer is yes—almost any application or service can be found or made available through cloud services. The caveat is that just because you can doesn't mean you should. Practical considerations such as security, performance, availability, and more are critical factors in the decision (as identified previously in Figure 2.3).

In conclusion, the models as discussed show there is no mandate that the upper-level services make use of any lower-level services. In fact, there is only optional support of upper-level services by those below. The reason for this is important to understand because it illustrates key considerations you need to make when considering your use of both public and private cloud environments. For example, consider the SaaS and PaaS offerings from Salesforce.com. In March 2009, Salesforce.com released information about their customer base and infrastructure at that time showing that they supported around 55,000 customers on about 1,000 servers, with near only 50 servers dedicated to actual database storage. Salesforce.com achieved this by running a proprietary codebase, proprietary database, and proprietary "multitenant optimizer" that efficiently slices and dices the data.

Consider that Facebook[2] uses a similar approach. Facebook makes good use of open source software: Tornado, Cassandra, Thrift, XHProf, Hive, MySQL, Memcached, and more. Like Salesforce.com, most everything else is customized and highly optimized for the services Facebook provides.

In contrast, SmugMug[3] offers a SaaS solution for managing and editing pictures, and its infrastructure uses the IaaS and PaaS capabilities of Amazon Web Services (AWS). In 2006, when SmugMug decided to use AWS, it was an independent, self-funded, profitable, and debt-free company with 1 programmer and only 15 employees. They offer a subscription-based online photo-sharing company, have more than150,000 paying customers, and safely store more than 70 million photos on their behalf. It took SmugMug just one week from writing the first line of code to being fully operational on Amazon S3. The company copied more than 80 terabytes of existing customer photos to Amazon S3, representing more than 70,000,000 original images and 6 display copies of each.

SmugMug also backs up new photos to Amazon S3, resulting in an additional 10 terabytes of data added to Amazon S3 each month.

Before making the move to Amazon S3, SmugMug was depending on its own redundant storage and multiple datacenters, which soaked up a lot of capital. The company saved roughly $500,000 in planned disk drive expenditures in 2006 and cut its disk storage array costs in half.

Although cloud solutions may use other cloud provider offerings, it is not required, nor in some cases is it efficient to do so.

Let's also be clear that each of these lower-level architectures may or may not map easily to your historical architectural or management approaches. Private cloud implementations at large organizations often start as optimized data centers based on Information Technology Infrastructure Library (ITIL) types of environments. This illustrates a potential point of contention when enterprises consider using public or hybrid cloud solutions: Quite simply, the mapping of these lower-level architectural requirements is not always 1:1. Some providers say that ITIL is too heavyweight to allow for all the benefits of cloud services to be offered. However, the clear argument is that using an ITIL-like, or ITIL-light approach, allows for the consideration and inclusion of the essential management essentials for good IT usage. Enterprises will find that the years, perhaps decades, spent driving toward adoption and acceptance of their management will be challenged when using cloud services, much like the X.400 show down against SMTP, or how TCP/IP challenged DecNET, or for the business model folks, the VHS versus Betamax war. Later in this book, we discuss the standards swirling around cloud services.

Choices to use any or all of these solution types will be defined by specific requirements and data classification.

Key points to consider from this chapter are

- Cloud services involve multiple parties: a cloud service consumer, a cloud service provider, and sometimes a cloud service broker.
- Each party to a cloud service has specific roles and responsibilities.
- Each layer of the SPI model may use a lower layer as part of its architecture, but it is not absolutely required nor is it always optimal.
- Almost any application can be a cloud service, but that does not mean it is the right thing to do.
- Cloud service providers may go out of business and solutions can disappear immediately with them, or cloud service providers may modify or deprecate their offerings. This creates a very different experience to purchasing software licenses−more flexibility to continue using the software as is, regardless of the broader scenario.
- Look for cloud service providers that deliver and secure the service better than you can.

- Do not allow cloud services to preclude good IT practices.
- Simple approaches usually win over more complex ones.

Endnotes

[1] http://techcrunch.com/2008/01/14/coghead-20-built-on-adobe-flex-hosted-by-amazon/.

[2] http://video-jsoe.ucsd.edu/asx/JeffRothschildFacebook.asx.

[3] http://aws.amazon.com/solutions/case-studies/smugmug/.

CHAPTER 5

State of the Industry

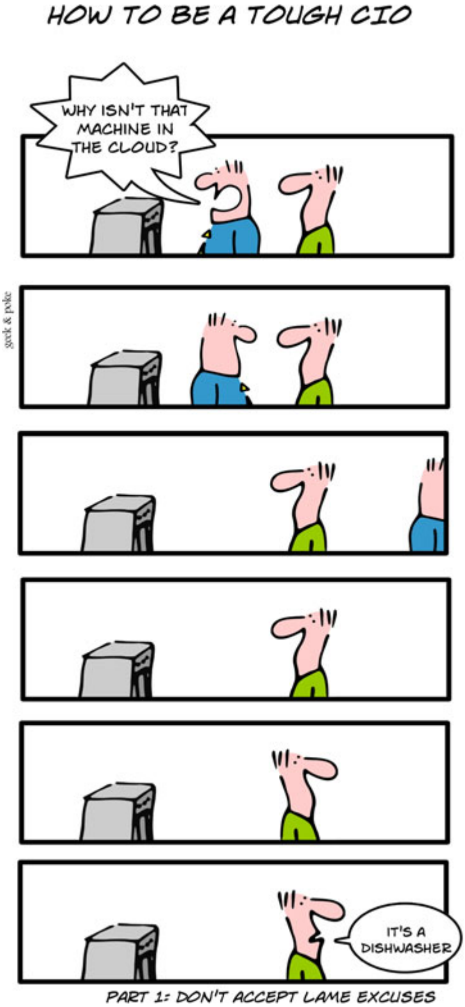
HOW TO BE A TOUGH CIO
WHY ISN'T THAT MACHINE IN THE CLOUD?
geek & poke
IT'S A DISHWASHER
PART 1: DON'T ACCEPT LAME EXCUSES

A tough CIO asks the tough questions—Geek and Poke

Reviewing much of the press and customer commentary on cloud computing creates a sense that to take advantage of these marvelous capabilities, we are looking at something so new that we need to rethink everything. What is important to do, however, is focus on the areas and disciplines that already exist in earnest, and do them right.

Many industries have regulations and recommendations on technology usage, yet quite often these and the legacy already available to us are sidelined as new opportunities are adopted. The "move to the cloud" isn't about capital expenditure (CapEx) versus operational expenditure (OpEx) or any of the other benefits. Cloud services open up whole new fields of innovation. *That* is the "cloud" world.

The changing landscape could become a Wild West, yet large players are investing heavily in standards, as well as approaches to defining levels of assurance in cloud services. Start-ups and smaller players are playing both sides of the fence. Some are racing ahead to deliver new solutions despite the lack of standards in place, whereas others are using standards as a marketing tool to try and ensure they are at the front of the line when the standards are ratified.

This chapter takes a look at all these aspects, as they are driving, and being driven by, the cloud service industry.

Cloud Services Create New Markets and New Threats

Throughout the narrative so far, we have detailed multiple examples of businesses that have gone through changes as a result of new market opportunities.

Although commonly cited cloud service providers such as Amazon, Google, and Microsoft are making headway with their own offerings, many other organizations are today viewing cloud services as an essential way to either gain or, at a minimum, maintain customers.

New Markets

For many years, companies have looked for hosting providers to support their web presence. From colocation facilitates through to more comprehensive application service providers, cloud is a natural evolution for vendors who previously provided these backbone capabilities. Now, these vendors are remaking themselves as cloud service providers by more closely aligning with cloud delivery models. Rackspace, Saavis, Hosting.com, and more worldwide are seeing their customers asking for greater flexibility over their usage patterns, billing models, and capacity management (i.e., the essential characteristics of cloud computing), considering that their managed services allow them to deliver against customer concerns with cloud services while maintaining a strategic link with customers. Vendors worldwide (including British Telecom, Telstra, AT&T, Verizon, and more) are a critical part of the chain for cloud services. These communications vendors are looking at the greater market for managed services and view cloud services as a way to create incremental revenue from their strategic assets of wired and wireless networks. Because of their direct link with customers through their control over the networks, they believe that they can mitigate common concerns around cloud services, offering greater levels of reliability, availability, integrity, and performance to customers. Many are offering infrastructure as a service (IaaS) with value-added platform as a service (PaaS) or software as a service (SaaS) solutions, attempting to head off any disintermediation, or customer migration to other vendors who do. Furthermore, most communications providers already operate under monthly billing models, with value-add services available for additional costs on an as-needed basis. Further, they operate at a scale that allows them to truly gain from the shared savings possible in multitenant architectures. However, the full costs need to be calculated to create a valid business case.

Hardware vendors offering infrastructure, such as servers, storage, and networking, are seeing a negative impact on sales. Although the global financial crisis has created some lag, the cloud delivery models have created another negative impact. As a result, most hardware vendors are aiming to provide hardware to several different types of customers: enterprise customers looking to achieve their own cloud implementations, for instance, and cloud service providers looking for hardware at optimal price points.

Traditional IT management solutions and consulting vendors such as Deloitte, Wipro, IBM, HP, BMC, and CA are all moving their consulting services and management tools to support cloud services architectures but are deliberate in doing so for a number of reasons.

For the management software vendors, the lack of standards means that adding support for cloud services that use proprietary APIs is a costly investment that may not deliver a return. As a result, start-ups that can afford to target specific vendors with a growing user base are filling the gap. Consider a couple of examples.

Aiming to address concerns with lockin alongside manageability, RightScale aims to support customers wanting to use multiple IaaS providers at scale. The RightScale SaaS solution allows for the complete life cycle management of virtual images and the associated monitoring of them. Deployment is possible to in-house or private cloud environments, Amazon’s EC2, Rackspace, Cloud.com, GoGrid, and more. In terms of support, RightScale can manage multiple virtualization solutions and solution stacks. This means they are more able to support varying customer requirements.

In contrast, CloudSwitch aims to provide a more secure and consistent solution via its solution model. CloudSwitch offers an appliance that sits within your environment and manages all security and migration between in-house or private cloud implementations through to Amazon’s EC2. In terms of support for virtualization, CloudSwitch is limited to a VMware Hypervisor. This is not only due to its later entry into the market, but also its need to create a secure environment in which to move systems, which it dubs Cloud Isolation Technology.

Quite simply, a significant amount of effort is going on in and around the management software vendors, and choice is generally limited by a simple question of whether the vendor supports your deployment model of choice. However, expect things to change rapidly.

As the market grows, start-ups will succeed, fail, or be acquired, and you may experience potential conflicts depending on the acquirer. This is a risk for any choice you make, but the changes occurring in the cloud market mean that planning requires a moderately different approach to solution usage. In this, your ability to operate in an agile mode will allow you to gain most success.

A similar issue exists for consulting, wherein established and trusted cloud service providers upon which consulting teams are willing to bet customer solutions are limited.

New Threats

As noted in Chapter 1, "Introduction to Cloud Computing," the threat models are also moving quickly as a result of cloud services, and Chapter 9, "Cloud Business Risk and Security," deals with how to manage security risks in cloud services. In terms of the market, however, security services in and for the cloud are evolving rapidly with the landscape, too.

An approach being rapidly adopted by organizations of all sizes is managed security services (MSS). Although sometimes used to refer to in-house solutions, MSS is more commonly used to describe an approach to security operations wherein a service provider (managed security service provider or MSSP) is tasked with operating an organization's network or information system security, or even complete security operations center (SOC), usually from an offsite location.

MSSP capabilities vary by provider, but may include the following:

- Intrusion detection and prevention systems (IDPS) and firewalls
- Email monitoring and malicious email and spam prevention
- Data-loss prevention
- Change management and patch management
- License management
- Ongoing log management and reporting
- Incident and emergency response

Most MSS providers are now casting their solutions as SaaS solutions, but they rarely line up with the essential characteristics of cloud services in terms of being on-demand in relation to the billing model. Despite that, their value is usually easy to assess financially relative to any existing in-house solutions. When organizations want to minimize the effort associated with 24x7 monitoring and initial response to security incidents, MSS are a popular choice. Symantec and McAfee offer broad and compelling solutions, alongside many smaller vendors focused on specific capabilities previously listed. In this environment, the rapid evolution of the threat model in relation to cloud services means opportunities for many other vendors to provide cloud-based security services. In addition, cloud service providers are increasingly including security solutions into their offerings, decreasing the need for customers to do it themselves. Google bought Postini in 2007 and quickly offered the email security services as part of Google Apps.

For any vendor, the critical test comes when a breach occurs. Although protection is a critically important aspect of the service, the ability to provide resources, including alerts, support, logs, and audit data is also crucial in supporting a breach investigation or forensics request. In fact, the need for similar requirements should be evaluated against all external providers, at least in relation to your internal security baselines.

Closely related to security is backup and recovery from data through to data centers. IT departments spend a considerable amount of time designing architectures that are fault-tolerant. Failover options ranging from one host to another, through to the implementation of complete backup data centers in some level of standby mode are common. Cloud services are creating new ways to deal with failover requirements. Cloud service providers are creating new and interesting ways to mitigate risks of data loss due to failure. Cloud service provider architectures are illustrating new ways to create resilient data centers that expect failure as the norm, rather than try to avoid it.

This is an area where multiple vendors have stepped in with a variety of data recovery and replication solutions that would not have been seen as contenders without using a cloud service model. Symantec, Mozy, Carbonite, and others offer near-instant and constant replication of managed data up to a cloud backup service. Usually targeted at individuals or small- to medium-sized businesses, these solutions use agents on devices to immediately replicate data into their secured storage in the cloud without user involvement. Most backup solutions use a variation on the standard approach to backup: In the event of a failure, you need to acquire a new machine on which to operate, installing whatever operating system and applications are necessary, and then use the service to recover the data.

New approaches can be seen in solutions from SugarSync, DropBox, and others that enable you to not only back up the specific data into the cloud, but also to nearly immediately synchronize that data across multiple machines and devices. In addition, these services allow you, with some varying degree of flexibility, to share synchronized directories/files with others. While these services can support a backup and recovery model, the synchronization services create a compelling differentiator. Many users today have more than one computer or device (e.g., a work and home computer, a desktop and laptop, a laptop and a smartphone). With this in mind, if one machine were to fail, a user could immediately move to another machine that is synchronized and immediately keep working using a reasonably recent and consistent version of the synchronized data. This is a compelling model for many users today.

Finally, consider the approach of a company like Backupify, which aims to preserve an individual's online life or critical business data online by backing up data from Facebook, Google Mail, Google Picasa, Flickr, Twitter, Zoho, and Google Apps (among other online services). When moving to the cloud, individual providers often refer to their own strategies for reliability and related backup policies. From a cohesive strategy to protect personal or business data, however, an approach exemplified by Backupify might be absolutely critical to preserve business continuity.

The goal here is to illustrate again how business models are changing, and how specific tasks or workflows inside your organization can be changed to optimize their effectiveness, cost, and reliability by using new approaches.

Standards

The world of cloud standards is changing rapidly, but not at the same speed with which they are needed, or at the pace that cloud services are evolving. As a result, you should expect the simpler standards and approaches to have most traction.

The importance of standards for your initiative should be made clear: Without standards, your ability to implement, migrate, and integrate your service from and across one provider to another will be incredibly difficult. This is the constraint of proprietary interfaces.

We've already touched on the Cloud Security Alliance as a critical component of securing cloud services. To expand our taxonomy and actually create commonality, the DMTF, formerly known as the Distributed Management Task Force, introduced documents to standardize terms used in cloud computing alongside a proposed set of standard and public cloud service APIs. The DMTF documents are

- "Architecture for Managing Clouds"[1]
- "Uses Cases and Interactions for Managing Clouds"[2]

In these documents, the DMTF lays out essential functions for cloud computing and the language that can be used to describe them. As the APIs are developed, however, it is clear how important these efforts are. Success was found by the DMTF in driving the Open Virtualization Format (OVF) for standardizing how to package and distribute virtual machine images. However, in the broader cloud services market, many vendors leapt into cloud services with proprietary solutions. Since then, some have offered their "standards" to the DMTF in order to support the API effort. VMware submitted the VMware Cloud API, and Fujitsu, HP, Telefonica, and Oracle all submitted various other cloud service APIs. This illuminates the absence of vendors like Amazon, who at the time relied on market dominance to maintain a de facto standard with their own APIs. Cloud consumer participation therefore becomes paramount if cloud consumers want the ability to more easily consume and migrate between cloud services.

The TM Forum offers the simple relationship between buyers/consumers and the actual standards, as shown in Figure 5.1

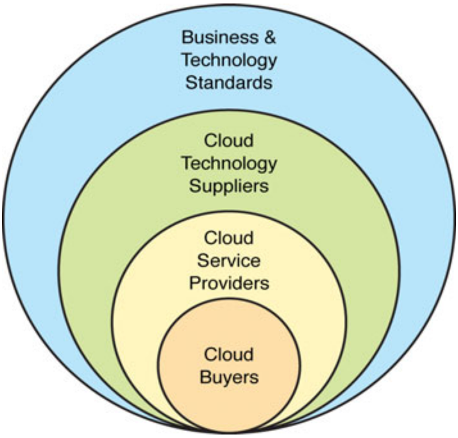

FIGURE 5.1
The TM Forum model of the cloud ecosystem
Source: http://www.tmforum.org/sdata/content/cloud/ecosystem.jpg.

Despite a firm desire to have standards, buyers (users) of cloud services in particular are generally in no rush to join standards groups for their own benefit. It takes time and money that is often seen as better spent elsewhere in business development. The expectation is that the right mix of proprietary and open standards will evolve as needed. Despite that, a common refrain from customers is that standards are not developing fast enough, or adhered to enough. The counter is obviously that their association with the standards bodies and development of standards is limited often by this same set of concentric rings.

Therefore, we advocate customers getting involved to see how the sausage is made when it comes to standards. This also serves as an indication that the standards are needed and supported.

When it comes to cloud standards, however, it's somewhat the Wild West. The best approach you can take here is to use standards as much as possible to consolidate and virtualize your environment so that you can automate all the mechanical aspects of your environment, regardless of whether they are internal or external. Figure 5.2 shows that proprietary approaches do provide benefits in the short term by allowing for business to develop and grow, but contrasts that with the longer-term need for standards to support a much broader ecosystem that allows for easy migration and choice.

This is an excellent time to make a point on virtualization. Although seen as a stepping stone to cloud computing, virtualization is generally most successful in supporting the drive to consolidation and helping to drive standardization. This is by the far the most benefit organizations will achieve through the use of virtualization.

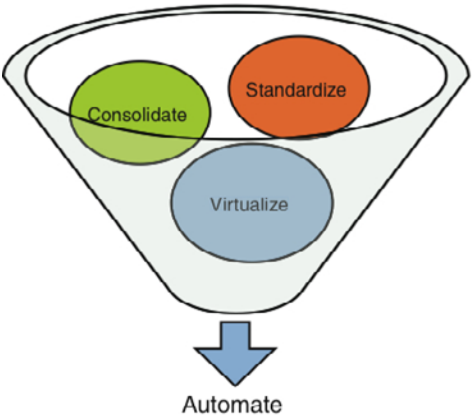

FIGURE 5.2
The simple approach to critical cloud management

For standardization, the market today is significantly fragmented but working toward common goals. Because cloud services as defined today are available in a vast array of environments and models, a single standard will never encompass them all in any satisfactory form. In fact, as shown in the sampling of organizations in Figure 5.3, there are multitudes specifically staking claim to different parts of standards requirements as they relate to cloud services.

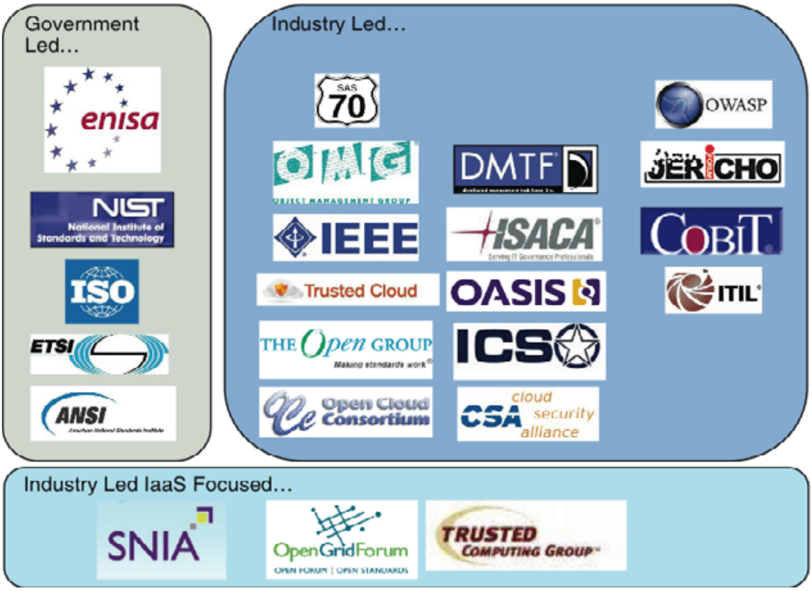
Government Led...
enisa
NIST National Institute of Standards and Technology
ISO
ETSI
ANSI
Industry Led...
SAS 70
OWASP
OMG Object Management Group
DMTF
JERICHO
IEEE
ISACA
COBIT
Trusted Cloud
OASIS
ITIL
The Open Group Making standards work
ICS
Open Cloud Consortium
CSA cloud security alliance
Industry Led IaaS Focused...
SNIA
OpenGridForum Open Forum | Open Standards
Trusted Computing Group

FIGURE 5.3
A sample of worldwide standards bodies working on cloud standards

The International Organization for Standards (ISO), the European Network and Information Security Agency (ENISA), the TM Forum, The Open Group, the Institute of Electrical and Electronics Engineers (IEEE), the Cloud Security Alliance (CSA), the Storage Networking Industry Association (SNIA), the Open Cloud Consortium, and many more—all are staking a claim in parts of the cloud space.

Each organization has something unique to offer, and despite sometimes vast overlap, all have a role to play. Although the standards might not be evolving at what could be termed “cloud speed,” it is essential to identify critical areas where standardization will benefit you. From that, you can better determine whether you are willing to accept de jure standards such as Amazon’s Elastic Block Store (EBS) API versus attempts to create de facto standards such as the SNIA’s Cloud Data Management Interface (CDMI).

Figure 5.4 shows the primary points of contention between proprietary and standardized APIs, and details the key concepts impacting cloud standards In one sense, proprietary APIs often see faster development over standards-based APIs; however, they create more vendor lock in unless widely adopted by the industry.

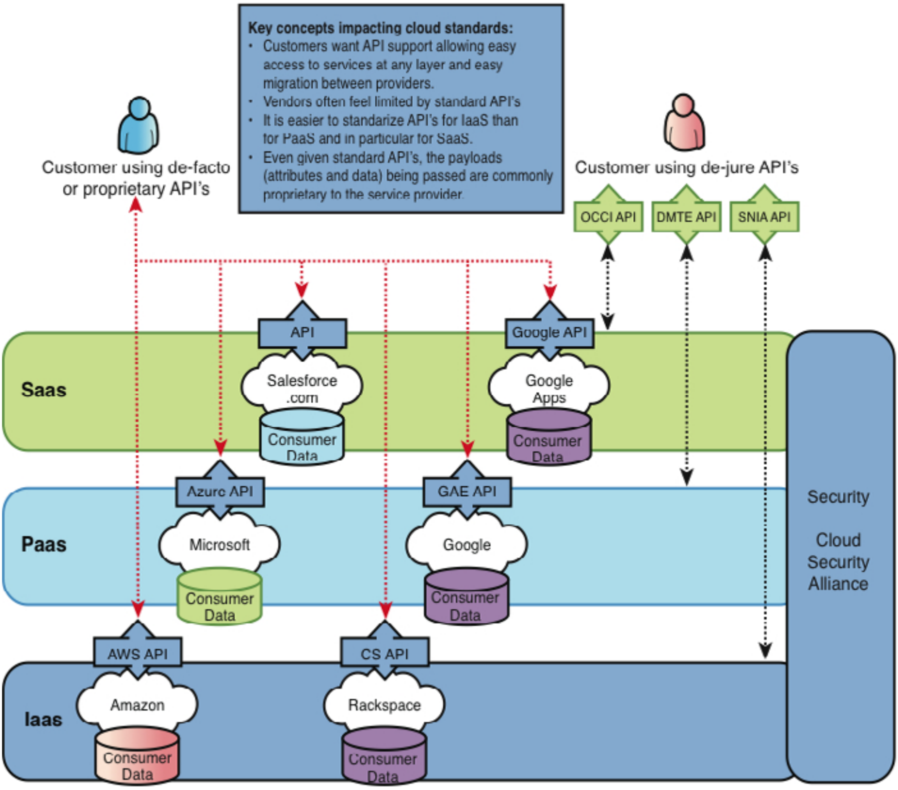

Key concepts impacting cloud standards:

- Customers want API support allowing easy access to services at any layer and easy migration between providers.
- Vendors often feel limited by standard API's
- It is easier to standarize API's for IaaS than for PaaS and in particular for SaaS.
- Even given standard API's, the payloads (attributes and data) being passed are commonly proprietary to the service provider.

FIGURE 5.4
De jure standards versus de facto standards

It is arguably hard to see the benefits if you just take the Amazon use case, but what this truly illustrates is the importance of the standards approach as more cloud services are consumed. The question then is what level of faith or trust (or in business terms, assurance) do you have in cloud service providers?

Assurance in the Cloud

Cloud service providers and cloud users look beyond the standards per se to gain some level of assurance that their governance requirements and risk concerns can be managed effectively. Cloud service providers should be pushed to demonstrate adequate controls and safeguards when they host or process data belonging to their customers.

Many providers are using the American Institute of Certified Public Accountants (AICPA) Statement on Auditing Standards (SAS) No. 70 in an effort to provide a level of assurance to customers that an audit has been conducted into the provider's control objectives and control activities, including controls over information technology and related processes. The focus of cloud service providers is generally to audit security-related controls. The benefits and risks of relying on SAS 70 must be understood, however.

Although used by thousands of public companies and businesses that provide critical financial capabilities to their own and other organizations, SAS 70 is an accounting standard developed more than 20 years ago. The intent was to create an auditor-to-auditor communication to reduce audit burden by customer auditors of shared services.

SAS 70 generates two types of (restricted) reports:

- **Type 1:** Reports on controls placed in operation
- **Type 2:** Reports on controls placed in operation and tests of operating effectiveness

Unlike prescriptive regulations such as the Payment Card Industry (PCI) Data Security Standards (DSS), SAS 70 is unambiguously flexible and therefore not a clear security certification or seal of approval.

So, SAS 70 Type 2 audits provide evidence that a provider has thought about and documented its controls and had a third-party audit performed to validate the effort. However, as said, this is not a security certification, but it does ensure that the provider is using controls based on the ISO 27001/27002 Information Security Management System (ISMS) standards published by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC). The focus therefore should be on other security standards and certifications that provide a better validation of ISO/IEC 27001 compliance because ISO/IEC 27001 requires that management:

- Systematically examine the organization's information security risks, taking account of the threats, vulnerabilities, and impacts.

- Design and implement a coherent and comprehensive suite of information security controls and/or other forms of risk treatment (such as risk avoidance or risk transfer) to address those risks that are deemed unacceptable.
- Adopt an overarching management process to ensure that the information security controls continue to meet the organization's information security needs on an ongoing basis.[3]

What this always incomplete model serves to remind us of, however, is that the potential landscape of cloud solutions is truly vast. The risk associated with using cloud services is increased by the lack of a standard taxonomy to describe how services will operate and how they are undertaking security-risk assessments.

So, although all these things evolve, consider the following:

- Require by preference ISO 27001 certification.
- Build your controls and assurance model using known models that together, support a comprehensive security evaluation:
    - Adherence to industry, national and regional data protection policies
    - PCI when using credit card data)
    - CSA controls matrix for comprehensive security
    - HiTRUST for healthcare security in relation to USA's HIPAA/HITECH.
    - Specific levels of Federal Information Security Management Act (FISMA) certification is required for US government initiatives.
    - ENISA cloud computing risk assessment, primarily but not exclusively for European cloud providers and consumers.
- Only accept a SAS 70 Type 2 based on a known security controls model, at a minimum.

These approaches to management and security are the baseline standards you should accept when evaluating a cloud service provider. Each option offers a method to assess the level of control and risk assessment a vendor has chosen to use, and outside of the SAS 70 audit, all provide a common framework and taxonomy to equally rate solutions against. Simply, despite the current popularity of SAS 70, most approaches require audits to be certified, and SAS 70 does not represent them at all.

We discuss this in more detail in Chapter 9.

Summary

In this chapter, the goal was to discuss the market landscape and reiterate with specific examples the rapid changes occurring to business models. The market is evolving rapidly, and that raises some cause for concern as you try to determine which solutions will allow you to capitalize on the opportunities we have detailed throughout the book so far. Critical aspects of how you will manage data and business logic to migrate across chosen services, while maintaining adequate levels of security, reliability, availability, and performance, were discussed. The choice of tools to manage these critical areas is clearly nascent, even from large-scale vendors, but they are adapting quickly.

Key points to consider from this chapter are

- Cloud services create new markets and new threats that need clear identification and management.
- Despite cloud services evolving from many technologies, standards related to cloud services are still nascent.
- Outside of industry and general regulatory requirements, certification of cloud services requires more than SAS 70 and should be based on ISO27k standards at a minimum.

Ultimately, evaluating cloud service providers will rely on how dedicated your organization is to using adequate risk analysis, and the final goal of this chapter was to introduce some methods to help determine the risk and assurance levels of each solution.

