CHAPTER 1
Why Cloud, Why Now?

> There was a time when every household, town, farm, or village had its own water well. Today, shared public utilities give us access to clean water by simply turning on the tap; cloud computing works in a similar fashion. Just like water from the tap in your kitchen, cloud computing services can be turned on or off quickly as needed. Like at the water company, there is a team of dedicated professionals making sure the service provided is safe, secure, and available on a 24/7 basis. When the tap isn't on, not only are you saving water, but you aren't paying for resources you don't currently need.
>
> **— Vivek Kundra, former federal CIO, U.S. government**

In 2009, I was invited to the IBM Impact conference in Las Vegas as a guest blogger and analyst. *Cloud computing* was a vastly misunderstood term at that time, and there were very few enterprises leveraging any cloud services other than a few of the mature SaaS solutions like Salesforce.com and Concur's expense management software. I witnessed some very intelligent senior IT people from various companies scoffing at the term *cloud computing*. I can still hear the lines: "We were doing this on the mainframe in the '60s" and "There is nothing new here, this is just a fad." At that time, my team of one developer was testing a prototype that was executing hundreds of thousands concurrent point-of-sale (POS) transactions to the cloud and back in subsecond response time on a virtual cloud server, costing us about half a dollar an hour charged against my CEO's credit card. I started to think about how much it would cost to implement the infrastructure, licenses, and professional services to perform a proof-of-concept on-premises. I also thought about how many months it would take to go through a vendor evaluation, the procurement process, and all of the steps required to make a capital expenditure that would have been required to buy servers from a large vendor like IBM. At the end of several months, I would finally have all the hardware, software, licenses, and professional services that my developer would need to test his proof-of-concept. My start-up would have been out of cash by then, and all I would have to show for it would have been a few lunches paid for and a nice golf shirt with the vendor's logo on it.

Instead of doing things as if we were a large company with all the time and money in the world, my team embraced the cloud and saw it as a competitive advantage. Our competition was two to three years ahead of us, but we felt we could provide superior products and services at a price point that could not be matched by companies that were purchasing and managing infrastructure and data centers. My developer was able to provision many different-size servers and test multiple configurations until he finally found the sweet spot. Our cloud services provider, Amazon Web Services (AWS), had made infrastructure management easy by abstracting out all of the complexities into a few simple application programming interfaces (APIs). We could build and deploy numerous server configurations in minutes and de-provision them when we were done. That is a drastic change from days past. Before cloud computing, the thought of asking one's boss to purchase three different types and sizes of servers in order to run a series of tests and theories on them to determine which one is actually the right one was not a feasible or career-enhancing proposition. Buying many different hardware configurations, tossing out the configurations that did not perform optimally, and then going back to procurement to buy more of the right configuration is an insane approach when dealing with physical machines. In the cloud, this is a best practice. Cloud computing resources follow a pay-as-you-go pricing model just like electricity and water. It is easy to test multiple configurations in a prototyping environment with very little investment.

Going back to my example, we used a simple management console to launch numerous virtual computing resources that were ready to run in five minutes. We would run our tests for two hours and then discard the virtual computing resources. Our total cost was 50 cents an hour or $1 worth of infrastructure to run this prototype. Then we would move on to the next server configuration and try another series of tests. We would do this three times during the day and rack up $3 in infrastructure costs. Here is a comparison of prototyping in the cloud versus prototyping in the same manner on-premises:

- **Scenario A (on-premises).** Buy three different severs at roughly $3,000 to $5,000 each, plus software, shipping, and installation.
  - Elapsed time to procure and implement likely to range between one and three months.
  - Outcome: Decide on which server to keep, buy more, get rid of the other two.
- **Scenario B (cloud model).** Developer creates three different virtual computing resources within minutes at $0.50/hour, using one at a time for two hours each (total $3.00).
  - Complete testing and make a decision in one day.
  - Outcome: Complete the entire scenario in one day of work for just $3.00 plus one person's salary. No wasted assets.

That is just one real-life example that made me a believer in cloud computing. As we continued our journey as an early stage start-up I was continually amazed at how quickly we could get work done at such a low cost. We owned no hardware and leveraged open source software. Since we did not have to manage data centers and physical infrastructure, we were able to focus on building product to generate revenue so that our start-up could live to see another day.

