

# Chapter 4

## IT Infrastructure: Hardware and Software

### LEARNING TRACK 3: CLOUD COMPUTING

Copyright Kenneth C. Laudon, 2009

Stern School of Business  
New York University  
44 West 4th Street  
New York, New York

Azimuth Interactive Inc.  
23 North Division Street  
Peekskill, New York 10566

Permission granted by the author.

#### Introduction

#### 1.0 Cloud Computing: Getting it Straight

#### 2.0 Leading Examples and Services

#### 3.0 Cloud Computing: Related Concepts

#### 4.0 Business Pros and Cons of Cloud Computing

Cloud computing is a vision of computing in the 21st Century in which most computer and IS functionality is located on the Internet rather than on your personal computer, iPhone and Blackberry, or corporate data center. Put simply, cloud computing is computing on the Internet. But this simple definition lacks detail and is not very informative even if it works as a grand summary of cloud computing. In fact, this definition is misleading even though it's commonplace in the popular press. Cloud computing is much more than computing on the Internet. Defining "cloud computing" has been as difficult as defining "Web 2.0."<sup>1</sup>

As we see later in this essay, cloud computing is not just a technological juggernaut that ineluctably is rolling over the landscape. It's also a business product line for the largest computing corporations in the United States who have a vested interest in ensuring that cloud computing does in fact become the primary corporate computing model of the 21st Century. Think IBM, HP, Sun, Oracle, Google, Yahoo, Amazon, and many others who plan to benefit from cloud computing. The estimated size of this marketplace is rather uncertain, ranging from \$80-\$160 billion in 2012, to as much as \$1 trillion in 2020.

1. In May 1997, NetCentric tried to trademark "cloud computing" but later abandoned it in April 1999. In April 2001, the New York Times ran an article by John Markoff that used the phrase "cloud' of computers". In August 2006, Eric Schmidt described Google's approach to its developing online software applications (SaaS) as "cloud computing." <http://www.google.com/press/podium/ses2006.html> Also in 2006 Amazon included the word "cloud" in its announcement of Amazon's EC2 (Elastic Cloud Computing). At that point the term became a quick reference to whatever Google and Amazon were offering-applications online, online development platforms, and online storage. IBM and other "Big Iron" firms, operators of large corporate data centers either self-owned or owned by clients, have expanded the term to include a wide assortment of data, resource, and center management services and software (Chris Sears, 2008).

Cloud computing is fashionable. Cloud computing is also a very imprecise slogan that gets mixed up with other social values like efficiency, productivity, and "Green computing" which sends some acolytes into a religious state, while others are skeptical, like Larry Ellison, founder and CEO of Oracle Computing, who quipped in November 2008:

"The interesting thing about cloud computing is that we've redefined cloud computing to include everything that we already do. I can't think of anything that isn't cloud computing with all of these announcements. The computer industry is the only industry that is more fashion-driven than women's fashion. Maybe I'm an idiot, but I have no idea what anyone is talking about. What is it? It's complete gibberish. It's insane. When is this idiocy going to stop? Everything is cloud computing these days and we've been doing it for years."

Shortly thereafter, Oracle launched a major cloud computing effort, and placed a prominent link on its Web page to this new effort.

Cloud computing promises a great deal, and at times these promises may overreach its realistic capabilities, and underestimate the risks and hence total life cycle costs.

### 1.0 Cloud Computing: Getting it Straight

Cloud computing is much more than just computing on the Internet—that happened already in the 1990s if not before. The reality is that cloud computing involves a number of different capabilities.

Figure 1-1 provides a graphical and more robust definition of cloud computing. In this view, cloud computing involves four different capabilities: applications servers, storage servers, platform services, and management services. Each of these capabilities plays a role in defining cloud computing. In the marketplace, providers of cloud computing generally offer these four different capabilities.