Endnotes

[1] http://www.dmtf.org/standards/published_documents/DSP-IS0102_1.0.0.pdf

[2] http://www.dmtf.org/standards/published_documents/DSP-IS0103_1.0.0.pdf

[3] http://en.wikipedia.org/wiki/ISO/IEC_27001

CHAPTER 6

Reflections on Cloud Services

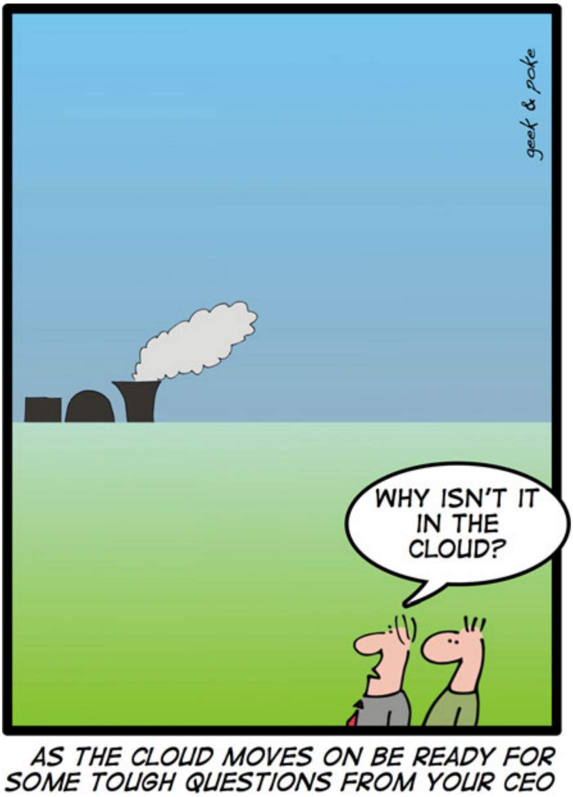
geek & poke
WHY ISN'T IT IN THE CLOUD?
AS THE CLOUD MOVES ON BE READY FOR SOME TOUGH QUESTIONS FROM YOUR CEO

A tough CIO asks more tough questions—Geek and Poke

So far, we have looked at the high-level models for cloud services and the key new attributes relative to traditional IT infrastructure or outsourced solutions. We have also looked at the worldwide changes occurring that have allowed these possibilities to exist. In the next chapters, we look at how to plan for using cloud services now that the common architectural approaches have been illustrated and a common taxonomy explored. In this brief chapter, we reiterate how important the understanding and strategic plan to use cloud services is.

The availability and speed of the Internet (broadband access) is increasing rapidly and significantly worldwide, and the impact is fundamental and immense. The rise of mobile or wireless networks that operate with data at their core has enabled incredible advances in the ability of smartphones, indeed any device, to integrate into a mesh of data services.

Moreover, all this is informing public policy makers worldwide, creating even more reason to provide services in the cloud for all types of businesses. It also illustrates that governments are expecting to move significant operations online using cloud services.

In 2009, Finland became the first country in the world to create a law guaranteeing that every citizen in the country a right to a broadband connection. Although the decision was that 100 megabit broadband would be a legal right by 2015, an interim step was defined for June 2010 that required at least a 1-megabit access option.[1]

In June 2009, Gordon Brown, the then U.K. prime minister, stated: "Whether it is to work online, study, learn new skills, pay bills or simply stay in touch with friends and family, a fast Internet connection is now seen by most of the public as an essential service, as indispensable as electricity, gas and water."[2] The U.K. government in 2010 made a formal commitment to move government services for all citizens online, with the concept of a personal web page for everyone who provides access to all required services relevant to that individual, from health care and licensing to local council permits.

In April 2009, the Australian government announced a plan[3] to create a national broadband network in conjunction with private industry. The network is intended to provide up to 100-megabit Internet access for at least 90% of all Australian homes, schools, and workplaces using fiber, and a minimum of 12 megabits to other more remote locations using wireless and satellite technologies.

It is obvious that more and more opportunities for business worldwide will be via the Internet.

In an environment where pressure is on IT to better support business goals and further improve cost efficiencies, cloud services have emerged as a gold ring. IT infrastructure is reaching a breaking point:

- In distributed computing environments, up to 85% of computing capacity sits idle.
- Consumer product and retail industries lose about $40 billion annually due to supply chain inefficiencies.
- An explosion of information is driving a 54% growth rate in storage shipments every year to accommodate the data.
- Security breaches are becoming more common.
- Challenges to automation mean highly skilled individuals are supporting basic operational tasks in data centers instead of innovating to serve customers.

The primary goal of this section is to ensure that both IT and business architects have a clear view and common taxonomy for discussing and delivering cloud services.

Alongside that is a need to survey an industry and set of ecosystems that are evolving to deliver and support cloud services.

Finally for this section, a word of warning: The industry is evolving exceptionally fast. People may get trapped if they rush into the cloud, but the opportunities are too great to ignore and having a plan, both tactical and strategic, is our critical goal for you.

In April 2009, the U.S. Government Services Administration (GSA) released a blanket Purchase Order (BPO) for cloud services. The BPO was written to allow U.S. government agencies to rapidly utilize a specific set of approved cloud services. On March 1, 2010, the GSA announced that it was withdrawing the BPO to rewrite the order, citing two main reasons, both of which should be critical points of consideration in your strategic cloud planning:

- The technology related to cloud had changed significantly enough that a review was merited.
- The security requirements in the original BPO did not meet the needs that had surfaced since the first release.

On the first point, Dave McClure, the associate administrator for the GSA's Office of Citizen Services and Communications, was interviewed on Federal News Radio on March 1, 2010,[4] and he stated: “It is a challenging market. Eleven months in cloud time is 11 years in computing time. This is a fast changing market in terms of the offerings, as well as the experience—both customer as well as provider. So, I think we have to make sure that we’re understanding that market as it evolves. Whatever we do, we want to do it right and we want to make sure that we’re focusing on the highest value solutions that we can for customer needs.”

On the second point, McClure stated: “Our initial offering that we put out was FISMA low impact. That was because of the market maturity at the time. I think there’s been some significant recognition that there is a bulk of government services needed for data in the low-impact area, but there’s a growing need for the moderate impact area. So, we wanted to change that requirement, as well, in the hopes of pulling in more products and services.”

Therefore, before you decide how and where to use cloud services, understand the speed of the market, and ensure you are cognizant of the need to evolve with the market at that speed.

In terms of understanding how cloud services are changing businesses, consider that Peter Horrocks, who took the role of director of BBC Global News in March 2010,[5] said it was important for editorial staff to make better use of social media and become more collaborative in producing stories. “This isn’t just a kind of fad from someone who’s an enthusiast of technology. I’m afraid you’re not doing your job if you can’t do those things. It’s not discretionary.”[6]

The huge shift of power, both in terms of producing and distributing content, has resulted in an arguably cataclysmic impact on traditional media organizations such as newspapers. The technology mentioned in the story on Horrocks included Facebook, Twitter, and RSS Readers, all cloud-related, and all part of the media and content production revolution. The point is to illustrate how critical it is to get ahead of these technologies in your business, because similar impacts are being felt in all verticals—public sector, government, finance, manufacturing, media, technology, and so forth.

It's time to consider our planning and management of cloud services.

Endnotes

[1] http://yle.fi/uutiset/news/2009/10/1mb_broadband_access_becomes_legal_right_1080940.html?origin=rss

[2] http://www.timesonline.co.uk/tol/comment/columnists/guest_contributors/article6506136.ece

[3] http://www.minister.dbcde.gov.au/media/media_releases/2009/022

[4] http://federalnewsradio.com/?nid=19&sid=1900641

[5] http://www.bbc.co.uk/pressoffice/pressreleases/stories/2009/02_february/26/horrocks.shtml.

[6] http://www.guardian.co.uk/media/pda/2010/feb/10/bbc-news-social-media.

PART III

Life in the Cloud—Planning and Managing the Cloud

IN THIS PART

CHAPTER 7

Introduction to Cloud Planning

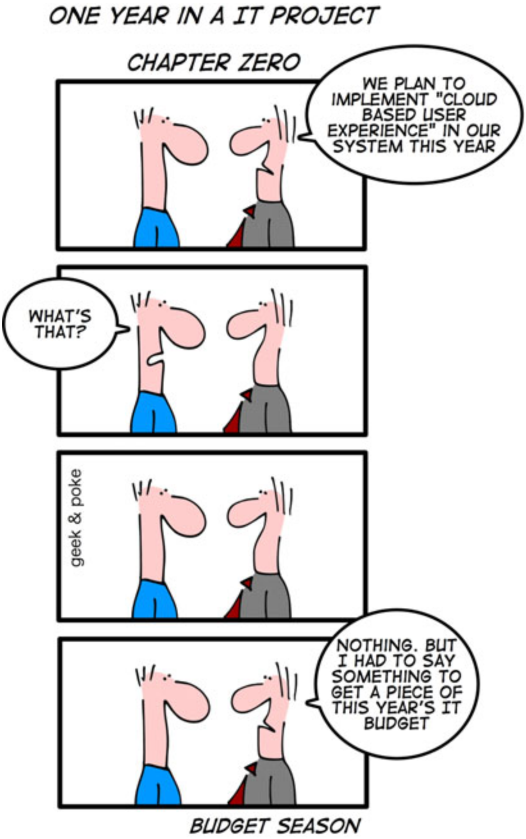
ONE YEAR IN A IT PROJECT
CHAPTER ZERO
WE PLAN TO IMPLEMENT "CLOUD BASED USER EXPERIENCE" IN OUR SYSTEM THIS YEAR
WHAT'S THAT?
geek & poke
NOTHING. BUT I HAD TO SAY SOMETHING TO GET A PIECE OF THIS YEAR'S IT BUDGET
BUDGET SEASON

One year in a IT project: The one and only important chapter—Geek and Poke

The potential for cloud services is larger than that of service-oriented architecture (SOA), and therefore the hype has been greater. However, just because the hype has been steep does not mean that the benefits of cloud services will come automatically. This chapter introduces the key topics and main challenges that you must address to benefit most from cloud services. We take a closer look at how to best manage them in subsequent chapters where we consider planning and managing cloud services in detail.

Cloud services, as with any newly hyped technology solution, tends to be rushed into by organizations that seek only a technical solution. However, for organizations to truly realize the benefits that cloud services offer, they must think about various management disciplines, including IT strategy, architecture planning, governance, and service management.

The maturity of these disciplines can be seen as a litmus test that indicates the impact and effectiveness that cloud services will have on your organization.

In essence, this necessitates that your organization approaches cloud services in a strategic manner rather than with a purely narrow project focus. These disciplines will assist you in finding the right balance between owning and executing a cloud infrastructure and the associated risks of using a cloud services vendor that offers the same service.

Your IT organization's raison d'être will need to transition from develop and build to architect, manage, and govern.

The transition of your IT organization's role will have to occur incrementally over time in alignment with your risk comfort level.

Enterprise Architecture

Enterprise architecture (EA) as a discipline has on the whole been adopted by large organizations. EA can be seen as an approach that is used to manage and align an organization's assets (for example, services, people, technology, and projects) and defines how they will support the organization's business goals and aspirations.

Even small and medium-size organizations that consider enterprise architecture as too heavyweight for their needs still need to think about planning, architecting, and managing cloud services strategically.

If your company utilizes EA as a discipline, it should be used as a major lynchpin when defining and executing your cloud strategy. EA will assist you in making cloud service decisions around enablement/alignment with the business and IT, instead of focusing purely on the technical needs. It can thus concentrate your organization's efforts on delivering value, both short and long term, instead of centering on purely a technical solution.

The allure of cloud services is compelling, but it is still imperative that any EA effort that includes a focus on cloud services must identify and prioritize the business and IT reasons for doing so (see Figure 7.1).

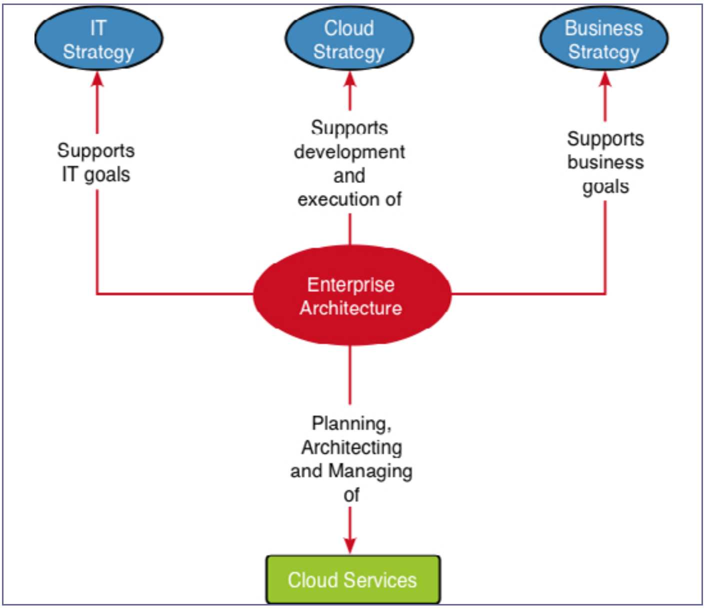

FIGURE 7.1
Enterprise architecture can assist with the planning and architecture of cloud services.

This not only assists in the development and prioritization of the phases of your cloud strategy roadmap but also lays the foundation for cloud governance decisions. You will need to develop a compelling, incremental, and pragmatic cloud strategy that identifies the critical aspects of your organization and how cloud services will support and deliver on their defined goals.

Without proper strategic planning, you will be leaving yourself open to consuming and producing cloud services that are costly, unjustified, and unplanned, which in turn can lead to so called silos in the sky, which will eventually lead to integration, security, and data-privacy issues.

Many enterprise architects will be looking to see what possible alignment they can utilize from their existing EA models when supporting the strategic and architecture requirements of cloud services. Figure 7.2 highlights example areas where EA can support and assist in the identification and planning of cloud services.

Although there are undoubtedly many areas where EA can assist in the planning and identification of cloud services, you still need to consider existing solution architecture and design disciplines to assist in the actual delivery and consumption of these cloud services.

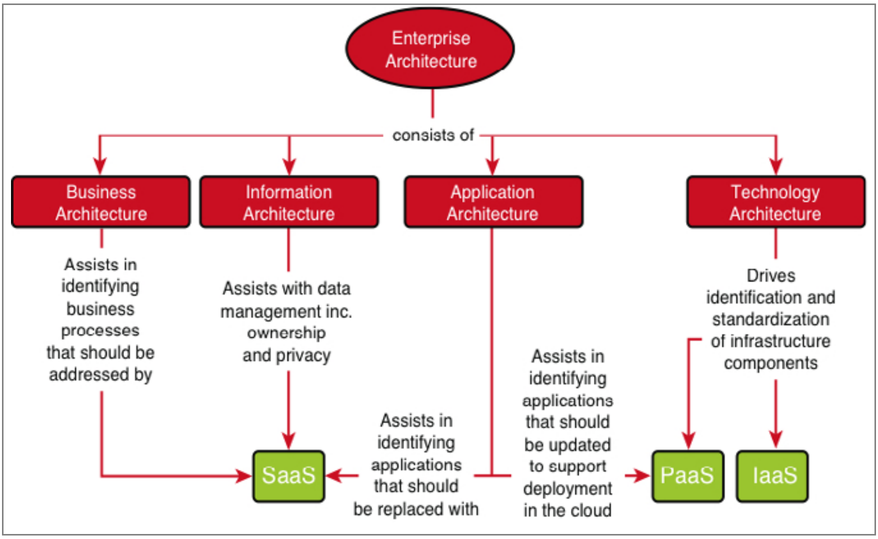
Enterprise Architecture
consists of
Business Architecture
Information Architecture
Application Architecture
Technology Architecture
Assists in identifying business processes that should be addressed by
Assists with data management inc. ownership and privacy
Assists in identifying applications that should be replaced with
Assists in identifying applications that should be updated to support deployment in the cloud
Drives identification and standardization of infrastructure components
SaaS
PaaS
IaaS

FIGURE 7.2
Examples of areas where EA can support cloud service planning

Solution Architecture and Design

Consuming and producing cloud services does not mean abdicating the need for good solution architecture and design practices. Your existing solution architecture and design practices should be reviewed to confirm whether it can appropriately take advantage of cloud services. One solution architecture and design paradigm that has garnered a lot of focus in recent times is service-oriented architecture (SOA).

Although SOA is not a prerequisite for executing your cloud strategy and is by no means the only solution architecture pattern that can be used for producing and consuming cloud services, it is definitely the architecture pattern getting most attention from the cloud community.

Service-Oriented Architecture

Much has been said of the linkage and cohesion of cloud computing and SOA. There have been many misconceptions that cloud computing is just an extension of SOA. Although they do share many core principles, and benefits, cloud computing is by no means an extension of SOA. They are categorically different technology strategies. The mutual benefits that cloud computing and SOA bring seem obvious at first glance. In effect, it is 2 + 2 = 5. On their own, each offers benefits, but bring them together and the sum is greater than their parts. Nevertheless, it is not hard to see that bringing together two enterprise technology strategies that focus on modularity and speed is very compelling. For example, SOA can deliver services driven by a set of business goals that can be seen as the enabling piece behind software as a service (SaaS) and aspects of platform as a service (PaaS), whereas cloud computing enables the dynamic provisioning of service-oriented services and service-oriented infrastructure components.

Cloud computing is in effect reviving interest in SOA, but there are still challenges that need to be addressed. At IDC's SOA and Beyond Conference 2010 in London, IDC stated that cloud services would actually boost the business spending on SOA as cloud computing is a catalyst for new SOA sales. Just because SOA is a good fit as an architecture doesn't mean that you can take advantage of what cloud services offer without effort.

A level of SOA execution maturity is required before an organization can take full advantage of the benefits that the combination of cloud computing and SOA offers.

With architecture and software development teams being asked to deliver more in a more expedited way, it is not surprising that organizations are turning to cloud services to help address these needs. Cloud services enable SOA teams to spend time innovating rather than maintaining infrastructure and configurations. One area of software development in which cloud services has really helped is testing; organizations can set up and tear down development and test environments rapidly.

As part of your SOA and cloud governance strategy, make sure that your SOA architects and developers do not give architecture and design short shrift. It is all too easy to fall into the trap that cloud services offer you limitless access to resources, which in turn gives architects and developers a false sense that architecting and developing scalable and performant SOA services is not important any more.

If your organization intends to develop for the cloud and not just consume external cloud services, you must be aware of the changes to your current architecture and design philosophy (see Figure 7.3). Your organization should familiarize itself with secure coding techniques, multitenancy data and application patterns, and cloud integration practices. Cloud services change the way in which business services and applications are architected, designed, composed, and managed.

One major requirement that feeds into a solution architecture and design is the capability to interact and integrate with other solutions, whether those solutions are built in-house or sourced from public clouds. A key success criterion will be how well the data stored in a public cloud service integrates with your in-house production systems. This level of integration is best achieved via SOA and an understanding of data and application interoperability within multitenant environments. Without this capability to integrate with your internal systems, any proposed solution will probably negate the benefits and cost savings that the SaaS solution offers.

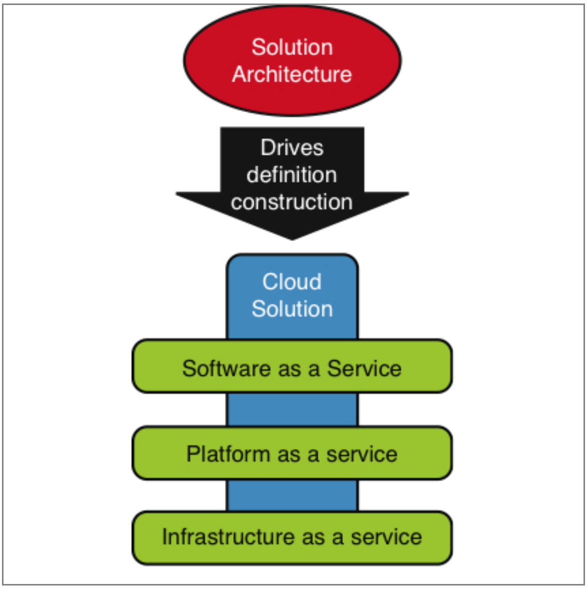

FIGURE 7.3
Which services and infrastructure components shape a solution?

A number of different approaches can be used to achieve this integration:

- Current investments in middleware products such as an enterprise service bus can be leveraged to facilitate the integration between in-house systems and public cloud services. Existing middleware vendors are slowly making available integration templates for popular public cloud service providers to expedite the integration effort.
- A category of cloud service brokers has started to emerge that offers the confusing acronym of IaaS (integration as a service), such as Boomi and Cast Iron. These cloud service brokers offer integration functionality in the cloud that provides application and data mediation between cloud services whether they are located on-premise or off-premise.

Data and Application Interoperability

It is important that both data and application systems expose standard interfaces. Organizations will want the flexibility to create new solutions enabled by data and applications that interoperate with each other regardless of where they reside (public clouds, private clouds that reside within an organization's firewall, traditional IT environments, or some combination). Cloud service providers need to support interoperability standards so that organizations can combine any cloud service provider's capabilities into their solutions.

Protecting sensitive corporate and customer data should be a priority if you're considering a virtualized environment that enables a cloud service provider to manage or store your organization's data.

Before you put your organization's data in the hands of a cloud service provider, demand that the cloud service provider demonstrate its data-protection and business-continuity capabilities. In addition, when you decide to move forward, make sure that your negotiated agreement is explicit about the cloud service provider's ongoing obligations to protect your data and holds the cloud service provider liable for failure to satisfy those obligations.

If your company operates internationally or in certain industries in the US (for example, financial services or health care), your negotiated agreement should require the cloud service provider to comply with applicable data-protection and privacy laws.

The negotiated agreement should also do the following:

- Review which data to put in the cloud and what legislation has to be followed.
- Incorporate the relevant portions of your privacy policies and obligate the cloud service provider to conform to them.
- State that your company owns its data, has access to that data at its discretion, and will receive the data upon the expiration or termination of the agreement.
- Describe the parties' responsibilities when it comes to recovering lost data.
- Clearly define data retention and destruction policies.

Identify in what geographic region the data is located (including replicated and backed-up copies), and ensure the cloud service provider warns you ahead of time of any changes in that situation. Country-specific regulations governing privacy and data protection vary greatly. In particular, a state or federal government could order a subpoena that can force cloud service providers to turn over its records. Such subpoenas have no judicial oversight, meaning that your privacy rights would be compromised. To help you grasp this issue at a high level, Forrester created a privacy heat map[1] that denotes the degree of legal strictness across a range of nations.

Data in the Cloud

Data is kept in the cloud in a multitenant environment. This means that one environment stores the data for several client organizations. The data is virtually partitioned to give each client organization a dedicated virtual environment. However, when discussing data security, we need to look at how the partitioning is done. Three possible approaches can be used:

- Separate databases
- A common database with separate client areas
- A common database with a common client areas

Let's discuss the pros and cons of each approach.

Separate Databases

This is the best approach for the client organization and the worst for the cloud service provider. Each client organization has its own database in the cloud (see Figure 7.4). Therefore from a security perspective, this is the best.

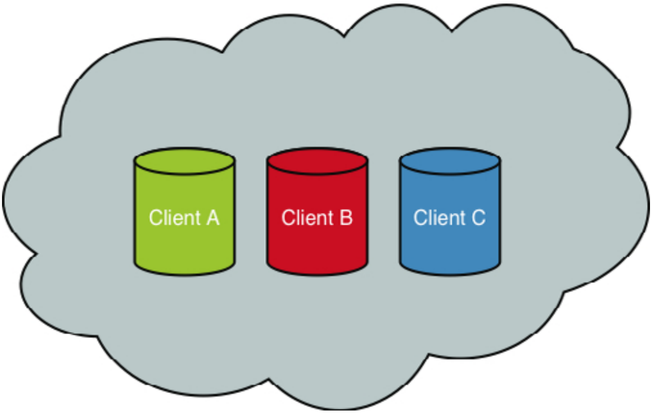

FIGURE 7.4
Separate databases for each client

Unfortunately for the cloud service provider, only a limited number of databases can be served by a given server, and maintenance and backup costs are higher than with the other two options. Therefore, all in all, this is the most secure, but highest-cost option.

Common Database with Separate Client Areas

In this option, a number of client organizations use a common database, but each client organization is allocated its individual areas to store data (see Figure 7.5).

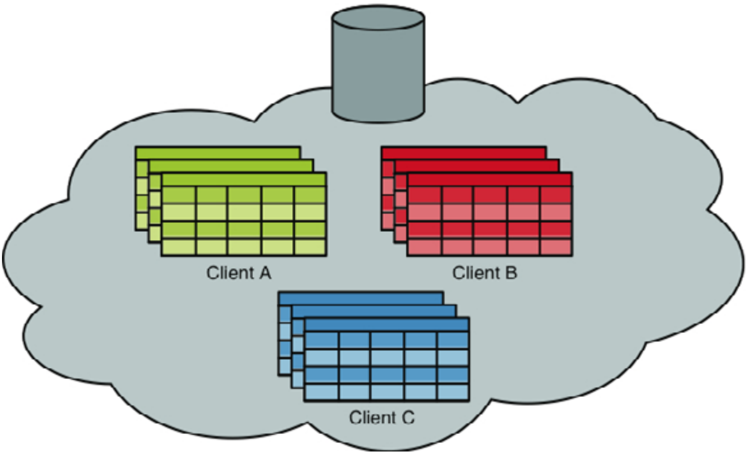

FIGURE 7.5
Client organizations share a common database.

More client organizations can be hosted on a single database server, making it cheaper for the cloud service provider to operate. However, in case of failure, restoring the data from a single client organization is more difficult. This approach is less secure than the previous one; skilled hackers may be able to access other client organizations' data because they have access to the database. All in all, it's less secure but cheaper to operate.

Common Database with Common Client Area

One database with a common client organization area to store data is used for multiple client organizations. Each client organization has a unique ID to identify which client organization owns the information (see Figure 7.6).

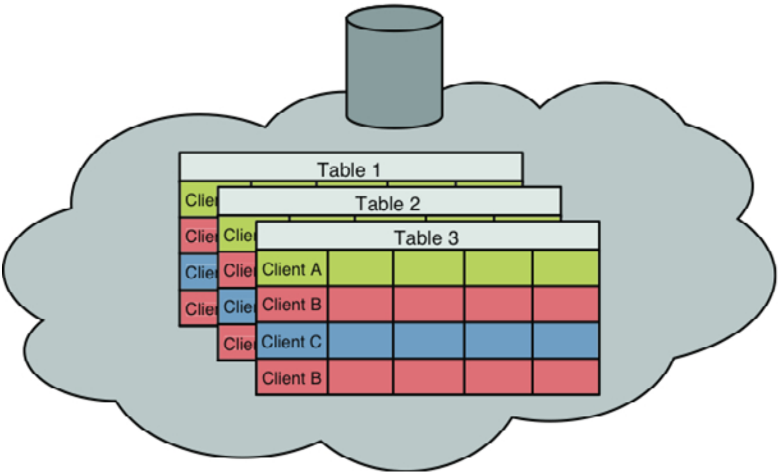

FIGURE 7.6

Tenants share a common database and data tables.

This approach allows the largest number of client organizations per database server, and the lowest backup costs, but restoring the data for a single client organization is extremely complex. Such operation requires a fair amount of resources, possibly reducing performance for other client organizations. This is obviously the least secure approach, although in some implementations, the client organization data is encrypted to reduce security risks.

Whether you intend to store data in a private cloud or a public cloud, to assess the risks associated with storing your data in the cloud, it is important to understand which approaches are being used by your cloud service provider.

Recovering Data from the Cloud

Another critical aspect of data management is a plan to extract critical data from the cloud service. Let's assume you have used Salesforce.com for the past several years. Your enterprise has grown dramatically, and now you want to implement your own customer relationship management (CRM) tool. How do you get your data back from the cloud? This question, although often forgotten when deciding to move to the cloud, should be an integral part of the contractual negotiations. It is similar to signing a prenuptial agreement; you hope it will never be necessary, but it's there just in case.

In the case of Salesforce.com, if you are an enterprise or unlimited edition customer (and for a fee, professional edition customers can do this), you can, from the Salesforce.com user interface, export all your data. The result is a zip file of CSV files containing the data for each object you have set up. Application providers such as SAP now have facilities by which you can upload those files in their system. However, if you decide to use one that does not, you will have to dig through the material, identify each field, map it, and extract the data. This takes quite some effort but is obviously feasible.

Our example scenario is important in that we have chosen to look at a SaaS provider. In this case, the SaaS solution also has another set of data that would be considered critical to most businesses: the business or workflow logic created during the use of the SaaS solution. Although businesses increasingly offer migration services as discussed, such migrations are generally not considered as part of the planning process—the first time around at least.

We can only stress the importance to have clear agreements right from the start to ensure that regardless of what you decide to do moving forward, your data is protected and you can recover it.

Managing Workloads in the Cloud

The key goal in managing workloads inside cloud services is to marry your data classifications and cloud usage models with your cloud service provider cost model. Figure 7.7 illustrates various types of workload that cloud services can accommodate much more easily than a data center with resources dedicated to specific tasks or application services.

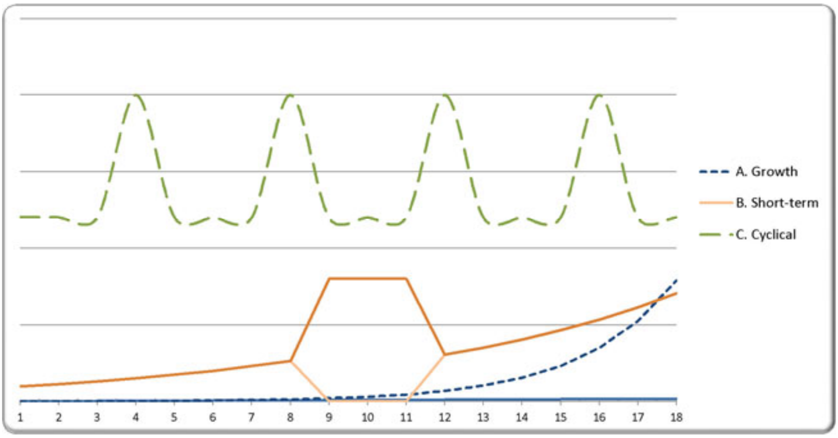

FIGURE 7.7
Workload flexibility as managed in the cloud

In each of these cases, a specific usage pattern illustrates an obvious match for cloud services. This also means that any cloud service you utilize for managing changes in workloads, regardless of deployment model, must meet the requirements for unlimited and rapid elasticity. Taking into consideration the data classification is important when you need to determine the right location and provider solution for moving workloads. More specifically, workloads that involve certain classifications of data, perhaps highly sensitive data, may limit you to using specifically secured cloud services. For example, if the data is related to payment cards, it may restrict you from moving the workload to a public cloud based on PCI compliance. If the data contains personal information of a German citizen, you may be restricted to maintaining the data inside Germany's national borders. Without data classification, you will not know where and when you can move to optimize your use of cloud services. For example, if you cannot classify data appropriately, yet utilize cloud services for backup, you will not know whether your backups will meet your compliance obligations! These issues are more apparent when considering public cloud options but cannot be ignored when architecting a private cloud.