Evolution of Cloud Computing

My first job out of college in 1988 was a COBOL programmer working at a steel plant in the South. We were migrating from an old Burroughs mainframe computer to a new IBM 3090 mainframe, which, to put things in perspective of the level of coolness in today's terms, is the equivalent of moving from a feature phone to a smart phone. The code of the first program I worked on was written the year I was born. It had been ported from one mainframe system to the next and was 23 years old at the time. When that code was written, a lot of engineering went into breaking up messages into very small chunks of memory because of mainframe memory limitations in the 1960s. Here we were in 1988 with this then-massive IBM mainframe with what seemed like endless amounts of memory and using code that was working really hard to break things down into 8K chunks. I thought this was absurd, but as a 23-year-old rookie, I shook my head and ported that code to the new system. Little did I know that some 25 years later, masses of people would be taking that same approach as they began porting their legacy applications to the cloud without even considering that the new target environment is drastically different and more powerful than the environment the legacy code is running on. We will discuss stories like this in Chapter 3 ("Cloud Computing Worst Practices").

Cloud computing is the result of many years of evolution dating back to the first computers. It is the natural progression from the centralized mainframe era, to the distributed client-server era enabled by the birth of personal computers, to the Internet era where the enterprise was able to connect to the rest of the world through a network of computers that spanned the globe. Back in the mainframe days, systems were centrally controlled and managed. The mainframe administrators were the powerful gatekeepers of all data and all systems. They were also often the biggest bottleneck because nothing could get done without going through them. When the PC was born, IT professionals were empowered and were able to distribute workloads across many work nodes without having to go through the once-powerful mainframe gatekeepers. This was both advantageous and dangerous. It was advantageous from the standpoint that systems were built and deployed faster, cheaper, and with richer features. It was dangerous from the standpoint that in return for

the gains in agility and flexibility, we witnessed a huge decrease in the effectiveness of governance and security.

Another way to say it is we gave up manageability and standardization for speed to market. The distributed nature of PCs in a client-server architecture created a "Wild West" effect, where applications could be deployed rapidly without the assurance of the proper security and controls in place. The net result was applications became more nonstandard and filled with gaping security holes that gave birth to the rise of security breaches, identity theft, and cyber-threats at levels never witnessed before.

In addition, management of the enterprise became a very complex and expensive challenge. In fact, one could argue that the birth of the client-server era was the defining moment where business and IT alignment took a turn for the worst. In the mainframe days, the mainframe and IT's sole purpose was to build systems for the business to enable business strategies. We built financial systems, payroll systems, and systems that drove the business's core competencies and automated operational processes. The PC-enabled client-server era allowed IT to create systems faster and cheaper but introduced new challenges like integration, interoperability, widespread patching, and much more. These complex issues led to a lot of IT-centric tasks that shifted large numbers of IT resources away from business enablement to IT maintenance. In addition, this era gave birth to a whole new breed of infrastructure, security, and operational professionals who spend the majority of their time working within IT silos dealing with issues and projects that do not drive revenue or increase profitability. In fact, much of this work increases opportunity costs for the business by consuming resources that could be directed toward increasing revenue or reducing waste.

Then came the Internet, which extended the enterprise to the outside world. Now companies could integrate their systems with their suppliers. Customers could go online and purchase goods and services in self-service mode 24 hours a day, 365 days a year. Software vendors could now deliver services as hosted solutions, eliminating the need to procure and manage hardware on-site. The Internet created a global revolution where any company or individual with an Internet connection could now do business from anywhere in the world, any hour of the day.

Once again, the level of complexity of systems increased dramatically. The level of control and governance decreased significantly. Applications became even more insecure, creating opportunities for people and organizations with bad intentions to attack systems and steal and sell data, which created a whole new industry of products and services to secure systems. When the Internet was being touted as a huge technology innovation, I remember the pundits fighting the movement while waving the security flag. This is the exact same thing we are witnessing today with the adoption of cloud computing. The same people and others like them are riding the security bandwagon in resistance to the next biggest thing in technology.

