Cloud Computing – Storm Clouds or is it Smooth Flying?

By Cary Whitaker

East Carolina University

Abstract

It once was that applications had to be downloaded and/or installed on one's own computer in order to function. As computer communication technology has progressed, there has been an increasing movement towards web-based applications, where the actual application functions takes place on a distant rather than local computer. With the increasing understanding of and need for network and personal security, where does cloud computing fit in? Can a system designed to require so much transport of possibly sensitive information be made to comply with current security principles and needs, or is cloud computing destined to be a niche market? I will explore the known security risks of cloud computing, and a few of the methods being developed to deal with them.

The World Wide Web has often been characterized as a cloud. Considering its nature, the ability to access information from the world over, and how ill understood the process by most Internet users is of releasing information to the internet and receiving information in return, the cloud perfectly describes the World Wide Web; a cloudy, nebulous object with hazy, permeable, and undefined borders. However, within the last few years, the cloud concept has taken on a somewhat more defined context; it's called cloud computing, and according to whom one listens to, it's the next big thing, the next technological revolution, or a vastly overrated business model.

There are many definitions of cloud computing, it pretty much all depends on how expansive or how narrow the person wants to make it. Infoworld offers 3 definitions: on the narrow side, cloud computing could be said to be virtual servers available over the internet. To the other extreme, it can be said to be anything used that is beyond the perimeter firewall of your network (Knorr and Gruman). Probably the best definition for the current concepts of cloud computing lies in-between: Cloud computing encompasses any subscription-based or pay-per-use service that, in real time over the Internet, extends IT's existing capabilities. Note that while certain free services, such as document conversion, technically meet the definition, they lack the revenue stream that is driving the industry.

Surprisingly, cloud computing is an old concept, at least in terms of the pace of technology. In 1961, John McCarthy theorized that someday computer technology may evolve to the point where computation power could operate on the utility business model, like electricity or water. The idea became popular, but died down in the early '70s because the hardware infrastructure needed for cloud computing was nowhere in sight yet. The concept renewed itself in the 2000s, taking its name from the use of cloud symbols to represent the internet in schematic

network diagrams (Sansom) . While there have been snippets of the idea over the years, the concept has only recently begun gaining traction as an independently viable industry.

There are many categories, or subcategories, of cloud computing. They are as unsettled as the exact definition of cloud computing itself, but we can start with SaaS and utility computing. Software as a Service (SaaS) allows clients to access and remotely utilize an application with their browser. Webmail is a good examples; the emails and interface are not kept on the user's home system, but rather on the remote service's server. Google Apps and Zoho Office are subscription-based suites that offer multiple applications to their subscribers. Utility computing is more in line with what was envisioned nearly 50 years ago. Amazon, like many online services, has IT capacity that far exceeds its usual daily needs. This is because occasional usage spikes occur where they need that extra capacity. Rather than letting all the excess capacity stand idle when it's not needed, Amazon found a way to profit from it. Amazon Web Services utilizes Amazon's excess capacity. It offers storage and computation capacity to its clients. Web Services in the cloud has been around awhile. It offers a functionality rather than a full blown application, such as the credit card processing services. Platform as a service offers a developmental environment to its users, where they can develop and market their own applications. The exploration and expansion of cloud computing has lead to some of these categories sometimes bleeding over into one another.

No market grows without demand for its product. So what is it about cloud computing that is appealing to people and companies? The main attraction are the cost savings. Cloud computing can offer cost savings of up to 50% for small businesses IT investments, although most will never reach that percentage. Still, 30, 20, or even 10% IT cost savings is nothing to sneeze at. Companies spend what they see as an excessive amount on the maintenance of their

IT infrastructure. IT staff, software licenses, software integration, troubleshooting, all take time and money. For new companies, the initial investment necessary for a working IT infrastructure is prohibitive, and a big barrier to starting a business, especially ones requiring a large IT infrastructure. With cloud computing, companies trade the intimidating up-front capital investments in hardware, networking, and software licenses for a much more modest up-front investment and operating costs that are based on monthly fees or pay-as-you-go utility fees (Korzeniowski). This significantly lowers the barrier to entry to the marketplace for new companies and oftentimes new products of existing companies, for those that have bought in to cloud computing anyway. Lowering the barrier to entry is almost always a desirable outcome. It promotes greater competition and innovation.