Governance and Management

As IT departments introduce cloud solutions in the context of their traditional data center, new challenges arise. Standardized mechanisms for dealing with life cycle management, policy management, licensing, and funding for shared cloud infrastructure are just some of the management and governance issues cloud service providers and consumers must resolve.

Traditional approaches to governance do not go far enough and do not suffice for enterprise cloud solutions, and therefore you will be required to imagine new ways for things to fail and to assess and govern your risk.

Therefore, extend your existing governance, risk, and compliance processes to cater for the new or elevated risks associated with consuming and producing cloud services. This will require you to increase the scope, maturity, and reach of your governance processes to cater for IT assets no matter where they reside, on-promise or off-site.

Whether you intend to consume/offer cloud services via a private or public cloud, your organization will need to acquire new skills to both govern and manage these cloud services. Example skills required include defining service contracts and service level agreements and assessing vendor viability, covering topics such as a cloud service provider's security processes, service management capabilities, and the approach to data governance/privacy. Organizations that have used outsourcing successfully will have an advantage here. The big question you must ask yourself is how your organization will adapt, govern, and capitalize on cloud services.

Cloud governance must not be deployed in a siloed manner but must interact with and complement the other governance disciplines with your organization. This is especially true of information governance, which includes the processes for classifying information and understanding risk so that policies can be put in place that specify which cloud-based services and applications are appropriate and which are not. This topic is covered in more depth in later chapters.

Security

One of the major challenges of cloud governance is making sure that cloud-based resources such as data are secure and that the appropriate policies have been defined and enforced. Many organizations are not comfortable storing their data and applications on systems they do not control, yet they accept risks in cloud services as part of the new approach.

Migrating workloads to a shared infrastructure increases the potential for unauthorized access and exposure. Consistency around authentication, identity management, compliance, encryption, and access technologies will become increasingly important. To reassure their customers, cloud service providers must start to offer a higher degree of transparency into their operations.

There are few places in the world where you would be confident leaving your keys in an unlocked car. Although some places require you to leave your keys with a valet, you generally think to lock up your car before you leave it anywhere outside your own garage. Yet leaving your keys in an unlocked car is what many organizations do when they put their valuable data into the cloud. Even if an organization considers the cloud service provider to be a valet (in this analogy), it must be clear that the organization is still responsible for the security of its data. Yet only some protections exist in terms of contracts and liability even as cloud services are being increasingly used. This topic is covered in more depth in later chapters.

Summary

This chapter highlighted a number of IT disciplines that organizations must update and refresh to cater for the nuances of cloud services. Doing so can enable your organization to extract the benefits and control the risks encountered when planning, consuming, producing, and governing cloud services.

Key points to consider from this chapter are

- Adopting cloud services requires the strategic planning, architecting, and management of cloud services.
- Traditional approaches to governance do not go far enough and do not suffice for enterprise cloud solutions.
- Consuming and producing cloud services does not mean abdicating the need for good solution architecture and design practices.
- Data classification is a critical, yet often ignored, part of good solution architecture and design practices in general but especially when using cloud services.
- Protecting sensitive corporate and customer data should be a priority if you’re considering a cloud service provider to manage or store your organization’s data.

The remaining chapters in Part III, “Life in the Cloud—Planning and Managing the Cloud,” examine in detail the disciplines discussed in this chapter. This lays the foundation for Part IV, “GPS to the Cloud…Where to Now?,” which focuses on developing a pragmatic cloud services adoption roadmap.

Endnotes

[1] http://www.forrester.com/cloudprivacyheatmap.

CHAPTER 8

Cloud Governance, Risk, and Compliance

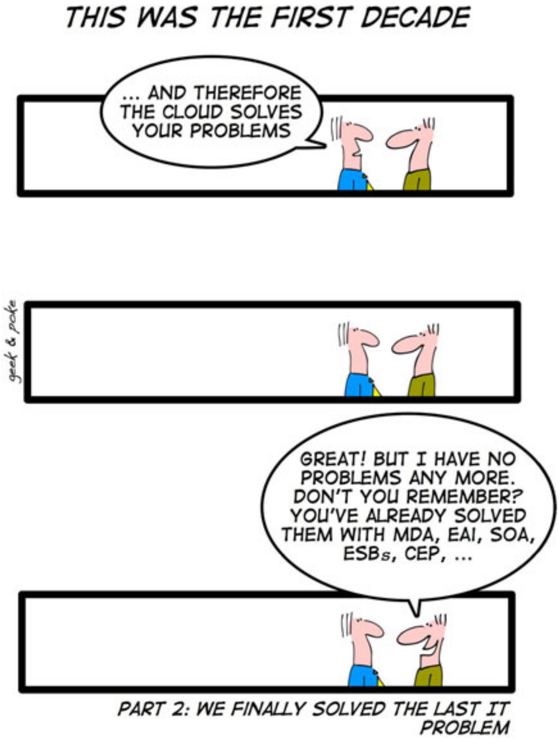
THIS WAS THE FIRST DECADE
... AND THEREFORE THE CLOUD SOLVES YOUR PROBLEMS
geek & poke
GREAT! BUT I HAVE NO PROBLEMS ANY MORE. DON'T YOU REMEMBER? YOU'VE ALREADY SOLVED THEM WITH MDA, EAI, SOA, ESBs, CEP, ...
PART 2: WE FINALLY SOLVED THE LAST IT PROBLEM

We finally solved the last IT problem—Geek and Poke

Although cloud services offer many benefits, these benefits are unsustainable without a level of planning. Without proper planning, control, and monitoring, obtainable cloud service benefits can be easily ignored, forgotten, or more important, undermined by the actions of individuals within your organization who make uninformed or reckless decisions.

Consider this scenario:

- The sales department within one division decided to use a customer relationship management (CRM) cloud service that offered quickly accessible and good enough functionality at a seemingly low price. Compare this to the long-established approach it had taken in the past of requesting server hardware, software, and configuration that would have taken months, if the request were even justified in the first place. Bypassing this perceived overhead and reacting much more quickly is obviously attractive to the business leaders.
- The IT department became aware of this decision only when the sales department requested to integrate its CRM cloud service with the on-premise billing system. The IT department surveyed other departments to discover whether additional cloud services were used.
- After some investigation, the IT department discovered that several sales departments within different divisions had taken the same initiative to utilize CRM cloud services from various CRM cloud service vendors. This led to software as a service (SaaS) sprawl whereby customer information was duplicated, stale, and incomplete. This type of scenario is not merely isolated to business divisions but can also be seen within IT departments.
- Instead of waiting weeks or months, IT software engineers who had a short timeframe in which to demonstrate a proof of concept felt using platform as a service (PaaS) cloud services was preferable to achieving their short-term goals rather than use the long-established approach they had taken in the past.
- Although in the short-term this seemed to solve their current needs, unforeseen business circumstances forced the IT software engineers to use the proof-of-concept solution as a production solution. On reflection, this seemed an acceptable risk.
- The increased competitive cloud marketplace coupled with recent recessionary pressures forced the PaaS cloud service provider into liquidation. This left the IT department with major issues. Utilization of the PaaS cloud service provider's proprietary API made the portability of the solution to the on-premise data center unachievable within acceptable timescales.

Chaos can ensue even in situations where organizations have made a formal decision to use IaaS services judiciously. Organizations can be left in administrative turmoil when many virtual machines (VMs) differ slightly from each other in undocumented ways leading to VM sprawl.

As highlighted in this scenario, it is all too easy for individual line of businesses to use its discretionary budgets to utilize cloud services without considering the effect this could have on its organization.

Rather than attempting to address previously undetected cloud services in use by your organization, a more proactive approach must be embraced.

As your organization considers or adopts cloud services, you will be required to make people, process, and technology updates to help ensure you continue to meet strategic objectives and drive cultural change, while at the same time manage the newly encountered cloud risks.

These risks require organizations to focus on developing better governance and standardization strategies. Therefore, governance, risk, and compliance disciplines play an important role in your cloud strategy and should not be considered as an afterthought.

Governance, risk management, and compliance are three disciplines that in the past have existed in silos within organizations. Leaving these disciplines in silos invariably leads to information, communication, and execution challenges that can circumvent their overall effectiveness. Figure 8.1 highlights key tasks from these disciplines that must be integrated and aligned to make sure that their maturity, influence, and value are in sync and do not undermine each other.

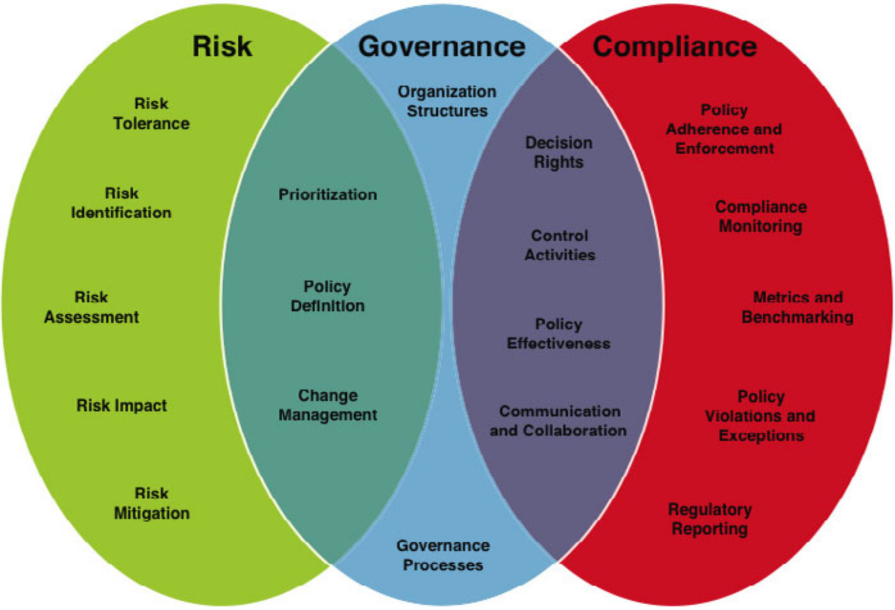

FIGURE 8.1
Governance, risk management, and compliance must be integrated.

Adopting a unified cloud governance, risk management, and compliance approach creates a proactive organized and efficient structure to identify, define, monitor, and enforce policies that enable the extraction and sustainability of the investment made into your cloud infrastructure.

Risk Management

Although it may be easy to be swayed into adopting and utilizing cloud services as a direct means to IT cost savings, cloud services present many unique situations for you to address. Therefore, you must consider and weigh up the risks against any monetary gains.

We've discussed many of the risks associated with cloud services throughout this book so far. Note that many of the risks are not new and can be found in organizations today, but cloud services heighten or elevate those traditional risks to greater levels. Critically, cloud services create new and varied risks.

The level and types of risks will vary with the type of cloud service models and cloud deployment models being considered.

Although security and data privacy risks have taken center stage in the press, it is important that your cloud risk management approach extends beyond security and privacy of data and include additional risk categories such as brand/reputation risk, financial risk, regulatory risk, and cloud service providers not meeting service level agreement (SLA) commitments.

Consider updating your existing risk management process to cater for the new risks that adopting cloud services will bring. This process must be agile and flexible enough to deal with the continuously evolving risks that cloud services are highlighting. Your risk management process should be well understood and communicated to the relevant stakeholders such that they can identify risks; analyze possible impact, probability, and timeframe; determine priority and manage through accepting, avoiding, mitigating, or transferring the risks.

To expedite your own risk identification and management efforts, use the many documents available from many of the standard organizations that are focusing on these efforts, including the following:

- The Cloud Security Alliance (CSA), focused on the central issues of cloud security, has created the "Security Guidance for Critical Areas of Focus in Cloud Computing." The document offers guidelines structured into 13 domains, of which 5 focus on governing in the cloud, as follows:
    - Governance and Enterprise Risk Management
    - Legal and Electronic Discovery
    - Compliance and Audit
    - Information Lifecycle Management
    - Portability and Interoperability

  The CSA continues to develop documents and tools (for example, "The Top Threats to Cloud Computing" and the "Cloud Computing Controls Matrix") that can assist you in making educated risk management decisions regarding your cloud adoption strategies.
- The ISO/IEC 27000 series (also known as the ISMS Family of Standards, or ISO27k for short) comprises information security standards published jointly by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC). This series provides best practice

recommendations on information security management, risks, and controls within the context of an overall information security management system (ISMS), similar in design to management systems for quality assurance (the ISO 9000 series) and environmental protection (the ISO 14000 series).[1]

- The European Network and Information Security Agency (ENISA) has documented key security risks and recommendations for cloud computing around information security.
- The Jericho Forum within the Open Group (TOG) offers a free self-assessment scheme to assist with the assessment of vendors to aid in deciding whether they offer the security solution you require.
- The American Institute of Certified Public Accountants (AICPA) offers a widely recognized auditing standard titled the "Statement on Auditing Standard No. 70" (SAS 70). This standard represents that an organization has been through an audit of its control objectives and control activities, which often include controls over information technology and related processes. It is advisable to ask a cloud vendor for its type II report as it is also includes a testing portion that type I does not cover. (Type I and type II are discussed in Chapter 5, "State of the Industry.")

When analyzing and assessing the risks of adopting cloud services, take into account the risks of staying with your traditional approach to delivering business and IT functionality. In addition, consider how the risk is affected when using different service deployment models. This opens the possibility of making the best use of both the public and private clouds. Last, compare different cloud service provider offerings when assessing or minimizing the identified risks.

Any risk assessment task must be the result of the collaboration between both the business and IT, including departments such as legal and finance. This collaboration enables the defining of risk tolerances and the acceptance of the consequences of the risks.

Traditional risk management makes it possible to transfer some identified risks to third parties, but when it comes to adopting cloud services, you can transfer the responsibility but you can't transfer the accountability. Work with cloud service providers to obtain assurances.

Having a structured approach to risk management can enable your organization to decide which cloud services to adopt, after taking into consideration the business criticality and the acceptable level of risk of the data on which the cloud service providers operate.

Governance

As a discipline, governance has been with us for many years, but with the advent of cloud services, the need and importance has been elevated for organizations to take governance more seriously. Effective cloud governance provides organizations with visibility into and oversight of cloud services used across the organization. To effectively manage and optimize your organization's investment in cloud services and improve decision making, effective cloud governance must encompass people, processes, and technology.

Cloud governance should not be confused with cloud management.

Cloud governance defines the actions that your organization requires to achieve your strategic cloud computing goals. Cloud governance includes the definition of

- Who holds the authority to make decisions?
- Who has the accountability for these actions?
- Who has the responsibility for the outcomes?
- How will the expected performance be evaluated?

Your cloud compliance efforts will be enhanced if your organization has clearly established and communicated expectations for cloud services and the policies that your organization must follow. This requires that existing processes and organizational structures be reviewed and updated to cater for a newly formed communication and collaboration model. For example, an existing IT steering committee could be updated to include representatives from both legal and finance to provide input into cloud considerations.

A good cloud governance model accelerates change. This will be critical in addressing the challenge of culture change and process disruption that cloud services will present to your organization. The basis for a cloud governance model within IT will shift from focusing purely on the underlying infrastructure components to the cloud services and their associated SLAs.

Traditional governance and management frameworks such as COBIT and ITIL should not be seen as a panacea for cloud governance but do describe the major controls and processes required to address this service delivery model, such as service portfolio management and supplier management. As with all frameworks, it is critical that these frameworks be customized for your dynamic and rapidly changing cloud environment. Initially, you will require more controls due to unidentified and unfamiliar challenges that cloud services and external cloud service providers will bring to your environment.

Compliance

Compliance is an iterative process that applies the mitigation strategies defined within the risk management process and ensures that individuals and systems adhere to defined cloud adoption policies.

These risk management strategies are analyzed so that the appropriate polices and control points can be internally employed as well as enforced on the cloud service provider. These policies derive from internal directives addressing previously identified cloud adoption risks, regulatory requirements, laws, internal standards, and best practices.

The demands of regulatory requirements—such as the European Communities Data Protection Directive 95/46/EC and related member states laws, the U.S. Health

Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standard (PCI DSS)—to name a few, highlight the need to implement controls around security and data privacy. These two areas of focus have garnered the most attention when adopting cloud services. Therefore, they are covered in more detail in the next chapter.

As mention earlier, you can transfer responsibility, but you cannot transfer accountability. For example, when it comes to data compliance, it is the data owner and not the cloud service provider who is responsible for compliance. California law SB 1386 is an example of why you cannot transfer accountability. CA SB 1386 requires that any state agency, person, or business that conducts business in California, and has computerized data that includes personal information, must disclose any breach of the security of the data to any resident of California whose unencrypted personal information was or is reasonably believed to have been acquired by an unauthorized person. Perhaps you have received a letter from a business that has your personal information informing you that such an event has occurred and that your personal information is at risk. This breach of security has been highlighted in recent times with the increase in news articles that report stolen laptops with unencrypted data.

Therefore, when using cloud services, you must understand where the data resides, how it is protected, and where it has been transmitted. This can be a challenging proposition because many cloud service providers are not forthcoming with this information. So, when you consider using cloud services, perform a gap analysis between the specific requirements identified in relevant regulations and the set of controls provided by the cloud service provider.

A key requirement of compliance is the continual monitoring of the cloud ecosystem to ensure that expected outcomes are being met. This continual monitoring enables a feedback loop to decide whether the policies and control points need to be updated to achieve your requirements and goals.

Policy adherence and monitoring must be made as unobtrusive as possible; otherwise, individuals will circumvent the process and utilize cloud services without the appropriate authority.

Therefore, automate as many decision and control points as possible to keep a low-touch approach and expedite resolutions. Otherwise, individuals will not be asking permission to utilize a cloud service; instead, they'll be asking for forgiveness after they have already done so. Beyond continuous might allow you to obtain access and network logs, most will not as these are not always segregated for multitenant environments in the same way that administrative applications and related reporting capabilities are. Cloud forensics will be discussed in more detail in Chapter 9, "Cloud Business Risk and Security."

Cohesive GRC

A cohesive cloud governance, risk, and compliance (GRC) approach eases the transition and ongoing adoption of cloud services. Such an approach enables you to reduce risk, maintain business alignment, and comply with applicable policies, laws, and regulations (and shows the business value of cloud service investments). This GRC approach provides the means to control, monitor, and adapt cloud services (on premises or off).

As highlighted in Figure 8.2, organizations typically carry out GRC tasks largely in a fragmented, siloed manner. This leads to duplicate, stale, and often conflicting information on which decisions are made.

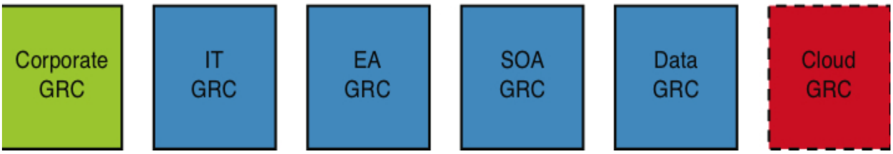

FIGURE 8.2
Cloud GRC must not be executed in a siloed manner.

Cloud GRC requires consistent and transparent decision making, which in turn requires a holistic view of the organization to fully understand risk, effectively monitor compliance, and adjust policies to meet changing requirements, market trends, and regulatory mandates.

With cloud services having such broad implications, it is counterproductive to define yet another fragmented GRC silo.

Figure 8.3 indicates that your cloud GRC approach must support and integrate with your traditional GRC approaches by defining the principles, policies, processes, roles, and infrastructure required to uplift your existing GRC approach.

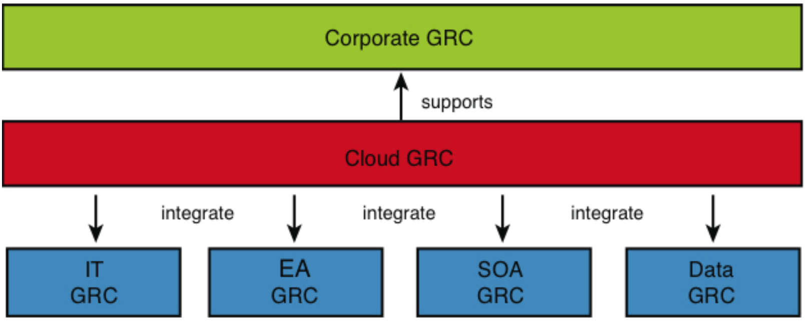

FIGURE 8.3
Cloud GRC requires an integrated approach.

Your cloud GRC must identify the unique nuances of cloud services faced by your organization and provide an approach for addressing these challenges and increasing the efficiency of cloud services.

Consequently, for your cloud GRC to be successful, it must be seen as an extension of your existing formal or informal GRC disciplines, such as information technology (IT), enterprise architecture (EA), and service-oriented architecture (SOA) GRC, as shown in Figure 8.4.

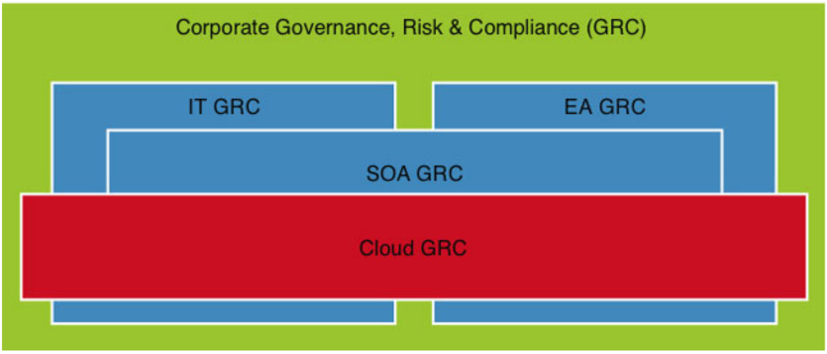

FIGURE 8.4
Cloud GRC relationships with existing GRC disciplines

Focusing purely on cloud GRC and not taking into consideration IT and EA GRC can lead to IT and EA challenges undermining the benefits that cloud GRC is attempting to achieve. Cloud GRC requires governing not only the execution aspects of cloud services but also the strategic planning activities.

Cloud GRC Model

As previously mentioned, security and data privacy have taken center stage in the press. It is important that your cloud governance strategy extend beyond security and privacy of data and include additional cloud governance categories, such as cloud service portfolio governance and the cloud governance requirements that apply to your employees.

Unfortunately, there is no single model of good cloud GRC because each organization has differences and nuances (for example, IT organizational structures, an organization's geographical footprint, the size of organization, and an organizations level of risk tolerance). Despite this, a number of key areas should be considered to frame your cloud GRC planning. Figure 8.5 highlights some key governance areas that your cloud GRC model should address.

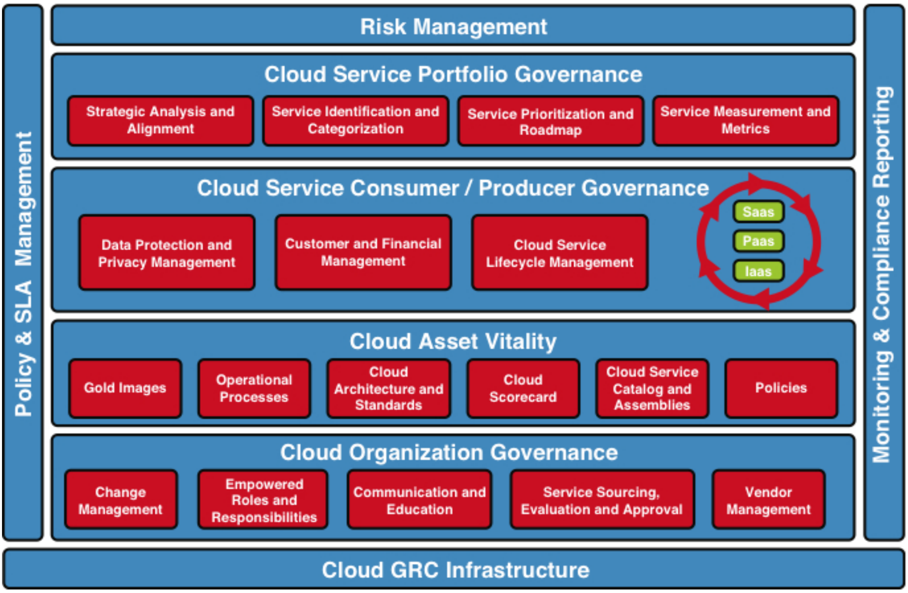
Risk Management
Cloud Service Portfolio Governance
Strategic Analysis and Alignment
Service Identification and Categorization
Service Prioritization and Roadmap
Service Measurement and Metrics
Cloud Service Consumer / Producer Governance
Data Protection and Privacy Management
Customer and Financial Management
Cloud Service Lifecycle Management
Saas
Paas
Iaas
Cloud Asset Vitality
Gold Images
Operational Processes
Cloud Architecture and Standards
Cloud Scorecard
Cloud Service Catalog and Assemblies
Policies
Cloud Organization Governance
Change Management
Empowered Roles and Responsibilities
Communication and Education
Service Sourcing, Evaluation and Approval
Vendor Management
Policy & SLA Management
Monitoring & Compliance Reporting
Cloud GRC Infrastructure

FIGURE 8.5
Cloud GRC model

Cloud Service Portfolio Governance

Figure 8.6 highlights the cloud service portfolio governance aspect of a cloud GRC model that makes sure any investments made in cloud services continue to add value.

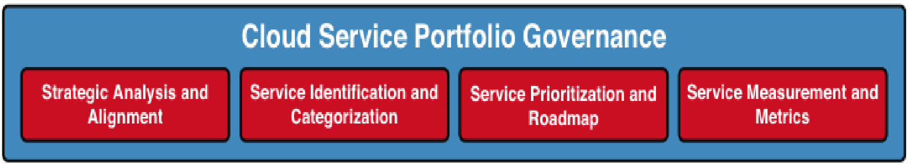
Cloud Service Portfolio Governance
Strategic Analysis and Alignment
Service Identification and Categorization
Service Prioritization and Roadmap
Service Measurement and Metrics

FIGURE 8.6
Cloud service portfolio governance focus areas

This requires that cloud service portfolios be monitored and regulated across several areas of focus, such as cloud service identification, prioritization, sourcing, and approval.

For example, cloud service analysis must be governed to make sure that the appropriate cost-effective services are identified in a timely fashion that align with the goals and priorities of the organization, while taking into account available resources. When identified, cloud services need to be prioritized and sourced from the appropriate cloud service provider, whether that provider is internal IT via a private cloud or a public cloud service provider. Identifying and selecting cloud service candidates should be based on a customized weighted criterion. Example criteria include SLA, cost, risk, complexity, security requirements, and business value.

Cloud Service Consumer/Producer Governance

Figure 8.7 highlights the cloud service consumer/producer governance aspect of a cloud GRC model.

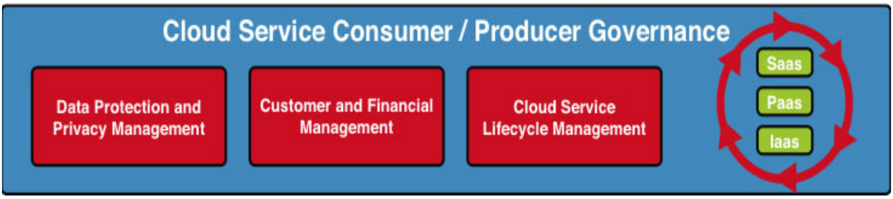
Cloud Service Consumer / Producer Governance
Data Protection and Privacy Management
Customer and Financial Management
Cloud Service Lifecycle Management
Saas
Paas
Iaas

FIGURE 8.7

Cloud service consumer/producer governance focus areas

It is imperative that both the cloud service consumer and cloud service provider be governed, because together they are key stakeholders in the overall lifecycle of a cloud service.

Without cloud GRC policies for cloud service providers, it is common to see cloud service sprawl. In this situation, duplicate cloud services are sourced from different cloud service providers and/or cloud services are built and exist in production but are not being used or managed. This can occur when cloud service consumers are unaware of existing approved cloud services and their associated contracts. Cloud services should be produced that align and support the defined business and IT objectives. Cloud services must be consistent in a manner that complies with the defined policies, standards, and guidelines to promote use and agility. Last, services must meet security requirements such as access control and encryption.

Effective service producer lifecycle governance requires proper policy management and enforcement to ensure that cloud services operate as intended, within parameters. This is critical for visibility into policy compliance and SLA metrics.

When consuming a cloud service, your primary focus is on governance of the operational aspects of the cloud service; this is assuming that the cloud service is approved for consumption in the first place. In additional to the traditional runtime aspect of a cloud service such as SLA definitions, areas such as billing calculation and contract-obligation approvals should be considered.

Cloud Asset Vitality

Figure 8.8 highlights the cloud asset vitality aspect of a cloud GRC model.

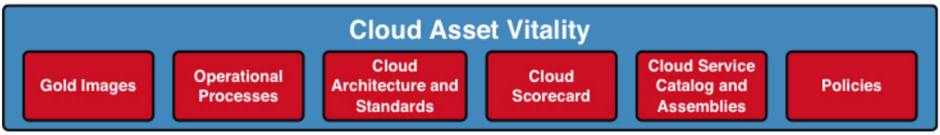

FIGURE 8.8

Cloud asset vitality governance focus areas

Governance models attempt to obtain the highest level of business value from the key assets that an organization has invested in. When talking in the context of cloud governance, key assets include services, solution assemblies, and the cloud reference architecture. These key assets need to be monitored to make sure that their definition and usage still addresses the current requirements and needs of the organization. This requires a strategy called cloud asset vitality, whereby cloud assets are routinely reviewed to make sure they are current, accurate, and most important, relevant.