What history shows us time and time again is that every new technology innovation is met with resistance. Early adopters and risk takers embrace the new technologies and become the guinea pigs for the enterprises that typically prefer to wait until the technology becomes mature. The trailblazers take advantage of these new technologies and create tremendous business value. As success stories start becoming more and more frequent, demand goes up. As demand goes up, issues like standards and security become prominent topics and major hurdles for mass adoption. Standards start to emerge, best practices are published, and vendor and open-source products start becoming widely available to fill in the gaps. Cloud computing, like the Internet several years before, is at that tipping point where many organizations are moving from the *why* question to the *how* question.

In Figure 1.1, the diagram uses the Gartner Hype Cycle terminology to describe how technology matures over time.

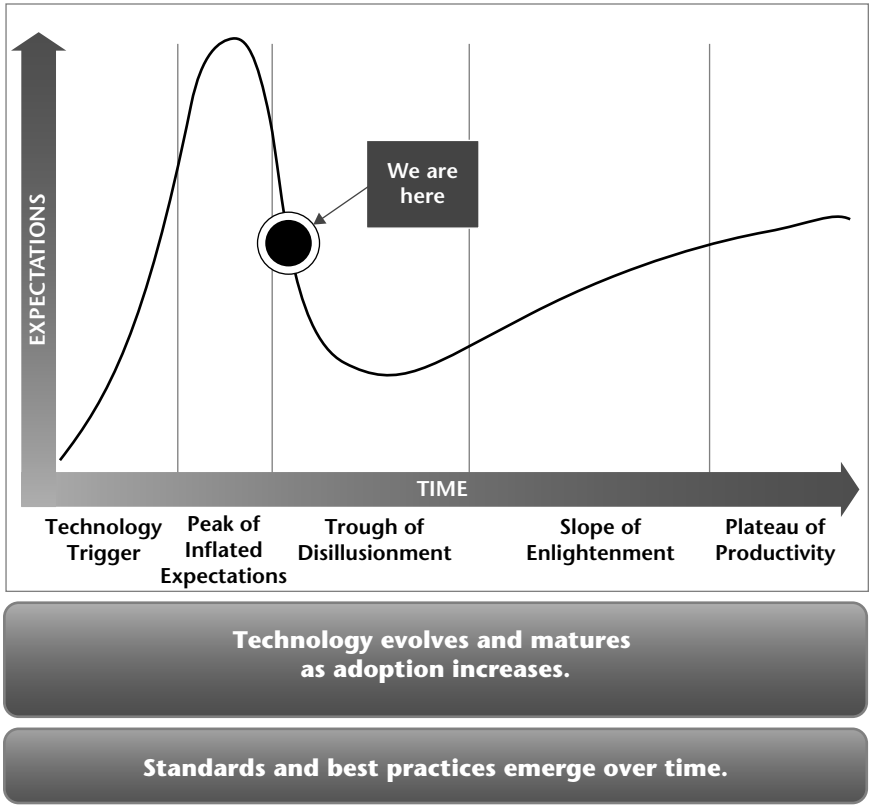

Figure 1.1 Cloud Maturity

As time progresses and more companies adopt cloud technology, the expectations move from hype and confusion in the early years and migrate toward broad acceptance as standards, best practices, and success stories emerge. Currently we are somewhere between the peak of inflated expectations and the disillusionment. As of this writing in early 2013, cloud computing is widely accepted by start-ups and small and medium businesses (SMBs), but large enterprises are late in adopting cloud computing. This is due to the complexities that come with years of legacy architectures, existing infrastructure and data centers, and organizational challenges.

The mind-set of large enterprises is changing rapidly in 2013 as many cloud service providers are delivering products and services that cater to enterprise-class clouds, where previously only commodity-class clouds were available. Commodity clouds were designed to commoditize infrastructure and offer it at low cost with the capabilities to achieve high scale and self-service capabilities. Enterprise-class clouds were designed to meet or exceed the security and service level agreements (SLAs) of the on-premises infrastructure they replace. Enterprise clouds are more expensive and complex than commodity clouds, but commodity clouds often do not meet the security, regulatory, and SLA requirements required within the enterprise.

Figure 1.2 shows how security maturity often lags behind in the adoption of new technologies, which delays mass adoption by large enterprises. The early pioneers and risk takers blaze the trail and eventually the best practices and security vendor solutions emerge from those early lessons learned. It appears that 2013 will be the year where enterprises start to embrace cloud computing in large numbers, as we are seeing huge turnouts to cloud computing conferences and a large increase in budget dollars allocated for cloud computing.