**Applications servers:** cloud computing involves making standardized software applications available to users over the Internet. Rather than purchase

Figure 1-1 Cloud Computing

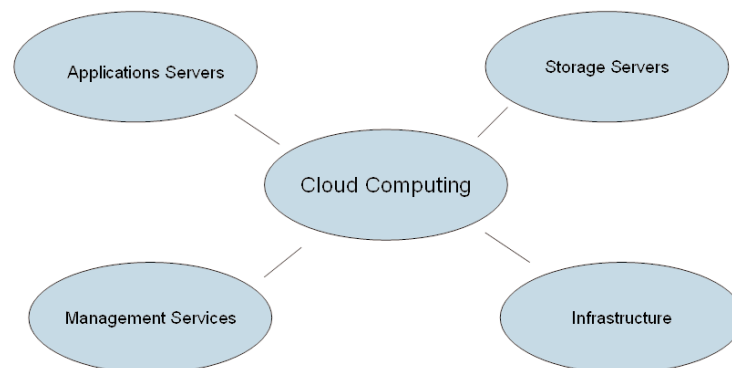


Figure 1-1  
Cloud computing involves the delivery over the Internet of four computing and IS capabilities: application servers, storage servers, management services, and infrastructure

software or build a software development platform, customers can use applications running over the Internet.

Application servers refers to specific server machines dedicated to that function, or multiple virtual machines, and/or application server software operating on numerous machines as demand dictates. Cloud computing also offers software developers the ability to create new customized applications which run on the cloud application servers. These customized applications are built either on the customers' computers and then uploaded, or using development tools on the cloud application server itself. The leading example of standardized applications running on the cloud is Salesforce.com which provides customer relationship management software online to thousands of business firms around the world.

Salesforce.com is an example of Software as a Service (SaaS), a concept which is intimately linked with cloud computing. . Other examples of SaaS provided by application servers are Google Apps (office applications on line), Google Sites (a online collaborative environment), and similar services offered by MicrosoftLive for small businesses. Other successful applications can be found easily on the Internet for video conferencing, , IT service management, accounting, IT security, web analytics, web content management, human resources, and e-mail.

Examples of entire application development platforms on the Web include Amazon's Elastic Computer Cloud (EC2) which provides a virtual programming environment; Google App Engine; Microsoft Live Mesh; and Sun network GRID. While each of these provide "platform as a service," they offer different services and capabilities.

In cloud computing, thousands and even hundreds of thousands of computers are located in cloud data centers. The largest application server installation is arguably Google which has an estimated 500,000 simple PC like blade servers to perform billions of searches a day.

There are several advantages for using cloud applications:

- Reduces the need for customers to purchase computers to run applications, along with the telecommunications and staff required.
- Eliminates software maintenance, upgrades and support
- Provides inexpensive scalability as businesses grow (or shrink)
- Enables "on demand computing" where customers are charged only for the capacity they use.
- Reduces the initial cost of purchasing software applications

The major disadvantages of cloud applications are:

- Reduced control by the firm over the functionality of the software
- Introduces dependency on an external firm
- Poses security risks for truly proprietary corporate information
- Introduces a new recurring expense beyond the control of the customer
- Exposes the firm to future switching costs if the relationship does not work out

**Storage Servers:** Applications need data, usually lots of data. In the past, firm data was stored locally on large disk drives, where it was used to support

local application servers. With cloud computing, the location of data changes from the corporate hard drives to Internet storage servers located in very large, energy efficient data centers available on the Internet. One of the more successful cloud storage services is provided by Amazon's S3 (Simple Storage Service). S3 is a part of Amazon's larger cloud service Elastic Compute Cloud (EC2). S3 provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the Web. It gives any developer access to the same highly scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites. The service aims to maximize benefits of scale and to pass those benefits on to developers.

Costs for large scale data storage start at 15 cents per gigabyte for the first 50 terrabytes. Telecommunications costs start at 10 cents per gigabyte in, and 17 cents per gigabyte out. The cost of storing 1 gigabyte of information in a typical corporate storage area network (a managed array of hard drives) is about \$2.00! One user of Amazon S3 is NASDAQ (the National Association of Securities Dealers Automated Quotations) for its Market Replay application. Market Replay records every transaction price on the New York Stock Exchange and NASDAQ, along with ten minutes of transaction data on the stock. The purpose is to ensure to regulators that NASDAQ trades were executed at the best possible market price in a given environment. This involves tracking trillions of transactions. To store all this information NASDAQ uses Amazon's S3, uploading 30 to 80 gigabytes a day.