Cloud Organization Governance

Figure 8.9 highlights the cloud organization governance of a cloud GRC model.

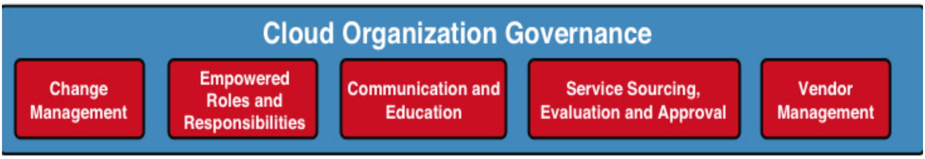
Cloud Organization Governance
Change Management
Empowered Roles and Responsibilities
Communication and Education
Service Sourcing, Evaluation and Approval
Vendor Management

FIGURE 8.9
Cloud organization governance focus areas

Cloud organization governance focuses on the aspects of cloud services that affect your employee/partner roles, including new responsibilities, communication, collaboration, and decision making.

Your organization needs to go through a form of culture adjustment to effectively adopt and execute cloud services. Catering for culture change should not be an afterthought. Instead, an appropriate change management strategy must be utilized to facilitate this adjustment, because culture change is not something that can be achieved overnight, especially when there is a fear that the use of cloud services may diminish employee/partner roles and perhaps risk their employment. Make sure you do not underestimate the impact of culture change. Instead, plan for it. After all, resistance to the change must be accounted for and expected.

Apart from dealing with the changes forced on your employees when consuming/providing cloud services, you need to address the organizational challenges as they relate to relinquishing a level of control when identifying and selecting an external cloud service provider.

Cloud service sourcing deals with identifying an appropriate cloud service provider (when possible) and managing the relationship between the cloud service provider and the consuming organization. This includes areas such as negotiating SLAs and agreeing on the levels of security and assurance of data privacy.

When it comes to cloud service providers, a big part of cloud governance centers on trust as you abdicate responsibility for various operational tasks that you will not have control over. If your organization currently contracts with third-party vendors for outsourcing, you will already have most of the skills and experience required to select and validate cloud service providers. Key areas such as history, reputation, and viability are vitally important, especially in this embryonic stage of cloud service vendors.

GRC Is a Process, Not a Project!

Although cloud GRC is required from day one, your cloud GRC model does not have to be heavyweight. Instead, you should deploy it in an iterative and incremental manner. Doing so enables your cloud GRC model to grow with your cloud initiative.

Cloud GRC must be seen as an ongoing process rather than a one-time project.

Therefore, your cloud GRC strategy must be part of an overall continuous improvement loop (see Figure 8.10) whereby you measure progress and update your cloud GRC model, as necessary, to perform any required course correction.

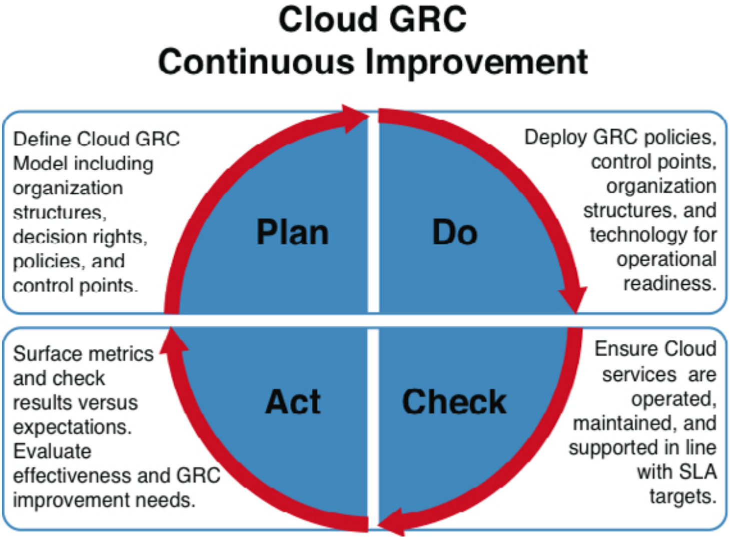

Cloud GRC Continuous Improvement
Define Cloud GRC Model including organization structures, decision rights, policies, and control points.
Plan
Do
Deploy GRC policies, control points, organization structures, and technology for operational readiness.
Surface metrics and check results versus expectations. Evaluate effectiveness and GRC improvement needs.
Act
Check
Ensure Cloud services are operated, maintained, and supported in line with SLA targets.

FIGURE 8.10

Cloud GRC continuous improvement loop

Summary

Initially, you must define and deploy a cloud GRC model that is appropriate for your maturity and your current needs and then evolve the cloud GRC model as your maturity and requirements increase. This should be subsumed as part of your overall cloud strategy roadmap. We cover a cloud strategy roadmap in Chapter 12.

Key points to consider from this chapter are

- The actions of individuals can undermine your cloud service adoption efforts. Therefore define an appropriate cloud GRC model to address your planning, control, and monitoring needs.
- Cloud services present many unique situations for you to address. Therefore, you must consider and weigh up the risks against any monetary gains.

- Effective cloud governance provides organizations with visibility into and oversight of cloud services utilized used across the organization.

Endnotes

[1] http://en.wikipedia.org/wiki/ISO/IEC_27000-series.

CHAPTER 9

Cloud Business Risk and Security

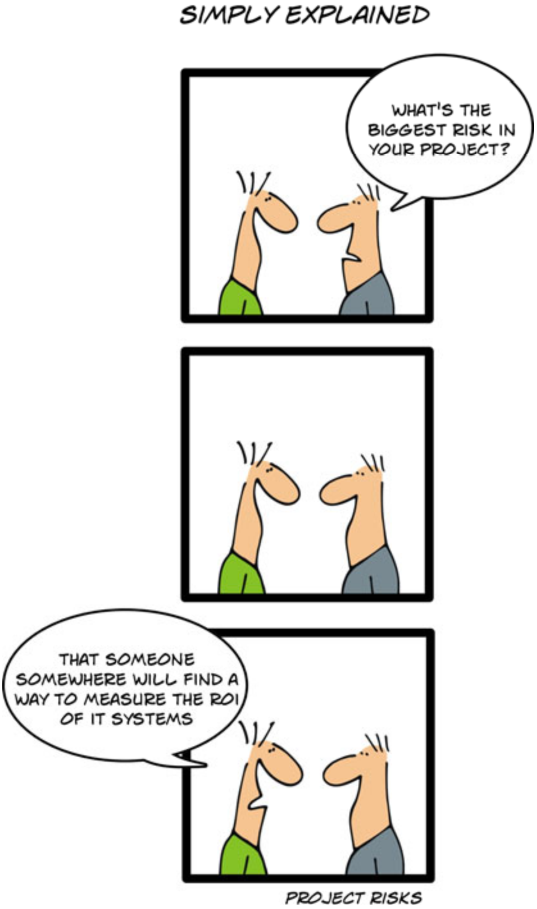
SIMPLY EXPLAINED
WHAT'S THE BIGGEST RISK IN YOUR PROJECT?
THAT SOMEONE SOMEWHERE WILL FIND A WAY TO MEASURE THE ROI OF IT SYSTEMS
PROJECT RISKS

Security is more than technology when considering project risk—Geek and Poke

As identified in Chapter 2, “Evolution or Revolution?,” security is the big concern that most potential consumers of cloud services raise when considering the leap. This is validated in many surveys.

Figure 9.1 shows the same high concern across North America and Europe in Forrester Research’s report “Infrastructure-As-A-Service (IaaS) Clouds Are Local and So Are Their Implications.”

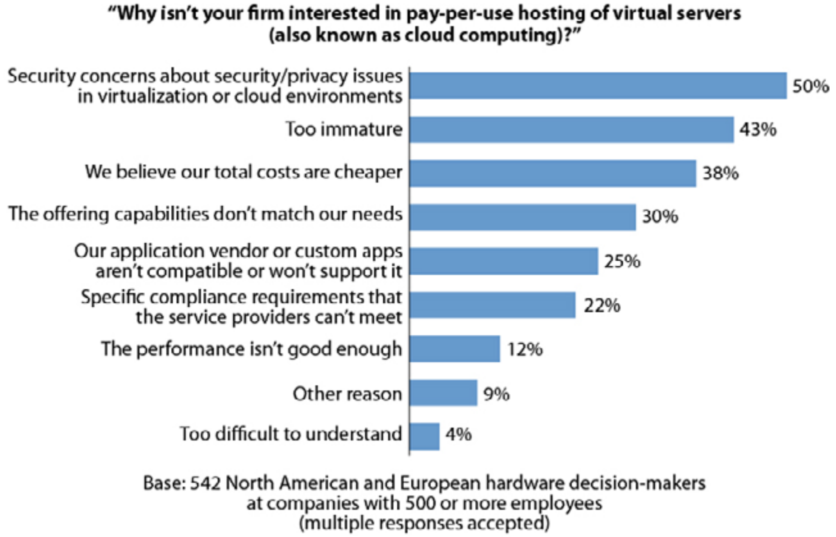

Source: Enterprise And SMB Hardware Survey, North America And Europe, Q3 2009

FIGURE 9.1

Why are we not using cloud services?

Source: Forrester Research, Inc., February 9, 2010

The top-level concerns listed in Figure 9.1 are examined in this chapter.

There are many reasons to consider security and associated business risk as critical in planning to use cloud services. Even so, this chapter does not attempt to define how you secure every type of cloud deployment. Instead, we seek to create a set of guiding principles that align with the governance and operational management goals discussed in this book. Using the commonly identified concerns associated with cloud services today, we offer guidance to best practices and approaches to mitigate them, in the context of a formal security program for your organization.

First, security is not one simple solution! As discussed in Chapter 8, “Cloud Governance, Risk, and Compliance,” GRC is a process, not a project. Security is the same. Although specific efforts can be undertaken to safeguard specific scenarios, the landscape of threats changes constantly, arguably at a faster pace than most aspects of governance, risk, and compliance (GRC). Therefore, the goal is to maintain a security program rather than address security in discrete projects or with individual products.

For cloud services, the word *security* does not actually help us to assess specific concerns and thus do something to mitigate issues (as you read in Chapter 2). Cloud service consumers are generally concerned with their privacy. Organizations are concerned with privacy management. Cloud service consumers, brokers, and even developers are concerned with data loss. Organizations are concerned with data theft. Therefore, to appropriately address risks, knowing the perspective or context of each participant in a cloud service usage model is critical.

This chapter discusses the threat landscape that cloud services add to your current security concerns and summarizes the approaches required to complete a security management program. First, we will look at some risks associated with cloud services, and then we will try to answer the question: Where do I start with cloud security?

The discussion then turns to critically important core areas of a security program that can change when using cloud services. Specifically, this chapter focuses on the following key areas of security management in cloud services:

- Risks associated with cloud services
- Where do I start?
- Framing the security discussion
- Risk and trust
- Data use and classification
- Identity and privacy management
- Infrastructure security
- Legal implications

Risks Associated with Cloud Services

Forrester Research's report "Infrastructure-As-A-Service (IaaS) Clouds Are Local and So Are Their Implications," provides more details on key security concerns with cloud services, as shown in Figure 9.2.

The concerns are of both a technical and business nature. Although it is perhaps obvious that most concerns focus on the ability of the cloud service provider to meet user needs, the reality is that these are key concerns that need to be addressed by the user. (Later in this chapter, we look at the division of responsibility between provider and user and consider the roles of each in terms of understanding and managing security of the cloud service.)

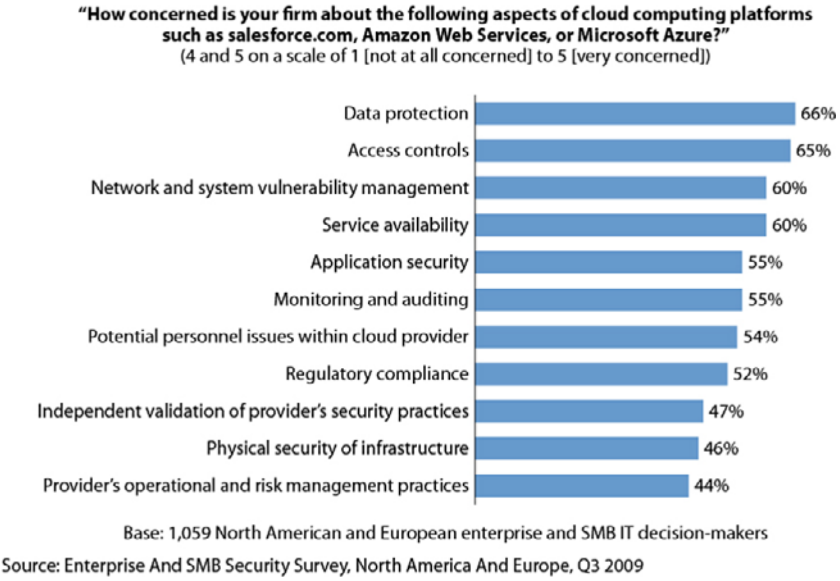

Base: 1,059 North American and European enterprise and SMB IT decision-makers

Source: Enterprise And SMB Security Survey, North America And Europe, Q3 2009

FIGURE 9.2

Top security concerns with cloud services

Source: "Infrastructure-As-A-Service (IaaS) Clouds Are Local And So Are Their Implications", Forrester Research, Inc., February 9, 2010

Looking specifically at the concerns raised in Forrester's report, data protection is the top concern. Data protection dictates a core yet often overlooked discipline for any organization considering the use of cloud services. That is data classification, which is often seen as an arcane and unnecessary exercise, yet it provides clear mapping of organizational data against compliance and operational risk, enabling you to better address that top concern.

Understanding what data you have and the level to which it is critical to your business allows you to better map identity management requirements to your workflows, application deployments, and so forth, addressing the second concern around access controls.

Overall, this list provides a key reference to gauge your concerns against others looking to best use secure cloud services. Although operational and risk management practices appear low on the list, the requirement for such initiatives actually apply and diminish many of the other concerns. Specifically, good management practices and tools help mitigate security risks. Concerns such as network and system vulnerability, availability, monitoring, and auditing all rely on well-defined processes that are defined in and managed by good management tools. If IT is using the right processes and tools, then IT should be better able to support the agility required by the business. If they are not, an organization might suffer in many ways, including from shadow IT.

Chapter 4, "Introduction to Cloud Services," covered shadow IT, which creates potential opportunities for the organization, but more often it creates a serious security risk.

For example, organizations have often faced new technologies being introduced ahead of IT's knowledge. While trying to be autonomous and delivery focused, employees often create security risks as a result. For example, consider the rise in the use of wireless (Wi-Fi) networks since 2000 or so. The lack of security and well-known default settings of common Wi-Fi routers created huge holes in what were previously thought of as secure networks. Even trying to make the networks invisible was a minimal issue to anyone intent on finding opportunities. The terms *war-driving* came about as hackers and their ilk roamed about looking for unsecured or easily compromised Wi-Fi networks, even if they were set up to operate in hidden mode by turning off broadcast mode. Sometimes the infiltrations were inconsequential, but at other times they had devastating effects on the organization (and even beyond, as the hacks spurred both legal and regulatory actions).

In addition, even when IT does roll out such technology, the corporate and IT policies often don't deal with all the critical security risks appropriately. In one famous example, TJ Maxx (TJX) was hacked in 2005. Because of the simplistic security in Wi-Fi routers, Albert "Segvec" Gonzalez worked with a team to crack TJX's Wireless Encryption Protocol (WEP) key and gained access to the corporate network. The hack was not made fully public until 2007, by which time at least 45 million customer credit card numbers were captured along with other personal customer data. During the period of May through August 2008, Gonzalez and ten others were indicted in the TJX case. Gonzalez was involved in the even worse hack of Heartland Payment Processing Systems, which came to light in 2009. In this case, it is estimated that more than 100 million customers were compromised. In this case, despite being required to comply with the Payment Card Industry Data Security Standard (PCI-DSS)—security controls mandated by the major credit card companies—Heartland was still compromised. Quite simply, TJX's security was not adequate. Whether this was a failure in process or implementation of the standard is a frequent point of contention. In this, your organization needs to understand that being compliant does not mean you are secure. To gain compliance and appropriate security requires an adequate security program.

The reality is that many organizations are not adequately securing their environment, and the environment is becoming more complex and threatening. When discussing its report on the "Top Threats to Cyber Security,"[1] the SANS (SysAdmin, Audit, Network, Security) Institute noted:

> Throughout the developed world, governments, defense industries, and companies in finance, power, and telecommunications are increasingly targeted by overlapping surges of cyber attacks from criminals and nation-states seeking economic or military advantage. The number of attacks is now so large and their sophistication so great, that many organizations are having trouble determining which new threats and vulnerabilities pose the greatest risk and how resources should be allocated to ensure that the most probable and damaging attacks are dealt with first. Exacerbating the problem is that most organizations do not have an Internet-wide view of the attacks.

Whether taking a measured step or blind leap into cloud services, most organizations need to improve their security knowledge and related security posture, taking a much broader view of potential attack vectors.

The key points of this discussion are

- Users are one of the biggest security risks you face, today, and tomorrow.
- Shadow IT is an ongoing risk and often introduced by employees who have no concerns beyond their own role in considering the risks of using the solution.
- Experienced teams often roll out new technologies, yet still the risks exist when traditional security practices are ignored or, when required, adapted to the new environment.
- Not implementing a security risk management approach to existing and new technologies will create problems for your organization.
- Attackers will go after things of value, and that is not always the money itself.
- A single security standard is unlikely to save you.

"I rob banks because that's where the money is" is a statement attributed to the American bank robber Willie Sutton. In this day and age, data has value, and that is what our foes are looking for. Thieves are targeting data stores because that's where the "money" is.

Given data has value, the consolidation of multiple organizations' data in a cloud provider's environment makes them a very attractive target for hackers, and their attacks are getting more sophisticated every day. Furthermore, the attacks today are not simple in nature or intention, with crime syndicates, government agencies, and corporate espionage, and more being clear and present dangers, as shown in Figure 9.3.

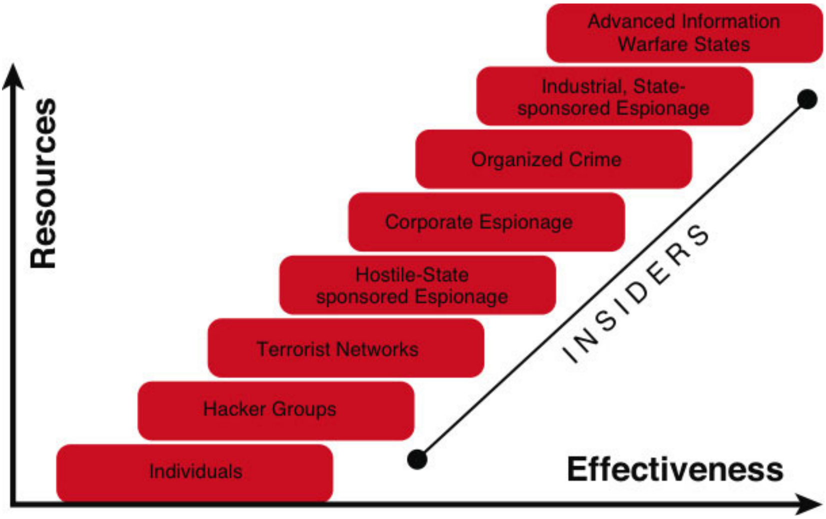

FIGURE 9.3
Threats to business and cloud service providers

Are these threats real? Absolutely!

There are multitudes of security companies out there today, more than 800 at last count. In that count, many pay for exploits and vulnerability details—for example, HP/3Com (Tipping Point), Mozilla Foundation, and Verisign's iDefense. As of May 2010, "iDefense offer as much as $15,000 (US), depending on the nature of the vulnerability, for acceptable well-documented research with reliable proof-of-concept exploit code."[2]

This represents the good guys.

On the other side, you face a fast-growing marketplace for exploits supporting the needs of the bad guys. Are you shopping for those vulnerabilities and exploits? You can join vulnerability sharing circles, or find free sites through ICQ, or look at sites such as Full Disclosure, Metasploit, Packet Storm, and InSecure.org. You can bid for exploits on various private sites dedicated to hacking auctions or even on the very public eBay. Here you can find general exploits, including nefarious hacker tools such as rootkits and rootkit construction kits. Attackers can hire botnets from many of the same sources (cloud service providers), today in the range of $8/system/hour or less. Botnets are massive numbers of PCs that have been compromised by hackers that can be used together to undertake activities under the control of the botnet owner. Activities such as massive SPAM delivery further attacks against corporate systems and more. Today, botnets are responsible for delivery of 85+% of spam (according to J. Green, McAfee Labs), and they have seen a rapid increase in the use of botnets as a means of other attacks requiring scale.

We could use a myriad of security breaches to demonstrate failures in cloud-like solutions. The key goal in this chapter is to provide you with key areas to focus on and approaches that will help you analyze the real risk to you organizations, as well as be able to answer the concerns of your customers, be they enterprises, individuals, or anywhere in between.

Where Do I Start?

First, take note: Although in this chapter we discuss serious risks that need to be considered and addressed in any plan to use cloud services, for a majority of organizations it is possible to find cloud services today that meet or exceed the security requirements they have. The question is more whether you can integrate them into your business processes and overall security program.

Public cloud service providers create cost benefits through scale. They also argue that as a result they can create a highly secure environment in a cost-effective manner, while at the same time maintaining 24x7 security operations.

Security from public cloud service providers spans the gamut from minimal to exceptional, and there are few real measurements you can make to guarantee the real security level cloud providers deliver versus what they offer.

The Cloud Security Alliance

For a look at security in depth for the cloud, the Cloud Security Alliance (CSA) created its security-aligned approach to cloud deployments. The alliance also chose to embrace the bulk of the U.S. National Institute of Standards and Technology (NIST) definition of *cloud computing* we discussed in Chapter 1, "Introduction to Cloud Computing." This is not something we expect small enterprises to really focus on, but reading the guidelines provides some level of understanding about the possible risks in your use.

The CSA is a nonprofit organization formed to promote the use of best practices for providing security assurance within cloud computing, and the alliance provides education on the uses of cloud computing to help secure all other forms of computing. The alliance consists of many subject matter experts from a wide variety disciplines, united in their objectives:

- Promote a common level of understanding between the consumers and providers of cloud computing regarding the necessary security requirements and attestation of assurance.
- Promote independent research into best practices for cloud computing security.
- Launch awareness campaigns and educational programs on the appropriate uses of cloud computing and cloud security solutions.
- Create consensus lists of issues and guidance for cloud security assurance.

The CSA definition of cloud computing as written in the "Cloud Security Alliance Guidance Version 2.1 (2009)"[3] is similar, but importantly provides a breakdown of the Saas-PaaS-IaaS (SPI) model, as shown in Figure 9.4.

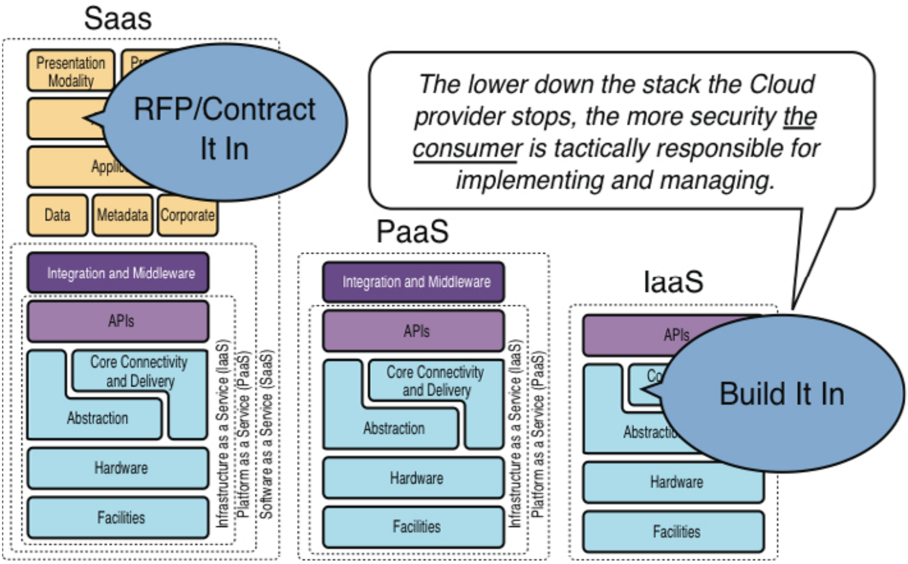
Saas
Presentation Modality
Pr
RFP/Contract It In
Applic
Data
Metadata
Corporate
Integration and Middleware
APIs
Core Connectivity and Delivery
Abstraction
Hardware
Facilities
Infrastructure as a Service (IaaS)
Platform as a Service (PaaS)
Software as a Service (SaaS)
The lower down the stack the Cloud provider stops, the more security the consumer is tactically responsible for implementing and managing.
PaaS
Integration and Middleware
APIs
Core Connectivity and Delivery
Abstraction
Hardware
Facilities
Infrastructure as a Service (IaaS)
Platform as a Service (PaaS)
IaaS
APIs
Co
Abstractio
Build It In
Hardware
Facilities

FIGURE 9.4

Cloud Security Alliance cloud architectures

In Figure 9.4, the CSA also references the level of criticality and responsibility for introducing security into each of the models, wherein a user of IaaS is more responsive for implementing and managing security than a user at the PaaS or SaaS level. The corollary being this: Cloud consumers are more dependent on the cloud service provider for security implementation and management when using a SaaS solution.

The CSA cloud deployment model also provides a clear picture of how each layer can be built upon the core components of the layer below.

Deployment models here suggest for a SaaS vendor to use a third-party IaaS or PaaS provider to build its solution upon. In some cases, it might use more than one. As with many approaches to cloud services, this creates different categories of opportunity and a risk.

The CSA advises the following:

> With so many different cloud deployment options—including the SPI service models (SPI refers to Software as a Service, Platform as a Service, or Infrastructure as a Service, explained in depth in Domain 1); public vs. private deployments, internal vs. external hosting, and various hybrid permutations—no list of security controls can cover all circumstances. As with any security area, organizations should adopt a risk-based approach to moving to the cloud and selecting security options.

This agrees with our thinking and will be a core thesis as we discuss your security strategy for using cloud services. Unfortunately, although security is clearly and constantly identified as the number one concern in adopting cloud solutions, the actual interest and focus levels appear to be incredibly low. Google Trends (again) shows the minimal, almost pathetic, level of queries about cloud security relative to cloud computing itself (see Figure 9.5).

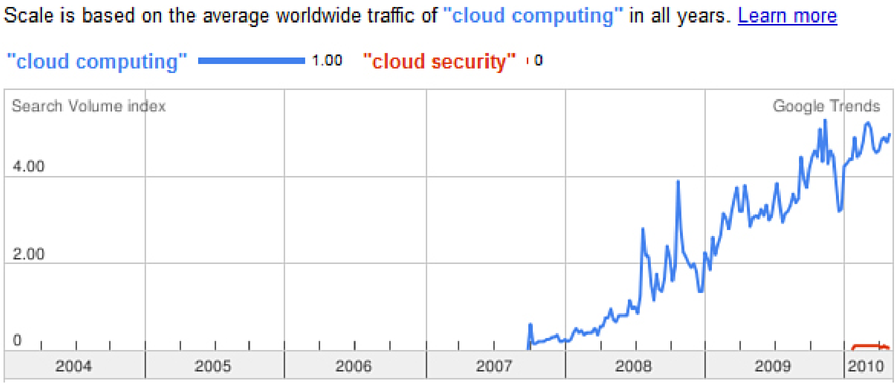

FIGURE 9.5
Google Trends for “cloud computing” versus “cloud security”

Cloud as a concept is massive, and determining the correct security strategy can be daunting as a result. Breaking down the approach to cloud services you take is one way to alleviate the concern, but to counter the security concerns we also need to break down the things we need to focus on. The 13 domains in the CSA guidance document are shown in Figure 9.6.

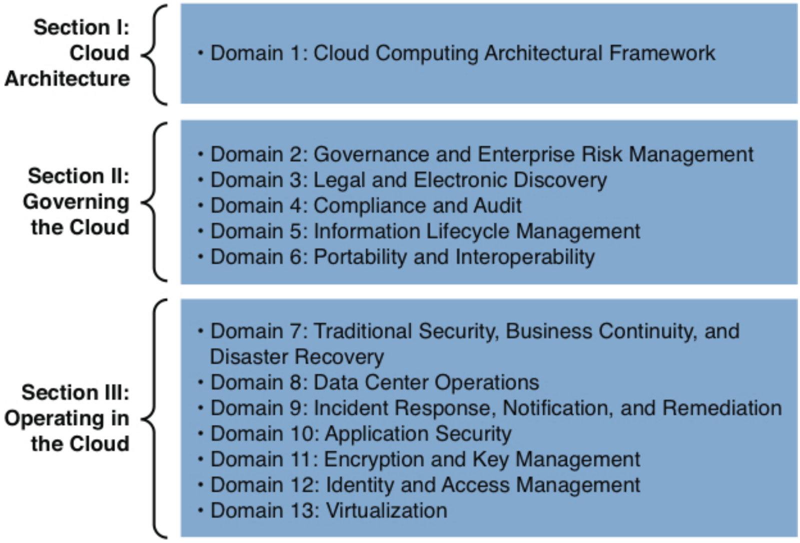

Section I: Cloud Architecture

- Domain 1: Cloud Computing Architectural Framework

Section II: Governing the Cloud

- Domain 2: Governance and Enterprise Risk Management
- Domain 3: Legal and Electronic Discovery
- Domain 4: Compliance and Audit
- Domain 5: Information Lifecycle Management
- Domain 6: Portability and Interoperability

Section III: Operating in the Cloud

- Domain 7: Traditional Security, Business Continuity, and Disaster Recovery
- Domain 8: Data Center Operations
- Domain 9: Incident Response, Notification, and Remediation
- Domain 10: Application Security
- Domain 11: Encryption and Key Management
- Domain 12: Identity and Access Management
- Domain 13: Virtualization

FIGURE 9.6

CSA guidance document domains

Users should find the CSA approach to be valuable, but others exist. The ISO27k standards provide a broad framework that is quite popular in large organizations and that (although not specific to cloud services) is an expansive way to ensure all aspects of a security program are considered. In addition, an ISO 27001 certification provides a better way to assess a cloud provider's security in terms of an information security management system (ISMS) than simple approaches like SAS 70, as we discussed earlier. Importantly, such certification is not a one-time exercise. To maintain a 27001 certification, the organization must both review and monitor the ISMS.