Enter the Cloud

Cloud computing combines the best of the mainframe era with the best of the PC-enabled client-server era along with the Internet era. Those old-timers from the conference I mentioned earlier were right. “We have been doing this for years,” as they said. What they missed, though, is that now we can do it at scale, using a pay-as-you-go billing model, at speeds never accomplished before, and all without ever buying any hardware or building any data centers. If managed correctly, cloud computing can give us back a lot of that central control and governance that we had from the mainframe days. At the same time, the cloud makes available to us a massive amount of distributed computing resources, gives us broad network access over the Internet, and bottles it

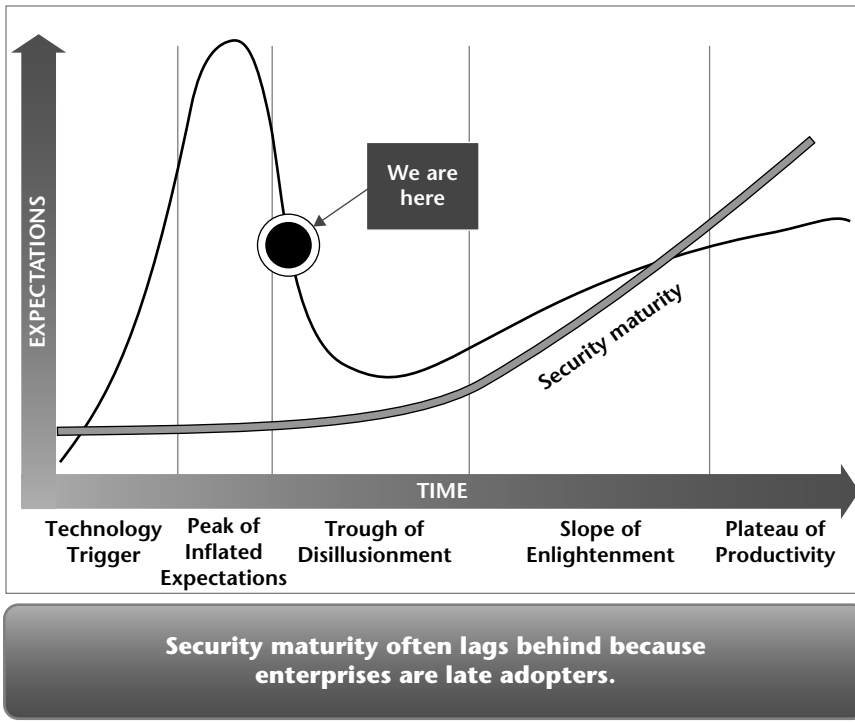

Figure 1.2 Cloud Security Maturity

up so we can pay for it as a utility like electricity or water. We pay for what we use and we turn it off when we don't need it.

It is true that many of the concepts of cloud computing have been around for years. What is new is that many of those lessons learned and techniques in computer science that have been perfected over the past few decades are now able to be simplified and automated and made available as highly abstracted on-demand services and offered at price points that are hard for the traditional on-premises or commercial software industry to compete against. The days of requiring customers to purchase and manage hardware and software licenses are long gone. Most customers now expect their needs to be met over the web either as an on-demand software solution (Software as a Service), a platform for quickly developing scalable solutions without all of the infrastructure costs (Platform as a Service), or a virtual data center for building scalable solutions at a lower cost (Infrastructure as a Service). These three cloud service models will be discussed in detail in Chapter 2.

When people tell me that the cloud is nothing new and that we have been doing cloud for years and years, I give them this analogy. "The iPhone is nothing new. We have been doing phones for years and years." My point here is, yes, we have been using telephones for decades and decades, but the iPhone is radically different from the rotary phone that I used as a kid, and it has made a major impact on businesses and on our lives. The cloud is to computing as the iPhone is to telephones.

Still not convinced? Here are some case studies of companies leveraging cloud computing to create business value. Each case study has a very compelling story of how great teams leveraged the cloud to get to market fast, scaled to incredible levels, and did it without buying hardware.