**Infrastructure:** In order to provide application and storage services online, cloud providers construct elaborate and large data centers—the infrastructure of cloud computing. This involves acquiring physical buildings for its computers, securing sufficient power resources for the computers, establishing fail-safe telecommunications links with the Internet, and providing data security services. Data centers used to be located in the same physical building as corporate headquarters, or in an adjacent community or state. This was necessitated by the expense of move large volumes of data over long distances, security issues and beliefs, and strategic considerations. Being "close by" is still a powerful phenomenon for many large institutions and firms, but the Internet has made moving data over long distances nearly as inexpensive as moving data to a corporate basement data center or across town.

One of the largest cloud computing infrastructures to date is supplied by IBM's Blue Cloud initiative. By 2008 IBM announced it was building "cloud computing services" into nine IBM data centers around the globe, each costing about \$350 million each. One data center in Singapore was specifically designed to be a cloud computing center. What's inside a Blue Cloud data center? Not much more than what's inside any IBM corporate data center. Thousands of Linux based blade servers designed to operate with minimal power and air conditioning requirements; some of IBMs new Series 7 mainframe computers that themselves run thousands of instances of various operating systems. Using proprietary software, these computers can be dynamically linked together for periods of time to boost computing power to trillions of computations per second (compared to a single PC capable of 3 billion instructions per second). Because these capabilities already exist in the data centers, it's difficult to understand "what's new" in the Blue Cloud initiative. What is new however is offering these services on a "demand computing" basis, and

charging customers for actual capacity and service used rather than for a fixed price. Instead of buying software licenses, for instance, customers can be charged on the amount of data or time used to process data with the software.

To ensure the operation and survival of the hardware, air conditioning systems, electrical distribution and fail safe backup systems, fire protection, and physical security round out the main components in a cloud data center. Air conditioning and computer power consumption are the leading cost factors in operating a cloud data center. IBM has taken a number of innovative steps to reduce the power consumption of its computers by 40% when compared to standard, stand alone servers, and reduced the floor space required by 50%. Cloud computing is "green computing," and much more sparing of electricity than millions of PCs on corporate desktops.

The technologies being offered to the community by IBM's Blue Cloud centers include IBM Rational software development tools, WebSphere Application Server software and DB2 database software running on IBM System x, System p and BladeCenter servers. IBM Tivoli systems management software will manage the cloud computing environment. In other words, the centers built by the large "Big Iron" manufacturers are typically populated by their proprietary software solutions even if these solutions run on Linux servers using "open source" operating systems.

**Management Services** Cloud computing is about more than hardware, software and physical buildings. It's about management of all of these elements, and in addition, other value added services that roughly fit under the rubric of "management services." For instance, a group of managers and lawyers will be needed to arrive at service level agreements with customers; another group of managers will focus on the maintenance and repair of the cloud components.

Managers of data centers provide to their customers the following services:

- System optimization
- Consulting services
- Hardware and software maintenance
- Continuity planning
- Physical floor management and planning
- Security management
- Infrastructure support (electrical and air conditioning)
- Pricing
- Accounting

An IBM Blue Cloud data center employs over 300 onsite employees, in addition to hundreds of off-site employees, who are required to administer the center and interact with clients.

## 2.0 Leading Examples and Services

Cloud computing is in its infancy. It tends to leverage the resources of very large hardware manufacturers, but even service oriented firms like Amazon have the potential to play a large role in cloud computing. Amazon's example may open the door to even more cloud services being provided by very large firms with excess computing capacity on their hands.