Nine ISO27k standards have been published so far:[4]

- ISO/IEC 27000 Overview & Vocabulary
- ISO/IEC 27001 Management System Spec
- ISO/IEC 27002 Infosec Controls Guide
- ISO/IEC 27003 Implementation Guide
- ISO/IEC 27004 Infosec Metrics
- ISO/IEC 27005 Infosec Risk Management
- ISO/IEC 27006 ISMS Certification Guide
- ISO/IEC 27011 ISMS in Telecomms
- ISO 27799 ISMS in Healthcare

Although the number and breadth of the ISO 27k standards might seem intimidating, the approach is comprehensive and solid. Having been developed over many years ahead of cloud computing, ISO 27k goes well beyond just cloud security concerns. Most large enterprises worldwide use ISO27k as the basis for their security programs. Providers would do well to observe this approach over SAS 70 as a preferred and more useful framework and taxonomy for discussing security with users.

In our discussion on standards bodies, we also looked at European Network and information Security Agency (ENISA), who have created a working group focused on cloud security. The initial investigation created an "in-depth and independent analysis that outlines some of the information security benefits and key security risks of cloud computing." The report is titled the "Cloud Computing Risk Assessment."[5]

The report included a strong recommendation that ENISA create an information assurance framework (IAF).[6] This IAF was designed to

- Allow users to assess the risks in adopting cloud services
- Compare different cloud provider offerings
- Obtain assurance from the selected cloud providers
- Reduce the assurance burden on cloud providers

Last, although cloud services have evolved from many existing technologies and methods for delivering services over the Internet, the experience of consultants and analysts at this time provides a greater level of experience to early adopters. Many consultancy groups and vendors offer cloud service implementation and security assessment services, while at the same time assisting in the development of the standards and frameworks just discussed. The intent is not to sell consultancy services, but more to state that there is a large set of potential risks that many consultancies worldwide have experience in managing, making their involvement to mitigate risks a worthwhile consideration, too.

The key points of this discussion are

- Introduce the criticality of a security program.
- Discuss the common factors and differences of each cloud deployment model that need to be considered in relation to a security program.
- Introduce a framework for considering security in the cloud through the CSA.

The next step is to understand which security solutions are applicable to your cloud service usage. This requires that we consider the architecture of cloud service models and the cloud deployment models,

Framing the Security Discussion

Traditional approaches to security have had to manage only the data center and ensure that any use of third-party services meet the same level of security policies.

Figure 9.7 offers a glimpse of the architectural components that must be considered in a security strategy. This figure shows the environment is already complex. Most organizations, especially those with a great deal of legacy technology or a large number of employees, already face challenges dealing with the exceptional array of disparate technology domains, organizational demands, business processes, and large numbers of people. All those factors make setting and enforcing security policy difficult. Although these environments can align closely to common private cloud architectures, the differences in security requirements when adopting private cloud services relate primarily to the architectural underpinnings used to optimize a data center.

Each layer of the cloud stacks, and end to end these layers offer potential threat vectors for against cloud services. Further, depending on the cloud deployment model used, the responsibility for security and ability of each party to control security for those services change significantly.

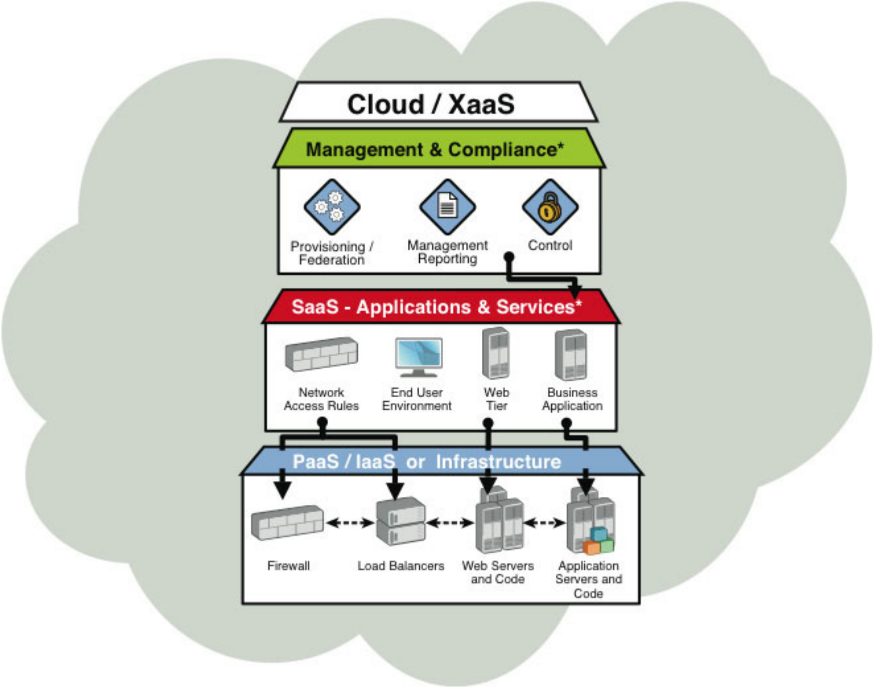

FIGURE 9.7
Managing security for business benefits

Most IaaS implementations use virtualization as the most prevalent technology to support multitenancy capabilities. Hypervisor protection has been the focus for security issues with the cloud since it became popular. With a short list of hypervisor vendors to choose from, there will be a small set of targets for attackers to work against. The terms *red pill* and *blue pill* have been used to describe and demonstrate specific attacks against a virtual machine manager (VMM).

In public cloud solutions, the provider takes responsibility for the security, performance, and availability of the shared infrastructure.

Virtualization can apply to each aspect of cloud computing services as we discussed earlier, and the intent that vendors have is that the core of virtualization solutions, the VMM, will provide *and* control access to shared resources: input and output, compute, memory, storage connections and network, as well as up the stack where possible, in terms of platform or application.

Introducing new service-based technologies introduces new access methods and APIs. Most applications developed today are primarily using web-based UIs and similar technology for API deployments. This capability is increasingly compromised at many different levels of the SPI options, or cloud deployment models. The critical concern is that many see management tools, secure virtual machine managers, and network access controls alongside appropriate isolation architecture as the only new control points that are required. Because APIs and UIs are now a primary vector for attacks, each subsystem for each cloud deployment model needs to be continuously tested. Vendors such as WebSense, Qualys, HP, and others provide tools to undertake automated testing and auditing of various cloud deployment models, but there is arguably no single comprehensive solution available today, so this will require investigation to match your needs.

Using a public cloud service provider creates a challenge because often its own terms and conditions, acceptable use policies, and end-user license agreements deny its customers the right to perform penetration testing. In many cases, such activity will result in being cut off from using the service. So, examine the solution's current process for its own in-house testing and whether it matches your minimal requirements.

In addition, consider the provider's approach to secure to its development. Can it detail a secure software development lifecycle? Does it separate out its testing, quality assurance, and developer environments? Also, can you answer these questions if you are developing your own cloud-based applications? Later in this chapter, we discuss responsibility and ownership of security issues, but for now consider that the architectural changes that result in cloud deployment models also require training for your staff, including developers.

The challenges increase markedly when dealing with public cloud solutions. If you are on the customer side, you now need to match your and the cloud service provider's security requirements and determine gaps. The need is for service and solution integration, accessed from multiple devices and data centers and requiring new management tools and policies. The gaps can be quite significant. If you are a cloud service provider, you need to consider that supporting more complex and demanding customers will require more comprehensive security compliance and certifications.

The key is to factor in the new risks you will face based on the cloud deployment models you plan to use. Table 9.1 presents the primary new threats common to most solutions.

TABLE 9.1 **Primary New Challenges and Threats Common to Most Cloud Solutions**

Challenge or Threat Areas	**Description**
Users	Perhaps the greatest challenge, if not threat to your use of cloud services, will be users. Education is key, but user adoption will depend on your ability to ensure that they are comfortable with the migration and any process changes that result.
Regulations	The complexity and scope of regulations is immense and creates risk with any cloud service provider solution. Involving your legal and human resources groups in the discussions is considered good practice. And if you have other groups dealing with regulatory issues related to processes, such as compliance, their involvement is merited.
Network and communications	Bandwidth availability, communications provider throttling, and issues with using the Internet as you pass through third-party networks (and geographies) create risk. Ensuring the adequate level of service from the providers supporting your use of cloud services is critical to ensuring that your performance and availability concerns are managed.
Identity	Small organizations will find managing identities across and within cloud services an effective way to gain rapid access to new services. Larger organizations that have not successfully implemented identity management will find significant risk in successfully moving to public cloud solutions, especially in terms of the identity lifecycle aspects such as self-service, access controls, authorization and deprovisioning. Smaller organizations are generally more willing to allow their identity data (such as a directory) to reside in the cloud. However, at a certain size, and in specific industries, security and regulatory concerns create the situation where having identity at the cloud service provider is an issue, and the approach shifts to using federation (linking of identity to remote authorization and authentication).
Privacy	You must ensure that the cloud service provider either understands or can meet any privacy management requirements.
Long-term support	The failure of a supplier as a business or to support critical business functionality long term.
Interoperability	Much of IT is about integrating or, minimally, extracting data to be used in other solutions. Interoperability is key to being able to adapt to changes in your business, and if a cloud service provider fails to provide for or advance its support for interoperability, preferably through standards, you will face risks longer term.
Long-term cost	Beware of incremental charges that drive overall costs up, possibly drawing close to or even exceeding the return on investment (ROI) of private solutions.
Support availability	24x7 is no good if it is only email-based.

Organizational change	Education is key to gaining support and success with your employees, customers, partners, and constituents.
Lifecycle management	As the cloud service evolves over time, ongoing assessments need to be made to ensure that both your current governance and compliance requirements continue to be met, security levels and standards are being maintained or improved, and that the cloud providers terms, guarantees, and management processes remain acceptable.

Figure 9.8 shows a number of the changes that occur architecturally when moving to public cloud solutions. Whether you are choosing SaaS, PaaS, or IaaS, the management and compliance aspects are critical to a successful security strategy, too.

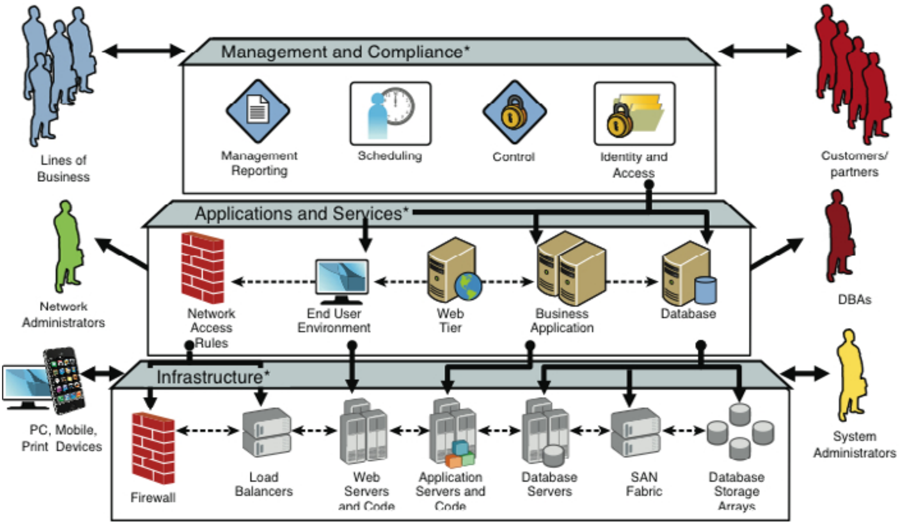

*Across virtual and physical resources

FIGURE 9.8
Managing security across IT and cloud services

Finally, we need to consider the use of hybrid cloud solutions, especially the expected common scenario of public and private cloud use (see Figure 9.9).

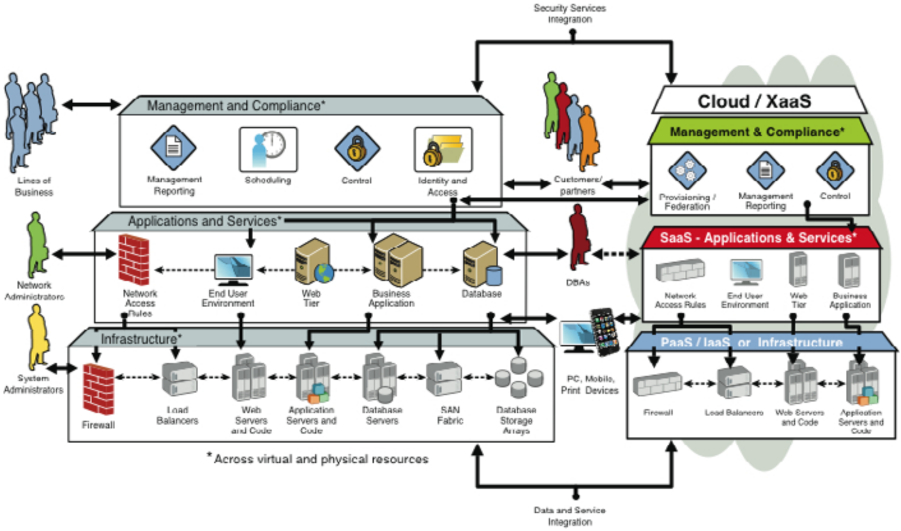

FIGURE 9.9

Combined security for business benefits across the cloud

We can see in this figure several places where security integrations can and should be considered:

1. **Internal data center or private cloud security**—Existing security practices and standards need to be observed. Use of virtualization is generally the primary new technology introduced to create the private cloud solutions, and therefore improving both system and security management tools to recognize this new architecture will be the goal.
2. **Third-party or public cloud security—**These security requirements are focused on the security practices of the actual cloud service provider. Internally, public clouds will be optimized with appropriate levels of multitenant, Internet-based delivery, and internal security, which will need to incorporate standard capabilities such as firewalls, network rules, access controls, and so forth. In addition, cloud service providers must ensure minimally that logical partitions exist between tenants' infrastructure, platform, and software components.
3. **Third-party cloud security**—We'll explain this one with several examples:

   Cloud security services for internal data center or private cloud security—Cloud service providers acting in their primary third-party mode offer security services from the cloud to the IT department's integration into their existing internal data center or private cloud. Examples include antispam/antimalware, data loss prevention, identity management, certificate management providers, network security services, and so forth.

Cloud security services for third-party cloud security—Cloud service providers act as partners or suppliers to other cloud service providers, such that the security is delivered in the same way.

Cloud security services for hybrid clouds—Both general vendors and cloud service providers provide solutions that enable you to manage your deployment across any combination of multiple public, community, and private clouds. For example, many start-ups are offering single sign-on (SSO) between internal IT access control systems, such as Windows Active Directory, and external cloud solutions, such as Salesforce.com or Google Apps.

Data integration requires clear understanding of the value of the data, its use, and its location when in the cloud. Service integration requires clear understanding of the value of the services being supplied, the data being used in or by the service, and the security mechanisms available to you.

There is no denying a lot is at stake.

Being compliant does not mean you are secure!

Being secure does not mean you are compliant!

Either way, when using either private or public cloud services, it is you who remains owner of the data. As a result, in any regulation, you are the responsible party in any theft of that data. Therefore, whether you are a large or small start-up or public sector entity, it is your responsibility to understand whether you are putting your data in the hands of a custodian who understands its role and can meet your compliance requirements or in the hands of a service provider that cannot or will not meet them.

The key points of this discussion are several:

- Introduce the criticality of a security program.
- Discuss the common factors and differences of each cloud deployment model that need to be considered in relation to a security program.
- Introduce a framework for considering security in the cloud through the ISO27k standards, the Cloud Security Alliance, and ENISA.

Risk and Trust

Richard Stallman, founder of GNU (the open source software foundation), spoke to the *Guardian* newspaper in 2008 about cloud computing and stated,[7] "It's a trap…. It's worse than stupidity; it's a marketing-hype campaign."

Stallman's point of contention is more on the business security side of concerns. He offers his theory about what will happen with the following comment: "Somebody is saying this is inevitable—and whenever you hear somebody saying that, it's very likely to be a set of businesses campaigning to make it true." So the situation is, given that we're now two years on, and the momentum for cloud services has done nothing but accelerate, is he correct? He is still to be proven wrong or right, but there is the concern that cloud service providers do not have your best interests at heart:

- Is cloud computing more or less secure than running you own software in-house?
- Do you trust your cloud service provider?

The reality is that cloud service providers are in business as much as anyone else. Their goal is to make a profit, and to do that they will support the needs of their customers in a way that drives that profit. The question then becomes this: Can cloud service providers offer the same or better service at an acceptable price point? If so, what is that price point?

From the security point of view, if you employ your own security professionals, you need to answer this question: Do you know how good your own team is, at every aspect of security versus a third-party (cloud service) option?

All types of organizations, large and small, may have a great security officer or group that takes care of all these risks for them, yet the challenge then is to decide whether a cloud service provider can offer the same level of security

Updates cause issues and are regularly made in instances where security risk exists. This means that when you code to a provider's API, much as integrators have done for many years to Microsoft technologies, Oracle environments, and so forth, those APIs will be deprecated over time, and new APIs will appear.

Similarly, expecting the browser client and associated plug-ins, add-ons, or whatever other integrated capabilities you are relying on to remain static over time is a mistake. Expect change, and given the speed with which change can (and often, *must*) occur in cloud services, you need to hire requisite personnel with appropriate DevOps capabilities if you plan for things to run smoothly. Salesforce.com, for example, provide a regular update window that identifies when it will have its system down for maintenance. A sample from Salesforce.com's summer 2010 schedule is shown in Figure 9.10. Consider that during these times, Salesforce.com may, at its discretion, introduce new functionality, new APIs, or even new security features, any one of which could break one of your critical business processes.

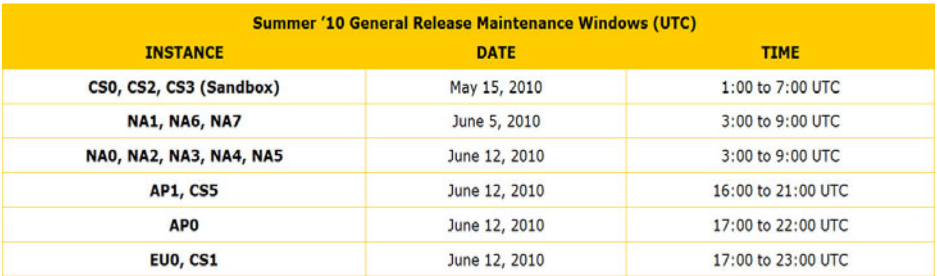

Summer '10 General Release Maintenance Windows (UTC)

INSTANCE	DATE	TIME
CS0, CS2, CS3 (Sandbox)	May 15, 2010	1:00 to 7:00 UTC
NA1, NA6, NA7	June 5, 2010	3:00 to 9:00 UTC
NA0, NA2, NA3, NA4, NA5	June 12, 2010	3:00 to 9:00 UTC
AP1, CS5	June 12, 2010	16:00 to 21:00 UTC
AP0	June 12, 2010	17:00 to 22:00 UTC
EU0, CS1	June 12, 2010	17:00 to 23:00 UTC

FIGURE 9.10

Salesforce.com's maintenance schedule for summer 2010.

Worse, because cloud service providers are not necessarily aligned to your timing requirements, their downtime may occur at a critical time for you, impacting your business ability to serve customers, auditors, or some time-sensitive requirement.

This is part of the control given up when using a cloud service solution. Given that these changes can occur at both irregular and scheduled times, cloud service providers must provide clear guidance for customers. Cloud service consumers must expect and plan for these types of changes, and generally, adapt to the schedule of the cloud service provider. In addition, cloud service consumers must be prepared to utilize any cloud service provider's beta or prerelease testing systems as part of their own change management processes.

So, who is really responsible here?

Figure 9.11 offers a broad outline of key security concerns and associated "responsible" parties. This is a common way to consider focal points and their likely control points when using various cloud deployment options.

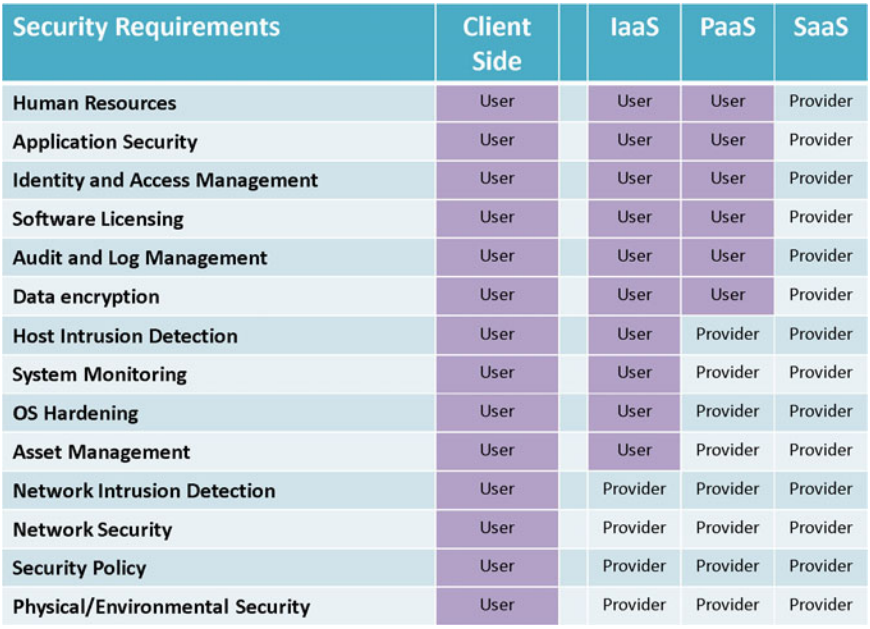

Security Requirements	Client Side	IaaS	PaaS	SaaS
Human Resources	User	User	User	Provider
Application Security	User	User	User	Provider
Identity and Access Management	User	User	User	Provider
Software Licensing	User	User	User	Provider
Audit and Log Management	User	User	User	Provider
Data encryption	User	User	User	Provider
Host Intrusion Detection	User	User	Provider	Provider
System Monitoring	User	User	Provider	Provider
OS Hardening	User	User	Provider	Provider
Asset Management	User	User	Provider	Provider
Network Intrusion Detection	User	Provider	Provider	Provider
Network Security	User	Provider	Provider	Provider
Security Policy	User	Provider	Provider	Provider
Physical/Environmental Security	User	Provider	Provider	Provider

FIGURE 9.11

Common areas of security responsibility in cloud solutions

Figure 9.12 shows where the responsibility lies in terms of control points related to security operations.

To mitigate these situations requires that you undertake the risk analysis considering the levels to which you are prepared to give up that control. In addition, especially for larger organizations, having a clear IT management strategy provides a framework for evaluation.

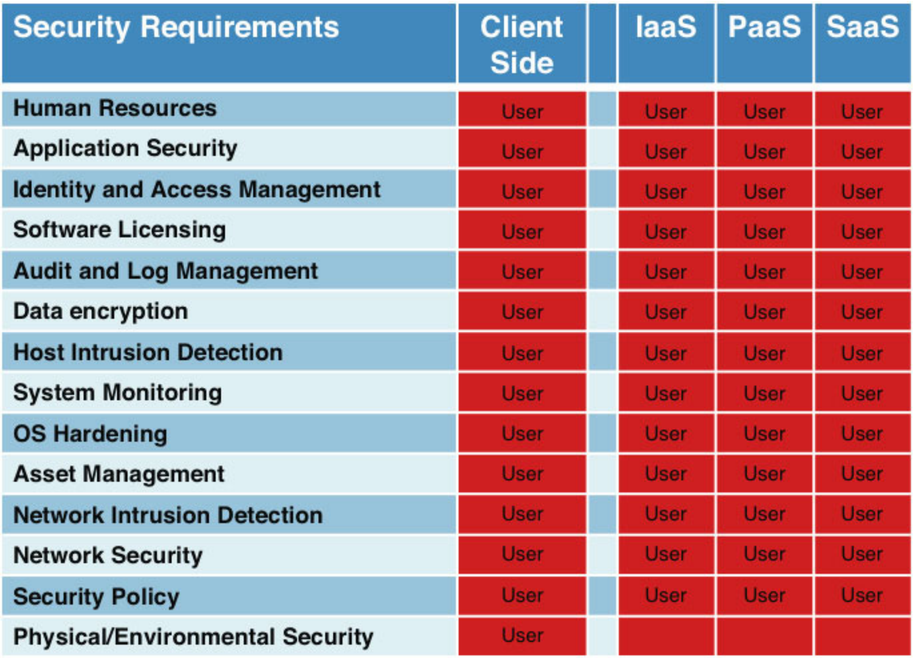

Security Requirements	Client Side	IaaS	PaaS	SaaS
Human Resources	User	User	User	User
Application Security	User	User	User	User
Identity and Access Management	User	User	User	User
Software Licensing	User	User	User	User
Audit and Log Management	User	User	User	User
Data encryption	User	User	User	User
Host Intrusion Detection	User	User	User	User
System Monitoring	User	User	User	User
OS Hardening	User	User	User	User
Asset Management	User	User	User	User
Network Intrusion Detection	User	User	User	User
Network Security	User	User	User	User
Security Policy	User	User	User	User
Physical/Environmental Security	User			

FIGURE 9.12

Actual responsibility and area of initial liability

An Information Technology Information Library (ITIL) strategy would guide you to consider the following:

- Use information security management to define, communicate, and evaluate compliance with security policies. Process access requests using access management.
- Use availability management to ensure proper confidentiality, integrity, and availability of user data, including legislative requirements.
- Use service architecture to apply architectures and standards for security to design a secure environment for the cloud service. Develop the cloud component parts using application development and service engineering.
- Risk management is critical for identifying potential risks inherent with the Internet and users not controlled by corporate policies.

The key points of this discussion are as follows:

- Ensure that you understand that risks exist and are too numerous to detail.
- Detail the level of responsibility you have in understanding and managing the risks.
- Offer an approach to dealing with risk analysis and risk mitigation.

To adequately undertake risk analysis when planning for cloud services requires clear understanding of the data used in the applications and workflows you plan to implement in any type of cloud deployment model.

Data Use and Classification

As we look at cloud services, in particular public cloud services, the consolidation of potentially vast amounts of valuable data will create a logical and attractive target for hackers. Compromising a cloud service provider that provides services for a multitude of customers is a real risk, and therefore before deciding to use a cloud service provider you need to assess its security in relation to your own and create a risk analysis of each option.

Managing compliance and governance in a secure manner while using cloud solutions where the physical location of things, primarily your data, is not necessarily known is quite difficult. You must address data security while being transported to the public cloud and while inside the cloud, and ensure compliance with security and privacy standards. This is job number one. Because many cloud service providers do not adequately understand the data they are dealing with, significant attention is required to ensure that you as a cloud service consumer do not create a compromise. Many regulations detail how certain data needs to be managed—from the European data protection directive, the Payment Card Industry Data Security Standards (PCI-DSS), through to the U.S.-specific regulations concerning healthcare personal information of patients in HIPAA/HITECH legislation. More sample regulations from around the world are shown later in Figure 9.14.

This means the burden, as always, falls on cloud service consumers to ensure they are in compliance with the regulations. All these regulations specifically mandate that contracts with service providers include appropriate controls, processes, and procedures in accordance with each regulation’s guidelines.

To begin, key questions to ask in relation to your data include the following:

- What is our data classification model?
- Are we using the right data classification model?
- Are we aware of all regulations and governance requirements that relate to data that impact our ability to validate compliance?
- Can we ascertain whether any data being considered for use in cloud services has any requirements for protection in relation to compliance?
- Does the cloud service provider offer adequate security and privacy controls to protect our data and meet our compliance needs?
- Does the cloud service provider use any third-party service (e.g., another cloud service) that impacts our ability meet compliance requirements?
- What methods are used to secure our data when in transit, in situ, or in use?

- What level of guarantee is available that the cloud service provider will not sell, misuse, or lose our data?
- How can we get our data out of the provider in the event of a change in relationship?
- Do the cloud service provider's contractual terms match our definition of *data?*

The potential loss of control when using cloud can be considerable, and to ascertain the risk, you need to either understand the scope or blindly accept that risk in the cloud. To achieve a level of understanding, you need a framework to represent your current state scenario and requirements within, and then be able to compare those to a desired state. Whereas utilizing the services of a consultancy that specializes in this type of analysis is one way to approach things, another is to review the publicly available frameworks. Many were already referenced in Chapter 4, including the Cloud Security Alliance (CSA).

As an example, the CSA maps its cloud model through a security control model and links to an actual compliance model defined by the PCI-DSS (see Figure 9.13).