Start-Up Case Study: Instagram, from Zero to a Billion Overnight

In October 2010, a photo-sharing application called Instagram was launched, and 25,000 people registered on that first day. Three months later, Instagram had 1 million users, and shortly after hit 10 million. At that time, the company only offered an iOS version of its mobile application, so it was only capturing iPhone users. A year and a half later, Instagram had close to 30 million users. When it finally launched the Android version, it acquired 1 million users on the first day. In April 2012, less than two years after it launched, Instagram was bought by Facebook for an estimated $1 billion. In September 2012, just shy of two years from its initial launch, Instagram hit 100 million users.

Wow! Three guys on a bootstrap budget were able to build a solution entirely on a public cloud. Imagine trying to scale that fast in a brick-and-mortar data center. In a physical data center, they would never be able to buy hardware fast enough to keep up with the skyrocketing growth. In fact, one could argue that if it were not for the cloud and the on-demand and auto-scaling capabilities, they may never have been able to achieve this success because they would have likely experienced outages as they ran out of capacity.

This story highlights the power of on-demand compute resources. These talented engineers were able to build an amazing, highly scalable architecture in a short amount of time. They did not have to manage data centers or networks or procure, install, and manage hardware. Instead they focused on application architecture and the user experience, two things they excelled at. For start-ups, the cloud is a no-brainer. For companies with an existing data center, it is more of a challenge, which leads us to our next case study.

Established Company Case Study: Netflix, Shifting from On-Premises to the Cloud

Netflix is an industry leader in streaming video content over the Internet. In 2009, 100 percent of all customer traffic was run through Netflix's own data center. By the end of 2010, much of that same traffic was running on AWS, Amazon's public cloud solution. Netflix's goal for 2013 is to have at least 95 percent of all services, including operational services, not just customer traffic, running in the cloud. On its technology blog the company stated its reasons for shifting to the cloud. The enormous amount of incoming traffic required it to rearchitect its solution. It decided that it would rather focus its engineering efforts on building and improving the business applications (Netflix's core competencies) and let Amazon focus on the infrastructure (AWS's core competency). Netflix also spoke about how challenging it was to predict traffic. Companies building on-premises solutions must buy excess capacity to handle spikes. That becomes a great challenge when traffic is not predictable. Netflix felt it was advantageous to leverage the public cloud's on-demand resources and focus on building in auto-scaling capabilities to ensure that it could consume compute resources at the same rate of its incoming traffic. According to the Netflix technology blog on December 14, 2010:

> Cloud environments are ideal for horizontally scaling architectures. We don't have to guess months ahead what our hardware, storage, and networking needs are going to be. We can programmatically access more of these resources from shared pools within AWS almost instantly.

Netflix also sees leveraging cloud computing as a competitive advantage. The company is able to scale at amazing levels* while controlling costs and reducing the risks of downtime. It also feels that the cloud is the wave of the future and leveraging the cloud will attract the best talent:

> It will help foster a competitive environment for cloud service providers, which will help keep innovation high and prices dropping. We chose to be pioneers in this transition so we could leverage our investment as we grow, rather than to double down on a model we expect will decline in the industry. We think this will help differentiate Netflix as a place to work, and it will help us scale our business.

Now we have discussed a start-up success story and an established company success story in the cloud. Let's take a look at how the government is leveraging the cloud.

*As of November 2012, Netflix accounts for 29 percent of all Internet traffic in North America.

Government Case Study: NOAA, E-mail, and Collaboration in the Cloud

The National Oceanic and Atmospheric Administration (NOAA) moved to a cloud-based e-mail solution—Google's Gmail—at the beginning of 2012. NOAA is a federal agency with over 25,000 employees whose mission is to understand and predict change in climate, weather, oceans, and coasts. NOAA has employees working in all conditions such as in the air, on land, and on sea. The employees rely heavily on Internet-connected devices and collaboration with team members and other agencies. To enable efficient e-mail and collaboration capabilities, NOAA chose a cloud-based solution that includes e-mail, instant messaging, videoconferencing, shared calendars, and shared documents. Migrating to these cloud services cut NOAA's costs in half and removed the burden of managing software and hardware updates in a highly distributed and device-heavy environment. NOAA's management staff claims that the cloud-based e-mail and collaboration tools are faster and easier to deploy than the on-premises solutions and the services themselves were more modern. Moving its e-mail and collaboration services to the cloud created great business value by delivering a better overall service at half the price with less effort.