**TABLE 1-1 CLOUD COMPUTING VENDORS AND SERVICES**

FUNCTIONAL AREA	BUSINESS PROCESS
Amazon Elastic Computer (EC2 and S3)	Broadly based cloud environment providing computing and storage; various Web APIs; AMI (Amazon machine image) programming environment. Utilizes Amazon's spare computing capacity required by its retail operations.
IBM Blue Cloud	Extensive cloud services from computing and storage to proprietary software
Google App Engine and Google Apps	Python programming environment for developing new applications development platform; data storage; office productivity applications SaaS. No value added services.
Microsoft Live Mesh	Small business storage, and Web based applications for collaboration; no value added services.
Sun Network.com	Broadly based cloud environment providing computing and storage; various Web APIs.
Oracle Cloud Products	Oracle Database 11G, Oracle Fusion Middleware, and Oracle Enterprise Manager available to run on the Amazon EC2 and S3 platforms.
Mozy	One of a large number of online backup and file transfer services available to individuals for free, or a nominal fee. Others include Carbonite, Xdrive, and ATT's Data Vault.

Table 1-1 describes some of the large vendors providing cloud computing services today.

### 3.0 Cloud Computing: Related Concepts

**TABLE 1-2 CLOUD COMPUTING: RELATED CONCEPTS**

FUNCTIONAL AREA	BUSINESS PROCESS
Cluster computing	Computers linked together generally in a local area network (same building) to provide redundancy, and dedicated to a small number of tasks.
Grid computing	Linking together of many remote computers and using their spare capacity to solve large computational problems, or to process bits and pieces of larger programs.
Utility computing	The provision of computing services over a network to remote customers much like an electrical or telephone utility works. This would include billing users based on the amount of service consumed, plus additional charges to pay for the infrastructure. First proposed by Leonard Kleinrock in 1969.

There were many precursors to cloud computing that at times become confused with cloud computing. Table 1-2 describes some of these other concepts and phenomenon.

#### **4.0 Business Pros and Cons of Cloud Computing**

Cloud computing offers business firms the potential for significant cost reductions for both software and hardware. Software should be less costly because its cost can be spread efficiently over a large number of users. Hardware should be less expensive because large data centers can keep their computers operating at much higher capacity levels than a single firm. The single firm is no longer required to purchase enough computing power for its peak loads, and instead it can count on the cloud center to handle whatever peak loads come along. From a cost and flexibility perspective, cloud computing makes a lot of sense.

On the other hand, firms tend to hang on to their mission critical and "strategic" applications which form the foundation of the business. The reason is that most managers are still wary about Internet cloud security, both the telecommunications links and the stored data at cloud centers. What happens if the cloud center goes down? What happens if the communication links are severed? For instance, on February 15, 2008, Amazons S3 Cloud Storage system went down for three hours. Many small and medium businesses were knocked out of service. This highlights one of cloud computing's biggest risks and nightmares. The question is: will the track record of cloud computing be as good as the track record of most corporate data centers (or just as bad)?

For this reason, large corporations have so far only off-loaded their non-critical and non-strategic applications to the cloud. Other managers are concerned about becoming a "captive" of proprietary software offered by some vendors such as IBM and HP. What happens if you want to switch to another provider?

In the years ahead, confidence in Internet security will probably grow, and firms will feel more confident about placing more mission critical systems on the cloud. However, some obstacles-like vendor dependency-- will not be so easily overcome. Firms will develop mixed strategies. They will retain their strategic defining systems and their mission critical systems, while offloading to the Cloud more routine administrative systems. The cost of buying excess capacity will be considered far smaller than the risk of a cloud failure. Small

and medium sized businesses may well be the largest users and benefactors of cloud computing. They are sufficiently small that they have no choice but to build scale by using cloud computing facilities.

### References

Sears, Chris. "Who Coined the Phrase Cloud Computing." IT Management and Cloud Blog, 2008. <http://www.johnmwillis.com/cloud-computing/who-coined-the-phrase-cloud-computing/>.

Ricadella, Aaron. "Computing Heads for the Clouds." BusinessWeek, November 16, 2007.

Lohr, Steve. "Cloud Computing Gains Steam With New IBM Gear." The New York Times, April 28, 2008.

Wall Street Technology. "Cloud Computing Begins to Gain Traction on Wall Street." January 6, 2009.

Worthen, Ben. "Is Amazon's Crash a Giant Crash for Cloud Computing?" February 15, 2008. [blogs.wsj.com](http://blogs.wsj.com).

Buyya, R., C. Yeo, and S. Venugopal. "Market Oriented Cloud Computing." The University of Melbourne, Australia. Keynote paper, 2008.