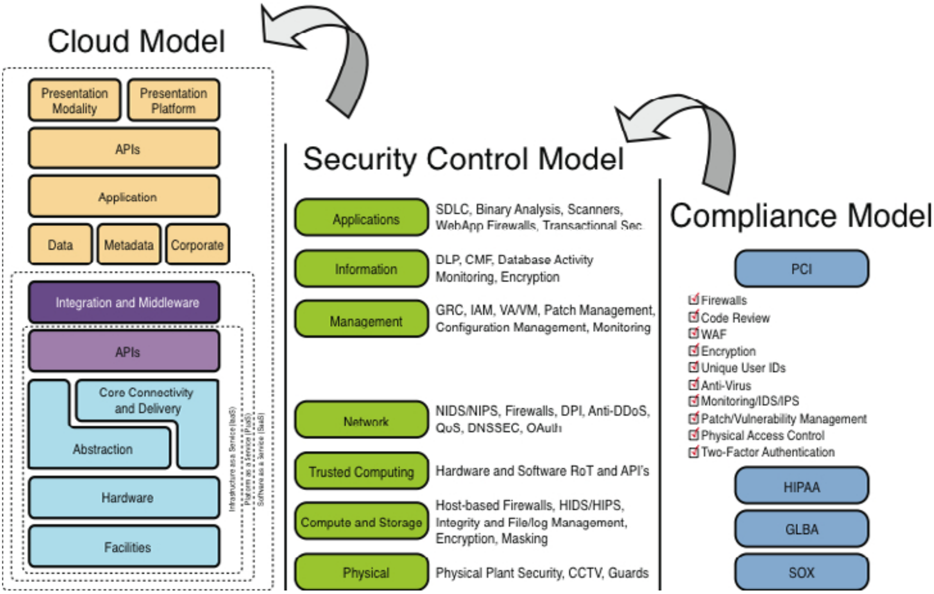

Cloud Model
Presentation Modality
Presentation Platform
APIs
Application
Data
Metadata
Corporate
Integration and Middleware
APIs
Core Connectivity and Delivery
Abstraction
Hardware
Facilities
Infrastructure as a Service (IaaS)
Platform as a Service (PaaS)
Software as a Service (SaaS)
Security Control Model
Applications
SDLC, Binary Analysis, Scanners, WebApp Firewalls, Transactional Sec
Information
DLP, CMF, Database Activity Monitoring, Encryption
Management
GRC, IAM, VA/VM, Patch Management, Configuration Management, Monitoring
Network
NIDS/NIPS, Firewalls, DPI, Anti-DDoS, QoS, DNSSEC, OAuth
Trusted Computing
Hardware and Software RoT and API's
Compute and Storage
Host-based Firewalls, HIDS/HIPS, Integrity and File/log Management, Encryption, Masking
Physical
Physical Plant Security, CCTV, Guards
Compliance Model
PCI
☑ Firewalls
☑ Code Review
☑ WAF
☑ Encryption
☑ Unique User IDs
☑ Anti-Virus
☑ Monitoring/IDS/IPS
☑ Patch/Vulnerability Management
☑ Physical Access Control
☑ Two-Factor Authentication
HIPAA
GLBA
SOX

FIGURE 9.13

CSA mapping from model to compliance

Audits, Logs, eDiscovery, and Digital Forensics

A closely related aspect of data management in relation to cloud is what log and audit data is available, and whether that meets your needs. This is not an issue that resides solely in the IT purview, but, depending on the value of the service and related data used, requires clear guidance from experienced security and even forensics specialists alongside any legal guidance.

On the management side, this data feeds your management and monitoring tools, which also provide an early warning system against breaches. However, the subsequent requirement is to define what minimal data sets will be required should there be a security breach or forensic audit that requires investigation.

Regardless of the type of cloud deployment you use, broader eDiscovery efforts require that data records are able to support the following characteristics based on ISO 15489 “Information and documentation—Records management”:[8]

- **Authenticity**—It is what is says it is—the creator is authenticated and it is time stamped.
- **Reliability**—It is trusted as accurate—facts can be depended on.
- **Integrity**—The records are complete and trusted, and protected and authorized changes are tracked and traceable.
- **Usability**—Records can be found, retrieved, and interpreted, in context of the actual activity and broader activities, and access activities are also logged and maintained.

The requirement to support these aspects of data record requirements applies equally to raw log data.

In terms of private clouds, the need will be to maintain existing security audit standards while increasing monitoring against threats across multitenant architectural elements. Shared storage is a common solution, but VMMs are increasingly a security control point as well as attack vector. In public clouds, can you guarantee that the employees of the cloud service provider have not accessed, modified, or created new data in your space? Can the provider guarantee that its solutions ensure your data does not transit or reside in unacceptable locations? This is a critical part of your security assessment of the cloud service provider’s processes. This includes the requirement to understand how the cloud service provider undertakes its own logging, access controls, and related security monitoring and enforcement.

Maintain vigilance by tracking online sites that allow for independent monitoring and analysis of systems and cloud services, such as cloudfail.org, SANS, and BugTraq among others. Your security program must include third-party references such as these for both independent and up-to-date data on threats.

From an overall perspective, you need to understand whether data is not available, non-verifiable, or verifiable (court ready). Consider the work of the ISO 15489 and the Cloud Audit group as a good start to help guide the efforts around private cloud and the conversation with public cloud service providers.

Encryption

Encryption is often pointed to as a solution for cloud computing security issues. This is a common mistake in understanding both the use of encryption technologies and the cost of implementing and managing them. Although encryption can be used to ensure data at rest or data in motion is logically protected, the fundamental situation exists today where encryption cannot address points of attack:

- Data in use (i.e., usable) must be in decrypted form. Although some research shows that processing encrypted data in the cloud is possible, this is an uncommon approach that requires a specialized architecture and approach to data processing.
- Encryption relies on keys, and a large data set with many users and uses requires a substantial investment in key management. Therefore, key management is the main challenge, rather than encryption, which is widely known and used where possible.

If a cloud service provider states that it offers encryption services, carefully evaluate the approach to encryption that the provider uses. Three key questions exist:

- Does the level of encryption in terms of key size and encryption algorithm meet your baseline requirements?
- Who manages the keys? Where are the keys stored? Who has access to the keys. How can the keys be exposed?
- Does the chosen encryption approach affect your plan to extract your data from the cloud service provider?

In terms of risk, encryption solutions offer several targets that require consideration. The key management system is a target and leads back to understanding how to protect keys and who is responsible for them. The more common target then becomes the applications and services that see unencrypted data. This means that application and service management tools and testing become critical. Question your cloud service provider on the level of testing it implements, and consider your needs to run your own application tests against the entire cloud service solution. This should align with your overall security program and testing requirements.

International Issues

Some definitions of cloud incorporate the concept of location independence. This means that using cloud services in a pure fashion should not require you to know where the processing is taking place, or even where your data is. In its most simplistic sense, that may be okay. But for many organizations, especially those facing audits, the capability to show where their data is and where it is being processed is actually a critical requirement. For example, PCI-DSS audit certification requires that data be processed in a secure facility and that transmission of the data be encrypted at all times. The European Data Directive mandates that personal data of European citizens not leave a country’s borders if it is not compliant with the directive. Many of these requirements actually result in conflict, both internally for you and internationally for governments. The U.S. Patriot Act prompted the Canadian government to not use solutions where resources reside physically in the United States because there were no guarantees of the confidentiality and privacy of Canadian data. In these examples, the scope of the requirements varies, and this is why your use of cloud services will vary by location, data classification, local laws, and industry regulations. Figure 9.14 provides a sampling of international laws that affect data and privacy management.

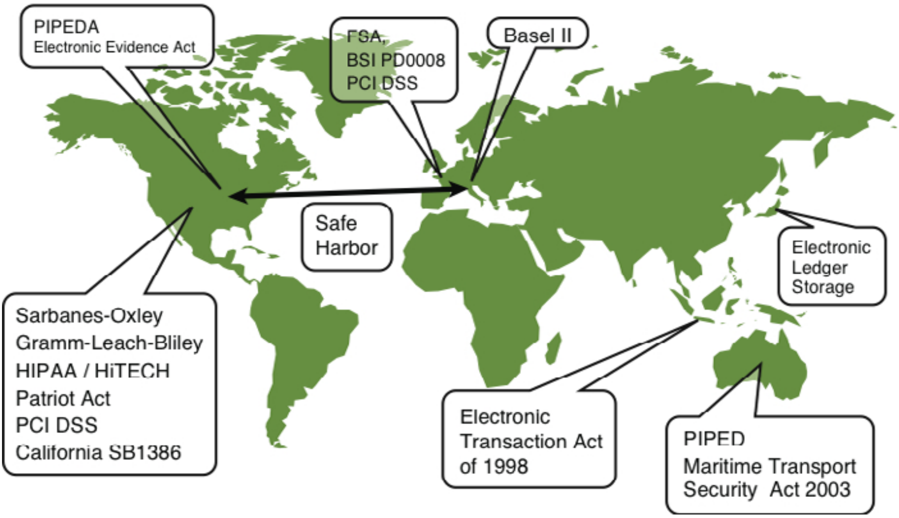

FIGURE 9.14

A sample of international laws impacting data and privacy management

Identity and Privacy Management

The huge amount of data that could end up in cloud service providers' environments closely aligns with the concerns about privacy and the location of PII. In that vein, this section will look at the needs for identity management to set up and control who has access to the information and services and then review privacy management requirements that result.

Identity Management in the Cloud

The entities affected by provisioning processes include the following:

- Employees
- Customers
- Partners
- Services
- Systems

First, let's examine the identity management requirements. Identity management pertains to the following:

- Who has access to what?
- Can users access what they need and no more?
- Are their access rights still valid?
- Who approved their access and when?

- Are their access rights consistent with corporate security policies?
- Do the users for these accounts still exist?
- Is the user information you have accurate and up-to-date?
- Can you prove all this?

For many companies, identity management is already a challenge due to massive or intricate complexity (see Figure 9.15). Organizations are faced with a myriad of accounts, profiles, and passwords and independent administrators and processes all linked through complex dependencies. Putting systems and controls in place to manage the identities and credentials is critical.

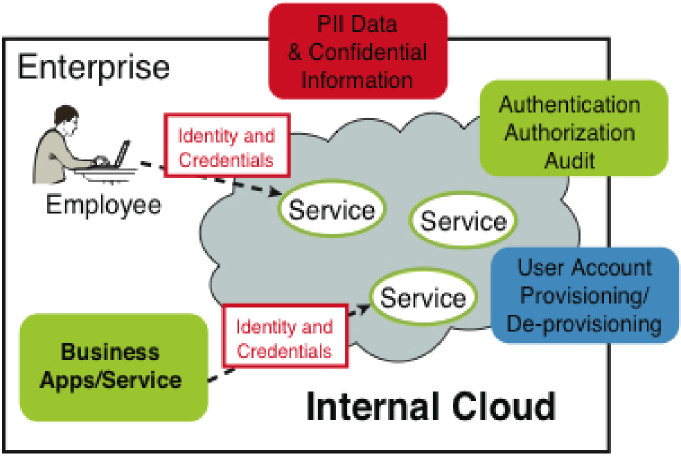

FIGURE 9.15
Enterprise identity

Understanding the number of workflows related to a complete identity lifecycle process that incorporates employees, customers, partners, and all the systems and services they have is challenging. Further, understanding the potential risks associated with PII can create legal challenges for organizations (see Figure 9.16). As discussed earlier, the legal requirements usually are mandated by specific national, industry-specific, or functional regulations, examples of which are shown in Figure 9.14.

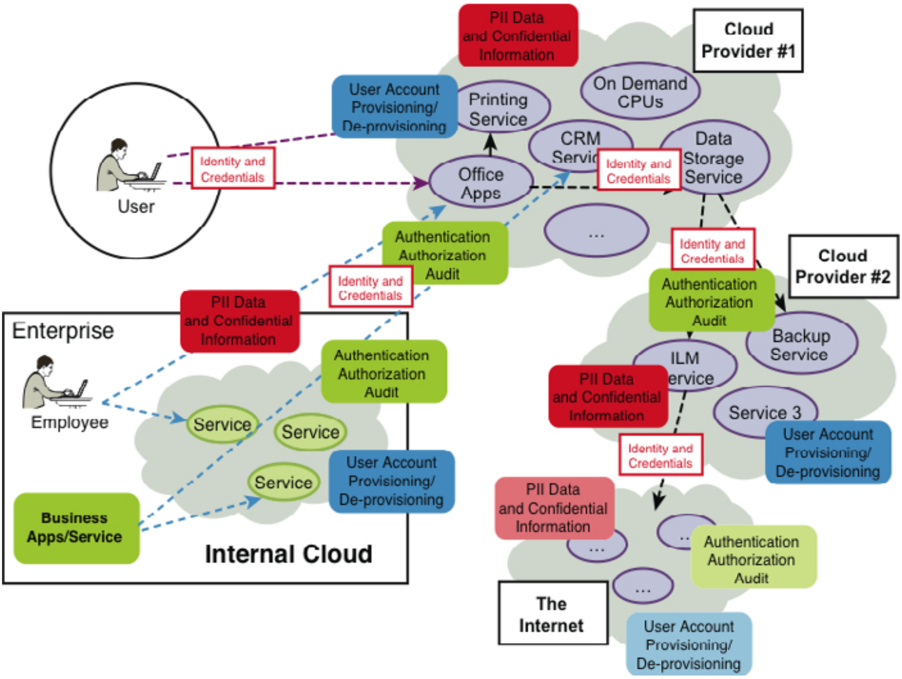
PII Data and Confidential Information
Cloud Provider #1
User Account Provisioning/ De-provisioning
Printing Service
On Demand CPUs
CRM Servic
Data Storage Service
Identity and Credentials
User
Identity and Credentials
Office Apps
Identity and Credentials
...
Authentication Authorization Audit
Identity and Credentials
Cloud Provider #2
PII Data and Confidential Information
Authentication Authorization Audit
Enterprise
Authentication Authorization Audit
Backup Service
ILM ervice
Employee
Service
Service
PII Data and Confidential Information
Service 3
User Account Provisioning/ De-provisioning
Service
User Account Provisioning/ De-provisioning
Identity and Credentials
Business Apps/Service
PII Data and Confidential Information
Internal Cloud
...
...
Authentication Authorization Audit
...
The Internet
User Account Provisioning/ De-provisioning

FIGURE 9.16
Identity across the enterprise and cloud

The impact of putting those processes and data into the cloud can be complex. Regardless of whether you use identity solutions such as federated identity or not, the amount of data that will cross between you and your cloud service provider is far-reaching and potentially creates greater risk for your organization. The use of third parties to manage remote servers has been around for some time. The use of cloud service providers that offer varied and disparate services over the Internet is an evolution of application service providers trying to optimize their offerings for the least cost and resulting security risks, and this means you need to seriously factor in the issue of privacy management in the cloud.

Ultimately, the processes of provisioning through deprovisioning, access control, authorization, audit reporting, and logging become more complex the more cloud service providers are used. Until standards are widely adopted to mitigate some of this complexity, the issue remains, and identity brokers are lining up to solve that problem. Unfortunately, many cloud service providers are still on the fence about adopting the existing standards, or even trying to move them forward because it takes resources away from other priorities. The competitive nature between vendors means that customers must therefore still incorporate identity into their evaluation of a cloud service provider or use some form of identity broker to minimize their identity integration.

Infrastructure Security

IT departments have spent years managing changes in technology approaches and then trying to deal with the failure of those solutions to maintain reliability, availability, and so forth.

Although these are not new requirements, cloud computing offers a chance to manage the issue of fault resolution in a different way. Rather than design to minimize failures, cloud computing vendors are increasingly designing to work around failures. Consider the early days of the Internet, when the Defense Advanced Research Projects Agency (DARPA) goal was to design a self-healing network that would do so by simply rerouting traffic around points of failure. Large-scale cloud service providers are taking a similar approach to the hardware and software stack, capitalizing on the lower costs of hardware and the ability of software to be spread across multiple hosts. Rather than create multiple data centers with multiple copies of data, the approach with cloud computing is to replace a failed part of the infrastructure only when it approaches certain thresholds, and essentially design a solution that has multiple copies of each capability, including data spread across multiple infrastructure components, and if one fails, another simply fills that role.

At the core of the Internet is the Domain Name Service (DNS), which provides the service to look up the addresses of systems around the world using friendly names (e.g., disney.com, times.co.uk, darpa.mil). Attacks using flaws in the implementations of DNS are increasing according to most top-level domain name providers. DNS poisoning is one approach, and it is often part of broader attacks against either providers or users. Attackers might be looking to create a denial of service against a provider or to redirect its traffic to a political website. DNSSEC is an increasingly popular approach used to create a secure Internet transport and should be one of the key longer-terms initiatives alongside IPv6 you consider. DNSSEC is an additional level of security at the core of the Internet that is implemented to ensure that DNS lookups cannot be intercepted and redirected. This type of attack against cloud and web services has seen some traction and should be protected against. A potentially simpler approach here is to manage your own DNS, which makes DNS cache poisoning harder (or at least slower) and ensures that staff can only get to sites authorized by the business. Even though this approach is easy for staff to work around, it provides a greater level of security in the event of a public DNS attack.

A common approach to ensuring only authorized users can access a cloud service is to use IP addressing blocks. This approach is used because it is cheap in terms of implementation time and processing requirements. However, although IP address locking is a simple enough way to ensure that machines from your network can access a cloud service provider’s solution, it is also a simple measure to defeat when an attacker is aware that the method is being used. A better way to provide this controlled delivery is to use virtual private networks (VPNs), which provide both authentication and encryption as one unit.

The need for change management is another essential aspect that needs to be clear across both private and public cloud services. The discussion around trust noted that when dealing with public cloud service providers, the risk exists that the provider might, at its discretion, introduce new functionality, new APIs, or even new security features, any one of which could break a critical business process. As such, the need to incorporate all parts of change from devices used to access the services, through to the cloud services themselves, need to be incorporated into your overall management processes.

Physical Security

Physical security is a consideration when assessing a cloud service provider's ability to prevent attack. Access to its data centers is not only possible via the networks. If an attack originates from inside the data center, the risk of infiltration is increased. As part of your security program, you need to determine to what level the physical security of the service provider requires your attention. This is based on your risk assessment in relation to the value of your data.

Legal Issues

As part of our discussion about risk, we discussed international issues. Such issues closely relate to the overall discussion of legal issues, too. How can any company using public cloud services ensure that it can comply with audit requirements from any regulation that may apply to it? The landscape is extremely broad, complex, and full of risk.

Private clouds create a level of assurance for many organizations in that the environment is mostly under their own control. In the event of an audit, or a requirement for forensic analyses, the bulk of any data, logs, and related information would be within the confines of the organization. Not so with public cloud solutions. The discussion about data location is of particular relevance here because most legal issues will require this understanding. However, IT professionals are rarely lawyers. It is critical at this point to reinforce the need for a change in how IT manages the delivery of business services. In this instance, the need is to change the makeup of the team to include legal representation.

The real challenge? Legal terms are a minefield for the uninitiated. Take the example of Amazon's EC2 SLA (http://aws.amazon.com/ec2-sla/). Running the grand total of 2 pages or 1,000 words, compare that to the Terms of Service at 18 pages, 8,700 words. The SLA is incorporated into the contract by clause 16.5, so it's part of the contract because the SLA is a policy

The TOS includes the following, presented in a box 10 lines high!

"PLEASE READ CAREFULLY—THIS IS A BINDING CONTRACT"

Although there are no major problems with it from a legal point of view, it can be seen as significantly one-sided terms, written in dense legalese. A number of express disclaimers of international sales of goods conventions stand out:

> This Amazon EC2 Service Level Agreement (“SLA”) is a policy governing the use of the Amazon Elastic Compute Cloud (“Amazon EC2”) under the terms of the Amazon Web Services Customer Agreement (the “AWS Agreement”)
>
> 14.2 “The parties expressly exclude application of the United Nations Convention for the International Sale of Goods to this Agreement.”
>
> 16.5. Entire Agreement. This Agreement incorporates by reference all policies and guidelines posted on the AWS Website, including all Additional Policies, and constitutes the entire agreement between you and us regarding the subject matter hereof and supersedes any and all prior or contemporaneous representation, understanding, agreement, or communication between you and us, whether written or oral, regarding such subject matter.

Note that this is a *policy,* not a *guarantee.* The SLA is so vague and full of exceptions that it’s not worth much to the customer, especially compared to the SLAs you will find from full-service outsourcing providers. Of course, comprehensive service levels and guarantees require real money and real penalties. Then, this is the crux of the matter: AWS is a cheap service. At this time, AWS does not refund money; it just provides service credits.

A number of standards efforts are underway to help you deal with the audit and discovery issues you will face. Prime among these is the CloudAudit.

Cloud service consumers expect and demand infrastructure and applications that are reliable and secure, with incidents resolved promptly and problems proactively addressed. Although many cloud providers will not allow you to negotiate their service level agreements, terms of service, or acceptable use policies, consider the following approach based on ITIL models:

- Record and agree on service expectations required and paid for by the business.
- Expectations should correspond exactly to service levels published for each cloud service option listed in a service catalogue.
- Use incident management to coordinate network, server, and application teams for rapid restoration of normal cloud service operation in line with the customer-purchased service level.
- Engage problem management and change management to identify cause and implement a solution that will prevent incident recurrence.
- Use event management to proactively alert IT about failing cloud components as well as provide data trends for proactive problem management.

All previous discussions round out the need for comprehensive and strategic security to be a core component of cloud planning.

Strategic Security

As mentioned in Chapter 8, and earlier in this chapter, ISO27k security standards provide the most comprehensive set of security guidelines for the management of information technology. This is a baseline from which to evaluate your security strategy in the cloud. Figure 9.17 shows the focus areas from an ISO 2001/27002 information security strategy model. The focus here is on information security strategy to support your security program. By mapping these to more focused approaches to security in various disciplines or industry verticals, you obtain some level of completeness in executing that strategy. Thankfully, some agencies and standards bodies do this for you. In the case of the CSA, they provide a controls matrix that maps back to this model.

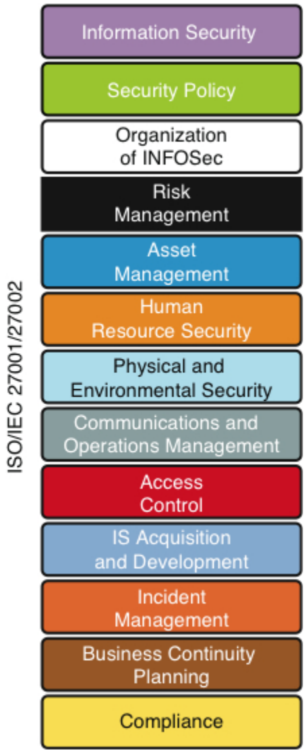

FIGURE 9.17
ISO 2001/27002 information security strategy model

From a small to medium-size enterprise standpoint, you do not need to start looking at cloud with the goal of 27k security certification; however, you should look for cloud service providers that have such certification.

Summary

It is clear that cloud security is a vast and confusing mess. The expanse of cloud service offerings does not facilitate a simple discussion until you focus in on a specific use of cloud services. Hence, in this chapter, we have reviewed the architectural models and security guidelines that are most able to provide you with some direction.

To achieve any level of success, spread awareness in your organizations (not fear). Security is all about managing business risk, so remember the following:

- Risk analysis and threat assessment is the key starting point.
- Match appropriate security/privacy measures to potential threats against your assets.
- Understand and communicate your data classification and data usage.
- There is such a thing as absolute security, but security can be overdone and underdone.
- Compliance does not always equal security.
- Evolve your plan regularly; update to address new vulnerabilities and threats.
- Security = People + Process + Technology.

Endnotes

[1] http://www.sans.org/top-cyber-security-risks/.

[2] http://labs.idefense.com/vcp/.

[3] http://www.cloudsecurityalliance.org/guidance/csaguide.v2.1.pdf.

[4] http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_tc_browse.htm?commid=45306&published=on.

[5] http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport.

[6] http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-information-assurance-framework/at_download/fullReport.

[7] http://www.guardian.co.uk/technology/2008/sep/29/cloud.computing.richard.stallman.

[8] http://www.iso.org/iso/catalogue_detail?csnumber=31908.

CHAPTER 10

Cloud Operations, Administration, and Management

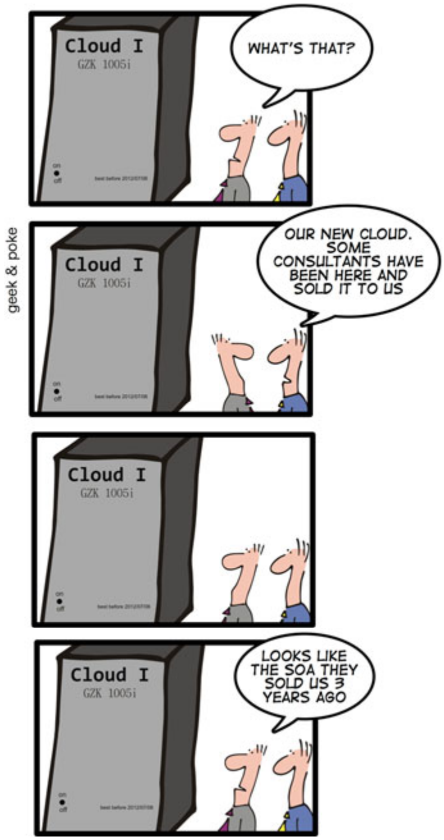
geek & poke
Cloud I
GZK 1005i
on
off
WHAT'S THAT?
OUR NEW CLOUD. SOME CONSULTANTS HAVE BEEN HERE AND SOLD IT TO US
LOOKS LIKE THE SOA THEY SOLD US 3 YEARS AGO

Cloud I—Geek and Poke

This chapter highlights the key aspects that you must be aware of when operating and managing cloud services. This will require a shift in your traditional approach to systems management to cater for a more dynamic environment that empowers end users. In the case of public cloud services, this also means a loss of control in certain aspects of the management capabilities. This will necessitate IT departments to standardize and automate operational processes, and address some unique infrastructure requirements.

Standardized Processes and Tools

Recently, more and more IT departments have undertaken the task of standardizing their IT processes. This can be seen with the recent surge of interest in frameworks such as Information Technology Information Library (ITIL), which documents a number of the key processes used by IT. Whether you consider ITIL a good or bad body of knowledge, it does highlight that IT departments have come to realize the value of standardization and consistency to control costs and to streamline current ad hoc processes. In addition to standardization, the adoption of cloud services will necessitate your movement to a service-centric IT delivery model using standardized processes and enhanced management and monitoring tools. In effect, good IT service management practices are more important in the cloud than before because you have less control, you need to control costs, and you must enable agility.

ITIL has sometimes been hailed as the savior to IT service management, while at the same time hailed as a major barrier to cloud services being adopted and operated in a pragmatic fashion. The truth tends to lie in the middle.

Yes, ITIL is a heavyweight framework. However, as with all frameworks, it must be customized to the needs of your organization. Too many times, ITIL has been taken out-of-the-box and implemented as is and then pushback invariable comes from many quarters of the organization stating that the processes are too complex, not relevant, and overly cumbersome.

You need to utilize IT service management processes that align with your needs and wants for adopting and operating cloud services. Just because your organization is considering utilizing cloud services does not mean that you have to throw away your existing operational processes and best practices. A more pragmatic approach is to understand the differing aspects of operating cloud services and update your processes accordingly. Conversely, if you do not have mature operational processes, utilizing cloud services can exasperate challenges you may already have in operating within your current IT environment today.

Whether you are utilizing ITIL or your own home-grown IT service management processes, you must review these processes through the eye of cloud computing. Given the need for a generally complete set of processes to introduce any new service, having at least a lightweight ITIL-like approach would significantly improve your ability to deliver and manage these new capabilities. For example, your IT service management processes must be automated as much as possible so that it does not hinder the speed at which cloud services can be provisioned. It is no good to have a fully automated cloud service provisioning solution if every cloud service request requires multiple levels of change management meetings before the request can be satisfied.

Companies today traditionally use mature management and monitoring tools for their current IT environments. However, can your existing management tools or cloud providers supply functionality to support your new service delivery model? Does it have support for monitoring and management hybrid clouds from a single management console? Does it have a consolidated dashboard where you can monitor service level agreements (SLAs) for compliance across multivendor cloud deployments?

Currently, many cloud service providers and cloud software vendors struggle to offer this type of functionality, but you must push your cloud providers/vendors to support management functionality such as this. In addition, if you intend to pursue a private cloud deployment, your own cloud service consumers will also ask you these very questions.

Operating Cloud Services

What are some example operational aspects of utilizing cloud services that you should be requesting from your cloud service providers, cloud software vendors, and from your own IT organization?

Figure 10.1 highlights five key areas for which it is important to understand their role on operating cloud services:

- Self-service
- Cloud service management
- Model-driven management
- Policy management
- Metering and billing

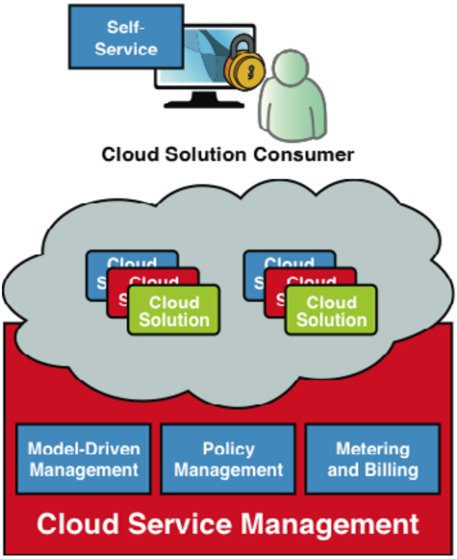

FIGURE 10.1
Managing cloud solutions

Self-Service

A key requirement for operating in a cloud environment is to enable the self-service aspect of cloud computing to allow service consumers the power to request the appropriate cloud solution from an approved cloud solution catalog.

The cloud solution catalog lists the available preapproved and configured cloud solutions with details of the pricing structure and different configuration tiers. The cloud solution catalog can include both private and public cloud solutions. The self-service portal should be able to handle large and complex organization structures and mediate cloud solution requests through role-based access so that users have access only to cloud solutions they are entitled to request. For example, geographically dispersed business units of your organization may have differing governance mandates, which can lead to contrary data privacy and security requirements. This can lead to location-specific cloud solutions available to the relevant users in different parts of the world.

In addition, the self-service portal must enable cloud solution consumers to access billing details, SLAs, service consumption dashboards, and management tasks such as starting/stopping virtual servers, resetting of passwords, and reservation of cloud solutions in advance.

Even though it is advantageous to automate the provisioning of cloud solutions on demand, there will be instances when the cloud solution request will need to verified against a number of predefined policies. (for example, budget approval, service level requested). Although it is advantageous to automate policy enforcement/verification, cloud solution requests will sometimes need some form of manual intervention. When this is the case, make sure your reply is swift and is not bogged down in bureaucracy.