We have discussed success stories in the private sector and the public sector. The next case study is an amazing story about a presidential campaign that built a billion-dollar e-commerce site overnight.

Not-for-Profit Case Study: Obama Campaign, Six-Month Shelf-Life with One Big Peak

Very rarely does one see the type of requirements that the Obama campaign's technical team was faced with. They very quickly had to build a suite of applications including an e-commerce fund-raising platform capable of managing over $1 billion that would run for only six months, have an enormous spike on the last few days, and then back everything up and go away. The team relied heavily on cloud computing solutions and used services from every service model (SaaS, PaaS, and IaaS). The team cited reasons like lower costs, speed to market, on-demand resources, and scalability as some of the reasons for its decisions. Its phone-calling application scaled to 7,000 concurrent users as it peaked on Election Day. The team spent roughly $1.5 million on web hosting and web services, but the amazing statistic is that over $1 million of that was for an on-premises hosting company that managed some of the social media and digital advertising while the rest of the 200-plus applications ran on less than $500,000 of cloud infrastructure and services.

Summary

Cloud computing has evolved from many decades of computing. Cloud computing is the biggest technological shift since the birth of the personal computer and the broad adoption of the Internet. Cloud computing is still in its infancy. Early adopters were mainly start-ups, small businesses, and risk-taking enterprises. As 2012 closed out and the year 2013 began, cloud computing has become widely accepted and enterprise budgets for cloud computing initiatives are growing at enormous rates. As with anything new and immature, cloud computing is still lacking in standards and best practices. The cloud vendors have occasional outages but their overall performance has improved over the years as their products and services mature. Incredible success stories like Netflix and Instagram are becoming more common each year. Enterprises are shifting dollars away from commercial software licenses and hardware investments in favor of a variety of cloud services across all three service models. The secret to success for enterprises will be picking the right cloud solutions to solve the right business problems. Understanding the three cloud service models—SaaS, PaaS, and IaaS—is crucial for enterprises to make the right investments in the cloud.

References

Barr, J. (2012, November 16). "AWS in Action: Behind the Scenes of a Presidential Campaign." Retrieved from http://aws.typepad.com/aws/2012/11/aws-in-action-behind-the-scenes-of-a-presidential-campaign.html.

Barton, M. (2012, April 11). "Cloud Lessons? Instagram Keep-It-Simple Mashup Approach." Retrieved from http://www.wired.com/insights/2012/04/instagram/.

Ciancutti, J. (2010, December 14). "Four Reasons We Choose Amazon's Cloud as Our Cloud Computing Platform." Retrieved from http://techblog.netflix.com/2010/12/four-reasons-we-choose-amazons-cloud-as.html.

Cutler, K. (2012, April 9). "From 0 to $1B in Two Years: Instagram's Rose-Tinted Ride to Glory." Retrieved from http://techcrunch.com/2012/04/09/instagram-story-facebook-acquisition/.

Fingas, J. (2012, November 8). "Sandvine: Netflix up to 29 Percent of North American Internet Traffic, YouTube Is Fast on the Rise." Retrieved from http://www.engadget.com/2012/11/08/sandvine-netflix-29-percent-of-north-american-internet-traffic/.

Gallagher, S. (2012, November 20). "How Team Obama's Tech Efficiency Left Romney IT in Dust." Retrieved from http://arstechnica.com/information-technology/2012/11/how-team-obamas-tech-efficiency-left-romney-it-in-dust/.

Hoover, J. (2011, June 9). "From Ocean to Cloud: NOAA Goes Google Apps." Retrieved from http://www.informationweek.com/government/cloud-saas/from-ocean-to-cloud-noaa-goes-google-app/230500174.

Rush, K. (2012, November 12). "Meet the Obama Campaign's $250 Million Fundraising Platform." Retrieved from http://kylerush.net/blog/meet-the-obama-campaigns-250-million-fundraising-platform/.

Thibodeau, P. (2012, September 11). "Netflix Guts Data Center in Shift to Cloud." Retrieved from http://www.computerworld.com/s/article/9231146/Netflix_guts_data_center_in_shift_to_cloud.