The majority of cloud service providers offer self-service functionality to enable end users to request resources and access management/billing details. However, over time cloud solution catalogs will mature to include hybrid cloud solutions that combine resources from multiple cloud service providers, whether they are all public cloud services or a combination of public and private cloud services. This will complicate and increase the need for self-service frameworks that can abstract the individual self-service portals from the different cloud service providers into a composite self-service portal. Today this can be achieved through the use of middleware products such as portals, process engines, and identity management that interact with the various cloud service self-service functionality through self-service/management APIs.

Cloud Service Management

The management of cloud services compared to traditional system management requires a shift in approach. You need processes and tools that can cope with a more dynamic, service-oriented, and complex environment. Cloud solutions form the consumable cloud service that the end user will utilize (see Figure 10.2). This can be as simple as a single infrastructure as a service (IaaS) cloud service or a complex mixture of multiple IaaS, platform as a service (PaaS), software as a service (SaaS) cloud services grouped together, which in turn utilize the underlying physical/virtual resources.

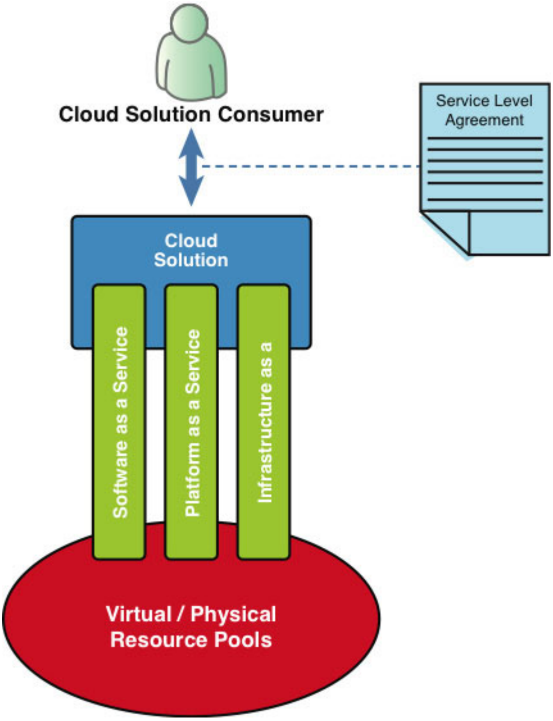

FIGURE 10.2

Flexible cloud solutions

Service Level Agreement

A cloud solution may or may not have an associated SLA. A SLA uses an agreed upon language that forms a contract between the cloud solution provider and cloud solution consumer that enables both parties to agree, to be understood, and to have a consistent approach in the definition and reporting of the quality of a cloud solution. As shown in Chapter 9, "Cloud Business Risk and Security," other contracts that may impact your cloud service usage are the terms of service (TOS) and the acceptable use policies (AUPs).

Many well-known public cloud solution providers do not offer SLAs. Your requirements will dictate whether an SLA is required. Some cloud solution providers offer a limited number of predefined SLAs that assist with the ability for cloud consumers to self-service their requirements. In higher-business-value areas, you will want to negotiate directly with your chosen cloud service provider regarding your exact requirements when defining an appropriate SLA. Do not underestimate the time and effort of this task.

Your SLA must contain a number of attributes that have an established set of metrics that will be monitored and measured to make sure that the cloud solution is meeting its contractual obligations. A key role for cloud service management includes actively monitoring and reporting on service levels against goals over a defined period of time (for example, response time, availability).

Just because you have negotiated an SLA with your cloud services provider does not guarantee the cloud service provider will meet the SLA.

Even with public cloud services that you have subscribed to, make sure that you do not shirk your own responsibilities in monitoring the public cloud services and that they are meeting your SLAs. This can be a challenge because you are sometimes under the whims of your cloud service provider in terms of what monitoring capabilities they make available. Some of the major public cloud service providers offer dashboards that offer basic monitoring information. However, the trend is to offer more of a robust set of monitoring data that can be customized and integrated into your own internal management and monitoring systems.

Figure 10.3 highlights what we believe will be a common scenario for many organizations. An organization’s data and security concerns will initially mandate that data must be hosted internally on a private cloud. At the same time, organizations will want to take advantage of public cloud services and build applications in the public cloud. These applications will need to access and store data that is within an organization’s private cloud. Ignoring any network latency issues, the ultimate cloud solution consumer will not care about the details surrounding the underlying infrastructure.

The cloud solution consumer will want to negotiate and agree on a single SLA with the private cloud owner. Before the private cloud operator can agree on an SLA with the cloud solution consumer, they will have to themselves negotiate with the public cloud provider on an SLA and take this into consideration. The use of dynamic cloud services and resource allocation helps meet the demands of these SLAs.

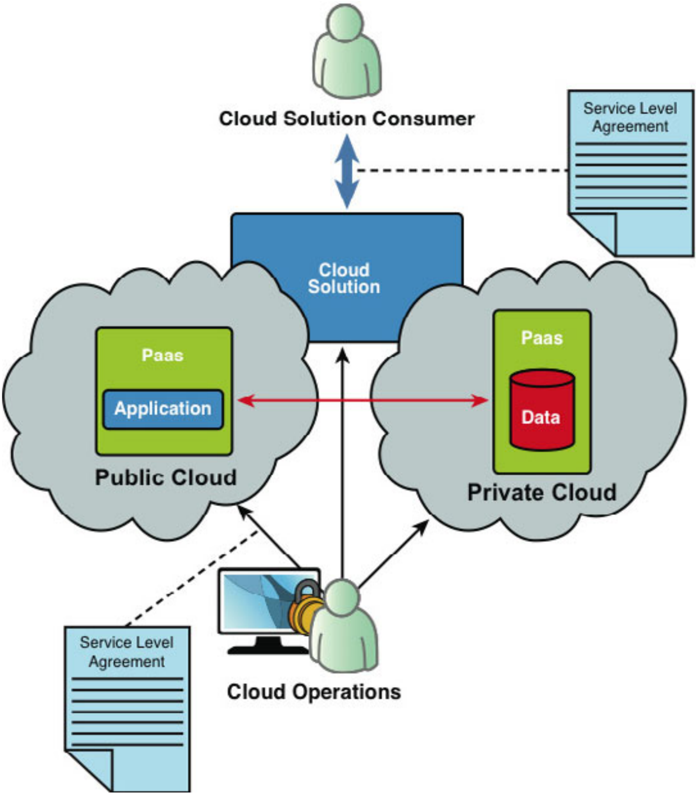

FIGURE 10.3

Flexible cloud solutions require enhanced management tools.

Automation

Even though today's IT operations utilize a level of automation, the optimal use of cloud services requires an increased level of automation. IT operations have long acknowledged the difficulty in deploying and maintaining new software, in provisioning and maintaining new servers with a variety of configurations, and the difficulty in adapting to changes in workload of the environment in a timely and consistent manner.

One of the main challenges when offering cloud services is that as much as possible must be automated. The number of moving parts and relationships within the cloud infrastructure make it unrealistic that human interaction would scale and more likely than not make configuration errors. Without an increased level of automation, your organization will not be able to control the associated costs and obtain additional agility. How else are you going to build so many servers and keep them updated automatically, reliably, and quickly?

Your IT department requires management and monitoring tools that enable the automatic provisioning and configuration of cloud services and resources in both a horizontal and vertical manner while detecting and overcoming human error. Provisioning deals with automation of the installation and configuration of operating systems, infrastructure software, applications, virtual servers, and hosts across different platforms, environments, and locations. Conventional management and monitoring tools do not enable you to increase access to resources and cloud services and automatically provision based on the current demand conditions.

Although virtualization is not a prerequisite for producing cloud services, it can be seen as a cloud automation technology enabler. Whereas in the past virtualization assisted organizations in controlling server sprawl, it now also serves as an enabler for achieving a level of dynamic allocation of resources. This functionality does not come without its challenges and does not preclude the need for management of these virtual resources. In addition, management tools in the virtualization space are still relatively immature. Do not be fooled by cloud service providers who state that by virtue of simply using virtualization you are executing a cloud strategy.

Even though it is possible to achieve dynamic allocation of resources utilizing physical resources, it can be more daunting and less flexible than virtual resources. As the market matures, expect to see more intelligence built in to the physical resources themselves (e.g., servers, storage, network) that will make the management of these resources in a cloud environment easier and more flexible.

Model-Driven Management

Although cloud solutions are made up from one or more cloud services, it may be more complex than you might imagine. Cloud solutions can be formed using a mixture of cloud services via different delivery and deployment models and therefore can have interdependencies crossing corporate and public boundaries. In addition, these cloud services make use of a number of dynamic physical and virtual resources.

This highlights that a large number of cloud solutions, cloud services, and resources need to be monitored and managed. In addition, the static and dynamic relationships between cloud solutions, cloud services, and resources are much too complex to track manually. Conventional management approaches and tools are ill equipped to address this challenge. Without access to information concerning these dynamic interdependencies, diagnosing problems and correlating problems in a complex, distributed environment is a huge challenge.

Model-driven management tools fully understand the configurations, relationships, and interactions between cloud solutions, cloud services, and virtual and physical resources.

In addition, the underlying model automatically maintains these dynamic relationships, including the additions and deletion of cloud services and physical/virtual resources. This enhances an organization's ability to perform monitoring and diagnostic analysis and to manage service levels.

Although using automated preapproved cloud solutions should remove the possibility of inconsistency and incompatibility, it is still prudent to have an approach to monitor configuration updates in case of human error. If this possible challenge goes unattended, it may lead to configuration drifts and security vulnerabilities that lead to lack of policy compliance.

After the cloud solutions and cloud services have been provisioned, it is important that their configurations be monitored. Real-time detection of updates to the configurations captures what has changed, when it changed, and who changed the configuration. This proactive approach to configuration monitoring enables a full configuration change history. Any updates to the configuration information can be compared either against a reference configuration set or against previously saved configuration snapshots. This approach enables an administrator to see the drift in configuration and track policy compliance over time. If a cloud service falls out of compliance, administrators can optionally define corrective action to bring them back into compliance.

Policy Management

Policy management is the demonstration of, and enforcement to, regulatory standards, industry standards, and internal best practices. To have your cloud services run effectively, it must adhere to these agreed upon standards and policies that promote efficient operating procedures. Example policy categories include the following:

- Security policies
- Regulatory policies
- Resource scheduling policies
- Configuration policies
- Quality of service (QoS) policies

After you have agreed, defined, and applied your policies, you must test for compliance. Compliance is assessed by way of defining policies that provide rules against which cloud services are evaluated, and where a corrective action may be taken.

Even though policy compliance is part of an overall policy management approach, it should also be seen as a cross-cutting concern. The enforcement of policies can cross many disciplines depending on what policies categories you define. Therefore, the enforcement of these policies is achieved using many different technologies. There will be times when it is not possible to automate the enforcement of every policy and will require human involvement.

Metering and Billing

Supporting provisioning, control costs, meet your SLA obligations, and define chargeback schemes requires the cloud infrastructure to have a comprehensive monitoring and metering system. Even if you currently have no plans of billing your private cloud consumers, metering is still a valuable endeavor in demonstrating resource consumption. In addition, it also highlights how your organization is using cloud computing and the value that each cloud service consumer is receiving.

Automated billing capabilities of the cloud infrastructure need to offer and support the new cloud service pricing models, whether that is via flat-rate or usage-based billing. However, this can be challenging due to the cloud infrastructure’s dynamic nature and the complexity to track usage based on resource consumption by each cloud consumer. This requires metering the appropriate metrics at the appropriate locations within the cloud infrastructure. This is another area in which model-driven tools come to the rescue.

It is important that billing and consumption details do not surprise your cloud service consumers. Rather than the consumer waiting for the monthly bill to arrive, a more real-time billing inquiry and analysis facility should be offered. Consumption and billing details should be provided via two means: an easy-to-navigate dashboard, and a billing API so that the billing metrics can be integrated with existing reporting and analysis systems.

If cloud service consumers can see their usage in real time, they can make analytical decisions about controlling their consumption accordingly. Real-time analytics include metrics such as service quality, identification of cloud service consumer, and patterns of usage over time. This assists organizations so that bills for excessive cloud service consumption will not surprise them.

Cloud service consumers attempting to identify an appropriate cloud service provider have a challenging time when trying to compare costs across multiple cloud service providers who offer the same service. For example, there are currently no real standards for measuring a cloud virtual server QoS, making cost comparisons challenging.

Cloud Service Provider Management

With the current maturity of the cloud marketplace, you should not be surprised if large majorities of cloud service providers disappear in the future. As highlighted in Chapter 8, “Cloud Governance, Risk, and Compliance,” when identifying a provider, review key areas such as their history, reputation, and viability.

Although you can take several steps to minimize the risk of your cloud service provider disappearing, other factors may still influence your need to move away from your current cloud service provider. Those factors include the following:

- Increasing cloud service usage costs
- Cloud service provider acquisition
- Decreasing QoS
- Transitioning functionality in-house

Transition to another cloud service provider requires the ability to move your cloud workload: data, configurations, application functionality.

Access to your data can be a major inhibitor to transition away from your current cloud service provider. Depending on the cloud services consumed and the amount of data that they have generated, you might be dealing with terabytes worth of information. This can lead to issues if you have to transmit this information over the Internet in bulk.

When negotiating a contract with your cloud service provider, consider the inclusion of data access via hard disk transportation.

Currently, most SaaS providers do not adhere to or support the export of data in a standardized format. This places the emphasis on you to extract, transform, and load the data into the new SaaS provider. Although this situation is not unique to cloud services, it is something that you should be aware of and consider.

Consider utilizing cloud services from providers that support standardized approaches to assist in the migration of your cloud workload, such as configurations, data, and application functionality. It is still early in the maturity of cloud services to have access to a single all-encompassing standard and toolset to address the transition of all types of cloud services from one cloud service provider to another. However, this should not stop you from identifying and requesting these requirements from your cloud service providers in a best effort to minimize future risk. For example, Open Virtualization Format (OVF) from the Distributed Management Task Force (DMTF) enables you to package together cloud services along with the appropriate metadata and then deploy it to another cloud service provider that supports OVF. OVF is supported by most large vendors and minimizes the time it would take to transition your workload to another provider.

Summary

Your approach and operational requirements will be largely dictated by the cloud services that you decide to make available and consume and the cloud service providers you choose. With cloud services being at such an early level of maturity, you will be faced with fluctuating risk, fluidity of operational requirements, and cloud service provider viability. So, you need to be aware of these risks, deploy standardized processes and tools, and have an approach to transition from a cloud service provider should the need arise.

Key points to consider from this chapter are

- Cloud services require a shift in your traditional approach to systems management to cater for a more dynamic and complex infrastructure environment that empowers end users.
- The adoption of cloud services will necessitate your movement to a service-centric IT delivery model using standardized processes and enhanced management and monitoring tools.
- The static and dynamic relationships between cloud solutions, cloud services, and resources are much too complex to track manually.
- Just because you have negotiated a SLA with your cloud services provider does not guarantee the cloud service provider will meet the SLA.

CHAPTER 11

Reflections on Life in the Cloud

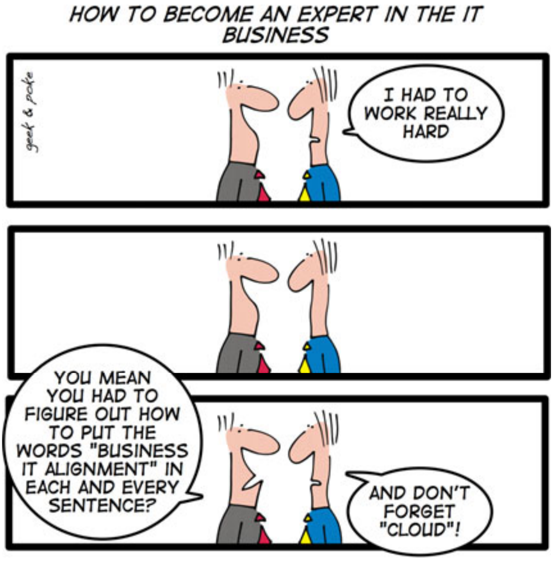
HOW TO BECOME AN EXPERT IN THE IT BUSINESS
geek & poke
I HAD TO WORK REALLY HARD
YOU MEAN YOU HAD TO FIGURE OUT HOW TO PUT THE WORDS "BUSINESS IT ALIGNMENT" IN EACH AND EVERY SENTENCE?
AND DON'T FORGET "CLOUD"!

How to become and expert in the IT business—Geek and Poke

The focus of Part III, “Life in the Cloud—Planning and Managing the Cloud,” was to highlight the different aspects and areas for consideration when planning your cloud strategy. The different aspects highlight that the benefits of cloud services do not come automatically and emphasize that simply deploying a cloud hardware/software product is only half the solution.

The IT industry has a history of persuading enterprises that a pure technological solution is all that you need. More recently, we have seen “service-oriented architecture (SOA) in a box,” and we are already seeing multiple vendors pushing “cloud in a box.” Although these products have a place, do not underestimated the time and effort needed to apply cloud services within your organization strategically. Otherwise, any gains garnered from cloud services will be wasted with duplicate efforts, out-of-control costs, and delays in project deployments.

Therefore, be wary of any cloud hardware/software vendor who says that he can deliver a cloud-in-a-box solution without having supporting strategic planning artifacts such as cloud maturity models, cloud reference architectures, cloud security framework, and cloud governance models.

Key takeaway points include the following:

- Garner support from senior executives and division leaders to assist with identifying, understanding, and prioritizing your business and IT reasons for adopting cloud services. You will also need this support when addressing risk management and culture issues. Do not underestimate the impact cloud services will have on your culture, and make sure you cater for this fear of change. Resistance to the change must be accounted for and expected.
- Use a risk management approach to identify your cloud service providers. Require your cloud service provider to comply with applicable data protection and privacy laws.
- You are required to understand where your data resides, how it is protected, and where it has been transmitted. You can transfer responsibility, but you cannot transfer accountability. Security and data policies are imperative, but also make sure you have a comprehensive cloud governance model that caters for more than just security and data policies.
- The adoption of cloud services necessitate your movement to a service-centric IT delivery model. Although this might seem daunting, many IT organizations are moving or want to move to a service-centric IT delivery model, and cloud service adoption is merely accelerating this. Automate and standardize your existing or future planned IT service management processes.
- Understand your cloud management requirements to drive your overall cloud infrastructure design. Investigate model-driven management tools to assist you in configuration standardization and the managing/monitoring of the static and dynamic relations between cloud services and the underlying resources.

Although all the perspectives and planning disciplines highlighted in Part III might seem daunting, you do not have to address them all at once. Whereas it is important for you to be aware of all these perspectives, just deploy the solutions or activities required to meet the goals of each roadmap phase.

As with any planning activity, it must not be seen as a one-time effort, especially with the relatively young maturity of the current cloud marketplace. Therefore, planning should be seen as a continuous effort in which you tweak your assumptions, adjust your thinking, cater for changes in the cloud marketplace, and execute the appropriate amount of your plan to support your current needs.

PART IV

GPS to the Cloud… Where to Now?

IN THIS PART

CHAPTER 12

Creating a Successful Cloud Roadmap

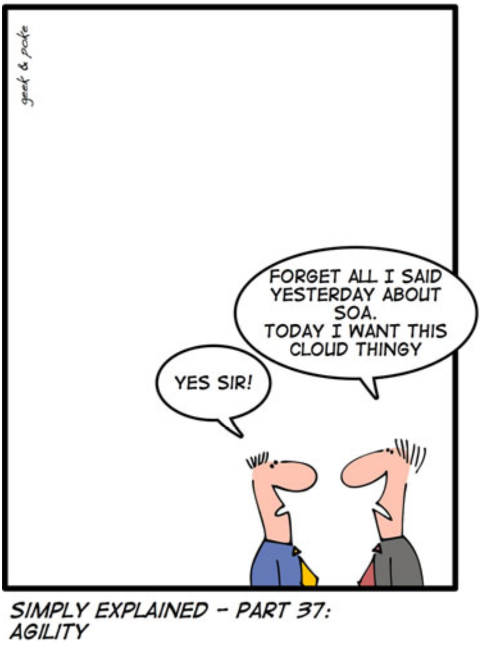
geek & poke
FORGET ALL I SAID YESTERDAY ABOUT SOA. TODAY I WANT THIS CLOUD THINGY
YES SIR!
SIMPLY EXPLAINED - PART 37: AGILITY

The end goal is a moving target—Geek and Poke

No matter what type of cloud services or deployment models you are considering as part of your overall IT strategy, you must have a cloud services adoption roadmap to guide your journey.

A cloud services adoption roadmap provides guidance that enables multiple projects to progress in parallel yet remain coordinated and ultimately result in a common end goal. The cloud services adoption roadmap consists of program-level efforts and a portfolio of cloud services. The program-level effort creates strategic assets such as the cloud architecture, cloud infrastructure, cloud governance, risk, and compliance (GRC) processes, and security policies that are leveraged across all the individual projects.

The program-level efforts provide and enforce the necessary consistency required to succeed at cloud service adoption. A delicate balance needs to be struck between too little control and too much control. With too little control, cloud services adoption will be haphazard at best. Too much control may stifle project teams, resulting in pushback or, in the worst case, outright defiance.

Initial projects drive the cloud infrastructure build-out and identify the initial cloud services. Follow-on projects leverage the cloud infrastructure and use the cloud services. Obviously, it is the follow-on projects that demonstrate the full value of adopting cloud services.

Without an adoption roadmap, there is a higher chance that you will encounter challenges that may undermine your organization's effort to realize the benefits of cloud services. This might lead to circumstances in which you are not making the most of the cost savings and agility improvements that you could have acquired.

Be aware that your initial cloud services adoption roadmap will be based on several assumptions—due to the cloud services market being in its infancy and your understanding of the affect cloud services will have on the culture of your IT organization. Do not be afraid to revise your cloud services adoption roadmap as your organization and the cloud marketplace matures. Make sure that any updates to the roadmap are executed in an expeditious manner to ensure that they have the best chance of addressing any unforeseen challenges and risks without delaying your overall program. The rest of this chapter focuses on a number of organization characteristics that will affect the way you address and plan your cloud services adoption roadmap.

Crossing Your Chasms

Geoffrey Moore's book *Crossing the Chasm* used the 1957 technology adoption model and highlighted the different organization types (innovators, early adopters, early majority, late majority, and laggards) and that technology vendors had a significant barrier to overcome between the early adopters and the early majority.

Moore's initial model recognized that organizations fall across the spectrum, but he focused on one chasm. Today, there is much lower resistance to technology adoption because of the Internet, so the gap and potential time lag between the early adopters and the early majority is now much less. In addition, today's organizations can fall across more than one category, depending on their size, structure, and appetite for risk.

We have used the original technology adoption model and identified a number of chasms that organizations may encounter, depending on their organizational type and culture to adopting cloud services. Our goal here is to help identify each chasm and the benefits and risks associated with crossing each chasm. You can then use these chasms to identify which organization types describe your organization and to start molding your own cloud services adoption route.

Use Figure 12.1 to identify the types of characteristics that you associate with your organization, which in turn will assist with the type of chasms you have encountered or will encounter when producing or consuming cloud services. Without addressing these chasms, cloud services will either lose momentum or run wild in your organization. As previously stated, your organization may exhibit multiple characteristics.

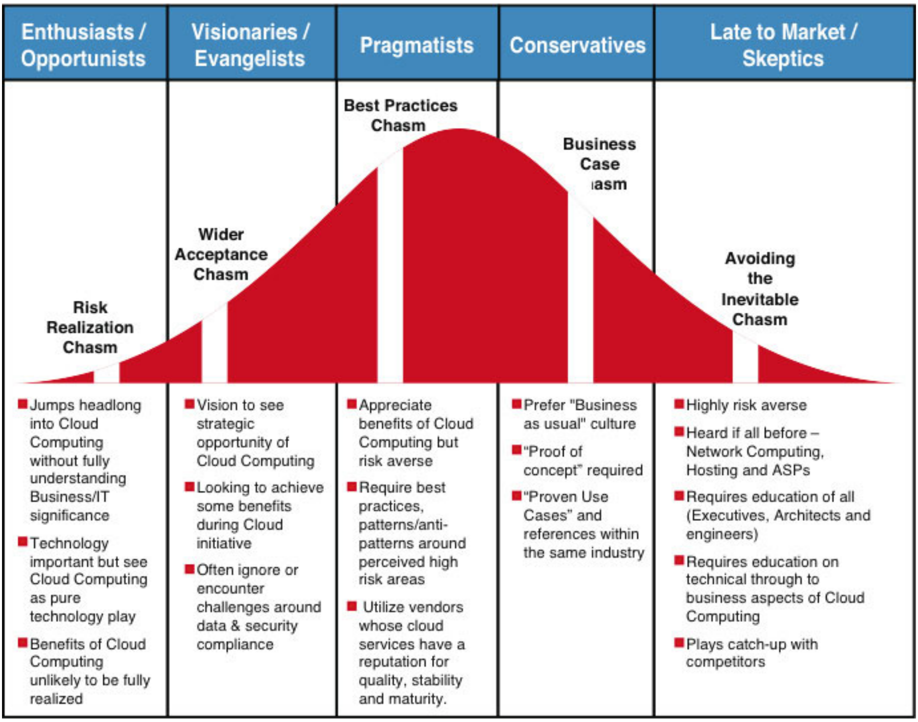

Enthusiasts / Opportunists	Visionaries / Evangelists	Pragmatists	Conservatives	Late to Market / Skeptics
■ Jumps headlong into Cloud Computing without fully understanding Business/IT significance ■ Technology important but see Cloud Computing as pure technology play ■ Benefits of Cloud Computing unlikely to be fully realized	■ Vision to see strategic opportunity of Cloud Computing ■ Looking to achieve some benefits during Cloud initiative ■ Often ignore or encounter challenges around data & security compliance	■ Appreciate benefits of Cloud Computing but risk averse ■ Require best practices, patterns/anti-patterns around perceived high risk areas ■ Utilize vendors whose cloud services have a reputation for quality, stability and maturity.	■ Prefer "Business as usual" culture ■ "Proof of concept" required ■ "Proven Use Cases" and references within the same industry	■ Highly risk averse ■ Heard if all before – Network Computing, Hosting and ASPs ■ Requires education of all (Executives, Architects and engineers) ■ Requires education on technical through to business aspects of Cloud Computing ■ Plays catch-up with competitors

FIGURE 12.1
Cloud chasms

As with any model that attempts to characterize all organizations into a small set of categories, there will be exceptions to the rule. In essence, however, this model describes in a generic fashion the current and future penetration of cloud services in terms of organization types.

Risk Realization Chasm

Enthusiasts/innovators tend to jump headlong into cloud services without fully understanding the nontechnical aspects and significance to both IT and the business. Cloud services are not a purely technology play, and enthusiasts commonly miss opportunities to fully realize all the benefits made available by the use of cloud services.

Enthusiasts must understand that technology is an enabler and not the final solution and that their organization requires a cloud strategy that caters for disciplines such as enterprise architecture, governance, risk management, and compliance.

Wider Acceptance Chasm

Visionaries/early adopters want a strategic advantage and have the vision to understand the strategic business opportunity and impact that cloud services offer and utilize it to gain a competitive edge.

Visionaries have enterprisewide plans and build a roadmap accordingly. Although pilot deployments highlight the benefits that cloud services offer, visionaries encounter resistance when attempting to deploy on a wider basis within their organization. This resistance tends to center on the data privacy and security concerns that the rest of the organization has with public cloud services. In addition to this resistance, another area of concern for visionaries is that the level of benefits highlighted during the pilot does not manifest itself across all the divisions that deployed cloud services. This can come down to various reasons, including the following:

- Additional challenges that the extra human involvement and interaction bring when crossing divisional boundaries.
- The cloud management infrastructure cannot accommodate the large amount of cloud services and the federated distributed nature required.

For visionaries to cross this chasm, they must appreciate the need for a cloud GRC model that addresses the decision and accountability required. This requires the visionaries to have the will power and authority to address the political and culture change aspects of cloud services. In addition, visionaries need to provide a cloud architecture and engineering framework that can be applied when consuming and producing cloud services. Lastly, visionaries need to focus on the operational and post-deployment aspects of cloud services, including collecting, aggregating, collating, and presenting key performance indicators to improve engineering, operational, and the business aspects of cloud services.

Best Practices Chasm

Pragmatists appreciate the IT and business benefits of cloud services but are less risk averse to what they see as a high-risk proposition. These pragmatists wants to make sure that cloud service products have gone through multiple released versions and utilize standardized APIs and have matured to the point in which they are supported by the majority of mainstream products. In effect, pragmatists want a reliable standard. In addition, the pragmatists want access to the best practices and antipatterns that the enthusiasts and visionaries have encountered.

Pragmatists are risk averse and thus tend to consume cloud services from known focused cloud service providers with a reputation for providing quality, stable, and mature cloud services. To cross this chasm, pragmatists need to understand consuming and producing cloud services best practices and patterns and to have access to detailed case studies and references from organizations within their own industry.

Business Case Chasm

Conservatives delay any technical major decisions and tend to be followers. Therefore, for the conservatives to cross this chasm, they will wait until cloud computing hits the bottom of the hype curve and moves up toward the slope of enlightenment. The two key areas to assist with crossing this chasm is education and cloud business case development.

Education is paramount to the conservatives, especially around cloud costs and benefits and the differences between utilizing cloud services and traditional data center costs and benefits. Conservatives require a business case specific to their organization rather than taking a boilerplate example and plugging in numbers.

Conservatives tend to be hesitant because cloud services are classified as a disruptive technology, which conservatives are against because it requires political and culture changing events to occur within their enterprises. Conservatives would rather concentrate on the current-year objectives rather than on multiyear disruptive technology project that would affect the "business as usual" culture.

Conservatives will wait until cloud services becomes an established standard not only in their industry but also across multiple industries. At the end of the day, conservatives require a multitude of cross-industry cloud computing references with a detailed return on investment (ROI) business case before investing substantially.

Avoiding the Inevitable Chasm

Again, education is paramount to the late-to-market organization in detailing the differences between cloud services and past and existing technology strategies such as hosting, application service providers (ASPs), and virtualization.

Some late-to-market organizations will wait until cloud computing hits its adoption peak and then try to play catch-up with their competitors. This can be an expensive and disruptive proposition because cloud services require a culture change, which in turn requires time to feed in some learning experiences.

Because time is of the essence for the late-to-market organizations, this imposes more pressure on the late-to-market organizations as they try to speed the culture-changing process through. Therefore, the late-to-market organizations actually have a higher chance of failure.

The late-to-market organizations either don't understand the full scope of cloud services or believe that there is a more risk-averse manner in which to achieve the same benefits that cloud services offers. Late-to-market organizations believe that cloud computing is just marketing or hype and that the benefits and approach to cloud services is no different from ASPs or virtualization.

Planning Your Journey

To build your roadmap, you must fully understand your current environment so that you can make an informed decision and justify why moving certain aspects of your environment to the cloud makes sense.

As highlighted in Part III, “Life in the Cloud—Planning and Managing the Cloud,” many areas apart from the technology aspects of cloud services must be considered as part of your overall cloud services adoption roadmap. Figure 12.2 highlights that your cloud services adoption journey and speed will be dictated by your risk comfort level, governance model, level of executive backing, size of organization, and the current maturity of your processes and procedures within your data center.

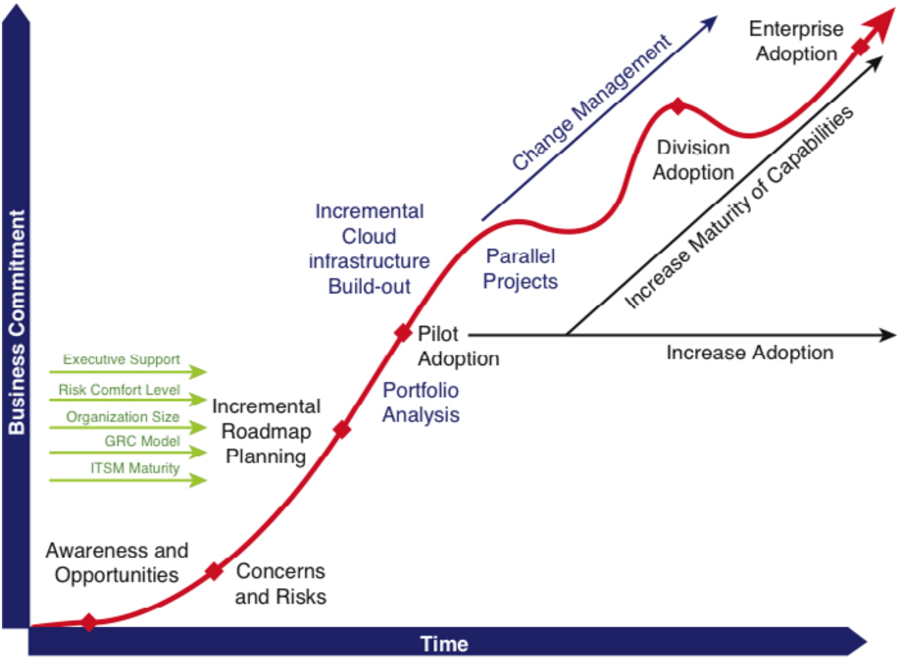

FIGURE 12.2
Business commitment, time, and journey will vary from organization to organization

Conferences, magazines, analysts, vendors, and (we hope) this book have raised your awareness and excitement about the opportunities that cloud services offers to your organization. As with any disruptive technology, however, a level of hype is followed by disillusionment as your concerns elevate when you start to understand the consequences and challenges that cloud services brings to your organization. You could say “silver clouds, dark linings.”

Many of the opportunities, risks, and challenges that your organization may encounter have been covered in earlier chapters. Fear can scupper your plans for strategically adopting cloud services. Therefore, as part of your overall adoption roadmap, do not underestimate the importance that change management will have in the successful deployment of cloud services within your organization.

Defining a cloud services adoption roadmap is no different from any other planning activity. As with all other planning activities, you need to fully understand the reasons why your organization wants to start utilizing cloud services. Depending on the process formality in your organization, this can lead to the development of a full-blown cloud services business case that documents the reasons and the cost, time, and quality benefits of cloud services.

There should be a combination of personnel from both the business and IT side of the house who should identify the strategic uses of cloud services for your organization. The path you take to adopt cloud services must be planned strategically and acted on tactically.

Your organization must understand the long-term goals and then make strategic decisions accordingly while at the same time delivering both incremental values to both the business and IT. In all likelihood, you will not fully realize the benefits of cloud services until you have executed a number of roadmap iterations. For example, noticeable cost savings due to economies of scale are unlikely to be met after one application has been moved to the cloud. Although you might not get all the cost benefits initially, these initial deployments give your organization the time to realign your operational and employee challenges.

You can then accelerate the deployment when internal resistance lessens. Until resistance lessens, a key success criteria for the execution of your cloud services adoption roadmap will be the level of executive support that you can garner and the willingness of your management teams to be innovators.

Make sure that your vision and expectations for utilizing cloud services are realistic for the maturity of your organization and the cloud services marketplace.
Because you will not realize all your goals on the initial project, it is probably best to view adopting cloud services on a multiyear horizon when focusing on a private cloud. Conversely, don’t fall into the old trap of analysis paralysis and attempt to answer every question and discover every risk before proceeding. This can easily lead you in fear of pulling the trigger. Instead, time-box your effort around assessing your current environment.

An important part of cloud services roadmap planning is understanding what your current environment offers in terms of functionality, costs, time, and quality. This is a real challenge to many organizations. For example, many organizations are not fully aware of how many servers (physical or virtual) are currently in production or what their individual running costs are (e.g., electricity, employee costs, license costs). Without an understanding of the current situation, it is much harder to make informed decisions about whether your current pricing model such as asset ownership is better or worse than a charge-per-transaction pricing model. Metrics such as availability, utilization %, and application network and data requirements are needed so that you can analyze and compare the expected results from utilizing cloud services against the current environment.

In addition to understanding your current environment, you must identify and assess the risks, barriers, and opportunities you expect when adopting cloud services. Doing so enables you to identify the specific cloud requirements that your organization needs.

These requirements drive the creation of the overall program and project plans focusing on resource needs, costs, risks, time, and deliverables.

When defining your future vision for the adoption and use of cloud services in your organization, consider your high-level goals and principles that will be utilized to guide the entire cloud services adoption roadmap. These principles lay the foundation for the cloud architecture, and most important, are the basis on which cloud governance decisions can be made in a defendable and repeatable fashion.

The gap between your current state and desired future state identifies key program-level activities that should be given priority when defining your cloud services adoption roadmap. As stated earlier in the book, cloud services are not a panacea, and therefore not all applications can or should be moved to the cloud. Review your project portfolio to decide which projects will be implemented in the cloud. This tends to be an easier approach than taking an existing application and updating its design to be implemented in the cloud. Consider your data privacy policies, project complexity, cost, risk, and business criticality. Involve as many parts of the organization as required to identify the appropriate projects and to alleviate its concerns. This includes involving departments such as security, audit, and legal.

We cannot recommend which cloud services your organization should adopt because there is no one way to leverage cloud services for your business. However, organizations have typically focused on public cloud services that provide noncore functions that cater for, for instance, email, collaboration tools, on-demand testing environments, websites that have seasonal peaks in website traffic, and customer relationship management (CRM).

These low-hanging fruit cloud services have been chosen primarily for their benefits around cost savings, lack of any real integration work required, lower network latency and security requirements, and minimal political pushback. As both the public and private cloud services marketplace matures, so will the confidence levels rise of cloud service consumers, and with it we will start to see organizations use cloud services in other major areas. The speed of cloud service adoption and the types of cloud services consumed are dictated by your organization’s characteristics, as discussed earlier in this chapter.

After you have identified and prioritized the appropriate projects, you can then derive your cloud service requirements from the needs of these projects, which can lead to the development of a cloud service consumption model that details which cloud services will be consumed (and how, by whom, and when). This leads nicely to your cloud vendor sourcing strategy; that is, whether you intend to fulfill the cloud service consumption model with internal services hosted within a private cloud, use a public cloud service, or use a mixture of the two that meets your requirements.

Summary

As stated earlier, we are in the early phases of cloud service adoption, and the marketplace has many hurdles to overcome. Nevertheless, organizations today are realizing the benefits that cloud services offer. As your maturity and comfort level increases over the years, adapting your cloud service adoption roadmap will be critical. This chapter has covered organization types and related characteristics when adopting cloud services and the potential pitfalls and risks you may encounter. This will give you a heads-up as to some of the near-term challenges you will need to address in your roadmap.

Key points to consider from this chapter are

- A cloud services adoption roadmap consists of program-level efforts and a portfolio of cloud services.
- Program-level efforts provide and enforce the necessary consistency required to succeed at cloud service adoption.
- Without an adoption roadmap, there is a higher chance that you will encounter challenges that may undermine your organization’s effort to realize the benefits of cloud services.
- Make sure that your vision and expectations for utilizing cloud services are realistic for the maturity of your organization and the cloud services marketplace.

The next chapter covers where the cloud services road seems to be heading, which in turn will require you to update your roadmap accordingly.

CHAPTER 13

Conclusion

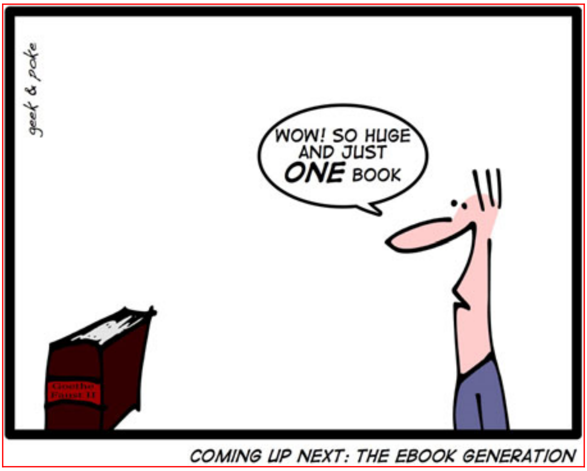
geek & poke
WOW! SO HUGE AND JUST ONE BOOK
Goethe Faust II
COMING UP NEXT: THE EBOOK GENERATION

Where to from here?—Geek and Poke

The transformations taking place as a result of the cloud approach to delivering services have been shown as immense and rapid. Throughout this book, we have approached the cloud from the perspective of an enterprise looking to use cloud services for the benefit of the business. In addition, we showed examples of business models that have radically changed and new businesses that have been created as a result of cloud services.

The rapid evolution of both reliable hardware and connectivity make potential information-related roles possible in many more locations than ever before. So, what will we see? As we look to the future and prognosticate, we envision several trends that will significantly impact business models, capabilities, and ultimately, all our lifestyles. These topics to discuss are as follows:

- Big data
- Big architecture
- Communications, networking, and the interconnectedness of all things
- Privacy issues

- The rise of the broker
- The rise of community clouds
- Rapidly changing billing models

Big Data

The amount of data we are storing and need to manage is exploding! Further, the types of data and metadata about that data are changing and growing almost exponentially. From the enterprise perspective, consider log files in enterprises that feed compliance programs and security systems. Consider intellectual property (IP), or media assets such as images, music, and videos. Individuals are increasingly loading videos online to share and creating backups of all their personal data, including pictures, music, and videos. The popular video site YouTube reported that changes in video quality to higher definition, and the introduction of new approaches altogether, such as 3D, would further expedite this trend.

Most large enterprises deal with big data and try to capitalize on the value by undertaking data warehousing and mining efforts. This data grows daily, with many Fortune 500 companies claiming growth between 20–40 percent annually. The amount of data generated outside is even greater. Eric Schmidt, CEO of Google, has stated several times, most recently at the 2010 Techonomy conference:[1]

> There was 5 Exabyte's of information created between the dawn of civilization through 2003, but that much information is now created every 2 days, and the pace is increasing.

Google is in the business of making money from big data. The dark lining here is the potential for malicious use of big data. Corporations worldwide already admit to not only doing credit checks on potential employees, but also searching the web for background information.

Similarly, the amount of bad and useless data also is on the rise. For example, according to the Symantec Intelligence Quarterly Key Findings in the "Overview of Internet security trends for April–June 2010"[2]

> Symantec observed 12.7 trillion spam messages during this quarter, accounting for approximately 89 percent of all email messages observed; this is a decrease from 13.8 trillion and 90 percent, respectively, in the previous quarter.

In this case, the data is garbage. Worse still, the amount of resources wasted at a technical level in terms of storage, networking bandwidth, and compute capacity for SPAM filtering, and sheer cost in terms of effort on a human level, is immense.

These trends are also felt in businesses today in terms of the surge in requirements to store that compliance data, security event, and financial reporting data, IP, and so forth. Much of this is machine-generated data that either needs to be stored or intelligently curated as it is stored. The healthcare industry worldwide is increasingly moving to electronic patient records. As these records increase in size with scans of X-rays, lab results, and more, the need to secure and manage them will see increasing focus, too.

In July 2010, Verizon launched a service that it hopes will become the nationwide backbone and foundation of a new wave of health information exchanges (HIEs). The sheer amount of healthcare information collected makes cloud services an obvious choice to handle this large amount of data. SearchCloudComputing.com discussed the launch, stating the following:[3]

> Verizon's HIE is typical of cloud services, in that it is built on already available technologies. Virtualized instances of Oracle Healthcare Transaction Base provided the back-end database. Verizon says it can spin up new instances wherever and whenever necessary to handle data location and compliance issues for healthcare providers.
>
> Verizon makes it clear that it can provide all relevant compliance needs, from HIPAA and HITECH to PCI DSS, and meet any security standards necessary. It says data is mirrored on each coast and connected via Verizon's IP backbone.
>
> The company says that it has deliberately made the service as flexible as possible. Verizon states that it can serve doctors' offices with two physicians and a few thousand patients as well as it serves the largest provider networks in the country, which can have thousands of doctors and many millions of patients a piece.

These are examples in which data is growing significantly, and the concern today is largely about how to manage the storage requirements. However, consideration should be given to the ability to use colossal data sets for processing, too. The ability to use an amalgam of those data sets is significantly enhanced through cloud services. We discussed the trends associated with cloud services and related search topics in Chapter 1, "Introduction to Cloud Computing," using Google Trends to show how fast and significant cloud computing was as a trending topic. As a targeted derivation of Google Trends, a tool called Google Flu Trends was released in late 2008. Google Flu Trends offers a great example of the exceptional and unique opportunities obtainable through big data. Google had identified a specific set of search queries that became more common during flu season, stating the following in its announcement on the Google blog:[4]

> Our team found that certain aggregated search queries tend to be very common during flu season each year. We compared these aggregated queries against data provided by the U.S. Centers for Disease Control and Prevention (CDC), and we found that there's a very close relationship between the frequency of these search queries and the number of people who are experiencing flu-like symptoms each week. As a result, if we tally each day's flu-related search queries, we can estimate how many people have a flu-like illness. Based on this discovery, we have launched Google Flu Trends, where you can find up-to-date influenza-related activity estimates for each of the 50 states in the U.S.

There are challenges. A study by the University of Washington (UW) into Google Flu Trends compared with actual CDC records was released in May 2010. The study was not wholly complimentary but did conclude that the approach had merit, as noted in a release by the American Thoracic Society, with comments from Dr. Ortiz:[5]

> Google Flu Trends influenza surveillance provides an excellent public health service, because it provides nationwide influenza activity data in a cheap and timely manner. Nevertheless, our study demonstrates that its data should be interpreted with caution and that other surveillance systems more accurately reflect influenza activity in the United States.

In essence, the UW study shows that the approach is not perfect but would certainly not be possible without a massive data set to draw from. The point of this example is to demonstrate the benefits aggregating billions of search results to produce meaningful and actionable information. Whether for research, planning, or personal effectiveness, this tool is available to all. Here's more from the Google blog:

> We couldn't have created such good models without aggregating hundreds of billions of individual searches going back to 2003. Of course, we're keenly aware of the trust that users place in us and of our responsibility to protect their privacy. Flu Trends can never be used to identify individual users because we rely on anonymized, aggregated counts of how often certain search queries occur each week. The patterns we observe in the data are only meaningful across large populations of Google search users.
>
> The CDC does a great job of surveying real doctors and patients to accurately track the flu, so why bother with estimates from aggregated search queries? It turns out that traditional flu surveillance systems take 1–2 weeks to collect and release surveillance data, but Google search queries can be automatically counted very quickly. By making our flu estimates available each day, Google Flu Trends may provide an early-warning system for outbreaks of influenza.

And this is an example where big data is possible and has significant value if allowed to be used, privacy and security issues aside, in a compelling way. In terms of data, you need to consider the following:

- What data would be useful to your business?
- Is it available?
- Do you have data or data services that could provide value to others and opportunities for you?
- Do your competitors have the same access, and is its use a new baseline or can it become a competitive advantage?

Big Architecture

Massive hardware, massive changes to data distribution architecture requirements, new usage models, and new operating systems are on the rise. This is what makes Google, Facebook, and other significant web properties possible. And these guys are buying into the concepts of big data and are using big architectures to deliver their services.

Today's systems need to be re-architected and rewritten to take full advantage of the cloud models. Admittedly, not all of them need to be, but if there is a path that enables better delivery, business advantages to be created, or new business to be created, it should be examined.

The tools and architectures used previously do not transport cleanly to cloud environments because the infrastructure does not operate in the same way. Tools that allow conversion from an OS-based solution to cloud environments will be less important and grow to do true analysis to recommend architectural changes, or disappear. This new thinking requires vendors who offer development environments to change their model. This will not happen overnight, and we can expect to see more efforts such as Hadoop, NoSQL, and the like take on established players in key markets up, down, and across the cloud stack. Microsoft's attempt to allow existing developers easy migrations to Azure through its familiar Developer Studio environment is another way that existing players will attempt to maintain control and dominance. Microsoft has a good handle on this in terms of tools and technologies. Competitors aiming to eat away at that need to take a look at Salesforce.com's effort to provide migration tools to the Force.com platform as a service (PaaS) for customers using Microsoft technologies.

All this largely ignores the changes we discussed in terms of mobility and device independence. An approach closely aligned with the concept of device independence is that of the cloud operating system (cloud OS). The thought here is that, to support user needs if all the data and the bulk of data processing occur in the cloud, devices should not have a complex OS and application stack.

Google is pushing this vision with the Chrome OS. This approach requires some thought on behalf of customers and users. For now, history shows a rather different result from these initial thoughts. Specifically, an attempt to create simple devices with central services has always devolved to more intelligent and complex device requirements as the next layer of requirements is understood. Simple Java Runtime Environments gave way to complex libraries and management requirements being downloaded later as device capabilities become greater and user requirements became more complex.

The parting thought is that a new approach will arise, and both existing and new architectures will increase in market share. Simple dedicated client environments will increase significantly, even if they run within virtualized client containers, but we can expect increases in the level of heavy clients, too. Factoring this into your thinking with respect to customer access to your services and the best approach requires a simple focus on how best to deliver in what timeframe, rather than critically relying on a winner.

For now, heavy clients will be pervasive, but in the future it will not matter which you choose.

Communications, Networking, and the Interconnectedness of All Things

Fundamental to the success of cloud services is the communications framework on which it relies. The importance of the network is driving significant innovation but at the same time is limited by legacy. The ability to make voice calls, listen to audio tracks, or watch streaming from almost any corner of the world in high-definition video on various devices shows the flexibility of the Internet's infrastructure.

One vision holds that many services will be hosted in the cloud over time and that simple, but powerful devices will be used to access that content from any place on the planet. Consider that most new smartphones first look for high-speed data connections prior to using their cell network service, for a number of reasons: reliability, traffic management, and cost. Moving forward, as communications network options increase to support those requirements, devices of all types will automatically look for the most optimal route to get to the services they require.

The approach to cloud services creates the ability to mash all these capabilities and create any version of communication control that you could want. Some want full control from their phone. Some want control from a centralized enterprise service. Others want everything in the cloud. As we move forward, this flexibility will become the norm.

Hybrids that incorporate devices and backends in the cloud are multiplying. Amazon launched its Kindle eBook reader with the boast that you can order any book and receive it within seconds. Apple launched the iPad in 2010 to much anticipation and some confusion (as there are obviously many potential uses for such a device). Many new devices will link to their own ecosystem but will also allow for integration into others, including enterprise or consumer cloud services.

From the enterprise perspective, prepare your communications infrastructure and related services to support employees in any time zone and from any location in a secure manner. This is not new for those who have virtual private network (VPN) access and 24 x 7 support; but for those who do not, this is the time to rework policies, procedures, and architectures for this model.

Work with national or global managed service providers when possible to gain the best rates for data over cell- and remote-device access. As each device is replaced, be it a phone, laptop, tablet, or otherwise, the ability to use cloud services is increased, as is the likely data requirements. Providers need to offer comprehensive plans for customers, from individuals to global concerns, to adequately address these needs; and as customers, you need to drive that model.

To leverage this enhanced mobility, organizations will increasingly adopt new business models. Examples of such include the following:

- Work from anywhere—not just teleworking, but full telepresence.
- Take terabytes of data with you, or access zettabytes or more from cloud services.
- User choice of devices as security options improve and availability increases.

Although analysts have differing estimates, they all agree that in the next decade at least seven to ten times the number of mobile devices will access data and browse the Internet than there will be PCs.

Privacy

The concerns about privacy as discussed in Chapter 8, “Cloud Governance, Risk, and Compliance,” will continue to increase. Quite simply, as noted in the discussion about big data, data has value and will be exploited in a way that does not conflict with the business value or trust. This suggests that use of personally identifiable information (PII) will increase to benefit businesses, and unwanted exposure of that data will likely increase in parallel. As a result, expect legislation to increase and the business environment to see greater complexity, especially in the United States, which provides no clear rights to data privacy. Without some basic tenants similar to the European Union’s Data Protection Directive, which details a fundamental right to privacy in many forms, the value of this data will continue to be exploited. These rights decrease the risks of potential PII exposure but cannot completely eliminate them.

Another viewpoint suggests that attempts to protect private data are too late. The data is already out there. This means that you need to consider privacy management in a different way.

Alternatively, consider a scenario in which the attackers are already inside your network. What liability do you face today if this were the case?

Whether you manage one, tens, hundreds, or thousands of computing devices, it is likely that they will or already have been compromised in some way. Antivirus, antispam, and all the other techniques we use continue to be compromised, and this security model is failing. Therefore, new approaches such as service isolation, virtualized sessions, managed cloud security, and more will become more prevalent.

From the compliance perspective, we noted in Chapter 8 that you need to avoid future issues by ensuring that your provider will not share any information beyond their current usage in the event of shutdown, transfer of ownership, or any similar type of asset disposition.

The Rise of the Broker

Integration brokers, as a model, have existed for many years. Covisint, a Compuware company, provides data exchange brokerage services alongside identity federation capabilities for the manufacturing and healthcare industries. Vendors such as HubSpan, Boomi, and Biz365 provide examples of data integration brokers that have moved into the cloud service provider market.

Identity brokers are another example. In Chapter 8, the problem of managing the identity lifecycle across cloud service providers was discussed. Today, the traditional enterprise identity management vendors have been slow to include cloud service solutions, but increasingly we see this in products from CA and Novell. The challenge is that these products are still architected for enterprise deployments and as such are not easily scalable for multitenant use. Identity brokers, therefore, are coming from smaller companies and start-ups such as Ping Identity, Conformity, Exostar, Symplified, and more.

Current models of infrastructure service vendors (ISVs) need to change to incorporate selling and integrating packages of secure password authentication (SPA) solutions. Because being onsite to deliver solutions will become less critical, however, this will see some channel conflict as larger vendors end up offering/hosting their own solutions, which can easily compete with their traditional delivery partners. This has already caused some issues, such as when EMC shut down its Atmos Cloud Storage solution, after less than a year, citing channel conflict (or confusion) as one of the primary reasons.[6]

The Rise of Community Clouds

The need for varied security, legal, performance, availability, functional, and geographical controls will increase the use of community clouds. Solutions such as these will come from group efforts or from vendors wanting to deal with a specific type of market.

Google offer "Google Apps for Government," designed to meet the specific requirements around policy and security from the U.S. public sector. In this case, Google has taken its Google Apps offering and created an operational delivery model based on a specific market requirement. Google gained Federal Information Security Management Act (FISMA) certification and accreditation, as many products and service providers before them have done. Along with standards adoption, this then allows for U.S. public sector groups to utilize the Google cloud offering with a much lower barrier to entry. However, from the cloud service provider perspective, this sort of approach means a change from its approach of a single solution offering that meets all needs. Now Google and other vendors who want to meet these needs have to manage two or more implementations of their cloud services. The U.S. government for its part is pushing ahead with catalogs of approved services allowing for an application selection self-service model to be made available to IT staff in various government agencies.

The financial industry is seeing increasing use of this. In April 2010, the *Australian Financial Review* reported that the Commonwealth Bank of Australia joined with other worldwide financial institutions, the Bank of America and Deutsche Bank, to pool their purchasing power in relation to computer software.[7] According to the article, "Banks have begun demanding pay-as-you-go IT services, and this could result in substantial changes in the operating structures of major IT corporations. The application of cloud computing on this scale could save billions of dollars in the banking sector." So approaches to dealing with software vendors are changing quickly and significantly, with great expectations that cloud delivery models will be available.

We are already seeing vendors such as Amazon, Google, Microsoft, IBM, and HP delivering cloud services to communities. But there will be initiatives such as NASA's Nebula[8] and the DISA's RACE[9] that are built and managed by the organizations that want to offer them to their respective communities as well, blurring the lines between private and community cloud services.

Many of these types of efforts are taking advantage of open source software solutions, and modeling their usage and access models on open source approaches. Further, NASA, along with other businesses including Rackspace and Citrix, are supporting efforts such as OpenStack,[10] which is intended to "allow any organization to create and offer cloud computing capabilities using open source software running on standard hardware."

Rapidly Changing Billing Models

Expect to see costs slow their rapid decline as both customers and cloud service providers realize that more capabilities in terms of security, reliability, and configurability require more infrastructure and more staff to deliver. Although the cost of infrastructure continues to decline, that decline is not at the rate with which providers have been dropping their prices, and the point comes where pricing decreases level out. The cost of support will require cloud providers to increase availability and reliability as they try to avoid this cost.

Consider that providers can also benefit from the flexible billing models, and that may create problems for you.

As a result, pricing models will continue to be explored, trying to find the right balance between growth and revenue maintenance. The sporadic nature of some cloud services will create an extension of charge models that offer longer-term contracts more discounts. Some vendors are already offering this type of contract, and preparedness plans wherein you reserve capacity for a minimal fee over time, such that your actual cost per use is also decreased. Talk to vendors about their plans for billing; it is likely to change and impact your cost-benefits analysis. Watch the changes as managed services providers (the telecom and communications vendors worldwide) change plans, offer new plans, and deprecate others. In many cases, the costs come down while the restrictions increase. Whereas these changes occur over months and years in that business, cloud service providers can change their offerings in minutes based on market circumstances (demand, competition, regulatory, or otherwise). Therefore, unless your contracts are longer term, too, you need to deal with these changes.

So, those are the trends we see. Now, why silver clouds and dark linings?

Silver Clouds, Dark Linings

The market for cloud services is emerging but growing explosively due to compelling business drivers and a combination of unrestrained enthusiasm. However, cloud services are not a silver bullet. (Or should I say silver cloud?) There are dark linings in those clouds, such as security, privacy, and reliability. Don't let the attraction of scalable, cheaper, and quicker solutions obscure these risks.

At the same time, don't let these dark linings stifle your adoption of cloud services. Instead, make sure that you understand and have a compelling, incremental, and pragmatic approach to address these challenges in a controlled manner whereby you can capitalize on cloud services.

This requires planning and the recognition that to succeed with cloud services, business and IT leaders must deal with the fact that the role of IT is changing within the IT supply management chain. This can enable your IT department to focus on and deliver more innovative capabilities rather than hardware and software updates.

As it is not possible, or sensible, to wholesale move your entire enterprise to using cloud services, here are the key questions you should be asking yourself:

- How will my business change as a result of these critical new capabilities and business models that the cloud supports?
- Who will be my greatest competitors and partners as a result of these changes?
- How do my business processes and solution portfolio relate to my financial and business models?
- What risks and change management challenges will cloud services expose in my organization?
- Can my technology group move at the right speed to take advantage of cloud services?
- Does my technology group have access to the right training, expertise, and tools to be successful?
- How utilized, effective, and cost-efficient is my current information technology infrastructure?
- What is my short-term and long-term plan to take advantage of the opportunity cloud services offer?

As stated at the start of the book, shift happens!
Make sure you are ready for it.

Our goal with this book was to create a high-level framework to highlight key concepts and plans for the cloud. Our intent was to focus on the needs of our business and technology brethren without needlessly digging into core technologies that either are already well understood or are in a state of flux. In addition, we did not want to delve into specific vendors that are well known or may not even exist after this book is published.

In conclusion, the concepts and approaches delivered are commonsense discussions in dealing with new business models and technology trends, using cloud service examples as the basis for discussion. We believe that we have delivered on those goals and look forward to further discussions as the cloud services market continues its rapid evolution.

Sincerely,

Archie Reed and Stephen G. Bennett

We look forward to your feedback:

Twitter:	@concisecloud
Email:	concisecloud@gmail.com
Web:	http://www.concisecloud.com/

Endnotes

[1] http://www.techeye.net/internet/google-ceo-warns-of-data-explosion-and-future-without-privacy.

[2] http://www.symantec.com/content/en/us/enterprise/other_resources/b-symc_intelligence_quarterly_key_findings_apr-jun_21072011.en-us.pdf.

[3] http://searchcloudcomputing.techtarget.com/news/article/0,289142,sid201_gci1516731,00.html.

[4] http://googleblog.blogspot.com/2008/11/tracking-flu-trends.html.

[5] http://www.thoracic.org/newsroom/press-releases/conference/articles/2010/google.pdf.

[6] http://www.atmosonline.com/?page_id=366.

[7] http://afr.com/p/business/technology/banks_seek_billions_in_it_savings_JBwXh6ZmMhhRLrGUMBZCSK.

[8] http://nebula.nasa.gov/.

[9] http://www.disa.mil/race/.

[10] http://openstack.org/.

REGISTER
THIS PRODUCT
informit.com/register