The next benefit is convenience. Before recent times, software needed to be purchased (or pirated), licensed, updated, and maintained by the user or the user's company. With cloud computing, none of that is necessary for the end user, because the headachy aspects of IT are taken care of by the cloud computing service. It licenses, maintains, and updates the software instead of the end user. The end user just rents access to the cloud in order to use the application(s) it needs, or it comes part and parcel of the service that the user has subscribed to. Brian Hayes, senior writer for American Scientist, stated "In contrast, software as an Internet-based service can be developed, tested, and run on a computing platform of the vendor's choosing. Updates and bug fixes are deployed in minutes."(Buck) For anyone that has performed an update or patch of software installed on individual workstations in a network, this probably sounds like some kind of fairy tale. This kind of convenience is difficult to argue against, especially for users that are not technically inclined (which are most of them these days) or companies just starting out or needing to tighten their belts.

Cloud computing services also take responsible for software compatibility issues. When installing a new program, care must be taken to be sure that the new program is compatible not only with the hardware, but also with the underlying operating system and some other programs already installed. When a software conflict does occur, much time and frustration must be invested to figure out the other offending party, so that the problem can be identified and resolved. Cloud computing services are responsible for this, so a company does not need to devote resources or lose time to software compatibility issues. Sadly, these conveniences usually means that a company does not need, or thinks that it does not need, as many IT personnel as before, so some of the projected cost savings from above would from slashed IT jobs.

The convenience does not extend only to SaaS. Online storage is another of cloud computing's key selling points, and perhaps its most widely accepted function currently. With online storage companies and individuals can backup their data when they subscribe to a storage provider. Mozy and Carbonite offer individuals unlimited backup services for less than five dollars a month (Rash). Each offer programs that back up your hard drive, then afterwards, when your computer is not in use, it backs up files that are new or have been changed, there is very little user interaction needed after the initial download and installation. This keeps the program from slowing an individual's computer down or using up too much bandwidth. Both also encrypt the data before transport, and while it is stored in the data center. It is important that the user keep track of the encryption key or secret password that he opted for in order to successfully retrieve his data when he needs it. While Mozy and Carbonite cater to individuals' and small businesses' back up needs, there are other cloud storage companies that service enterprise level companies. Backup providers like Iron Mountain offer significant scalability and high levels of security and redundancy (data is moved offsite to an underground facility and

mirrored to a second, separate data center). (Iron Mountain) For businesses, especially enterprise level businesses, need to securely backup their data anyway in case of disaster or catastrophic data loss, online storage is especially attractive. Backup can be fully automated, or it can be configured for bandwidth limits and known off peak times. Customers don't need to invest in the physical security and safety of their business backups, the online storage company takes care of that, as well as providing the security measures for confidential data transit. Online backup data is accessible from any highspeed-internet capable location, further adding to the online backup's advantage over traditional data backup methods.

With all of these advantages going for it, why is there a high degree of squeamishness many companies have about using cloud computing? There are a few outstanding issues that need to be resolved before cloud computing becomes accepted across-the-board in industry. The first is privacy. There are few companies, if any, that do deem some of their data confidential, whether because of proprietary concerns or because of legal concerns and laws.(Davis and Kennedy) There is some question of just who owns the data stored out on the cloud. Is it the customer that put it there, or the service that houses it? Also, what are the legal implications for storing personal or business information on a outside party's service? For instance, last year a judge ruled that email stored upon a third party platform negates the Fourth Amendment requirement for notification usually afforded to people whose information is seized; instead it is sufficient to serve the ISP with the warrant (Larkin). The emphasis on physical location of the information, rather than what the information is and who it belongs to, is outdated in concept, but still alive in current law. It is actually illegal for U.S. or European resident companies to store certain information outside of their respective geographic areas (the U.S. for the former, and the European Union for the latter). Also, HIPAA regulations must be met for cloud computing

companies to be legally authorized to handle medical information. These issues must be resolved before cloud computing can fully flourish.

Reliability is also a great concern of cloud computing skeptics. Simply put, it is a lot of companies and/or people putting their eggs into someone else's basket for them to look after and care for. If they drop the basket, how will the customers pick up the pieces? Ideally, should a catastrophic event occur, the service would be able to reassemble the information or restore it from backups. Last year though, Magnolia, a social bookmarking service, completely lost its data, even the backup data. Service outages can be crippling to a business, and cloud computing encourages condensed points of failure; if a cloud computing service goes down for whatever length of time, all of its customers are also bereft of its IT services. Even Google online services and Amazon's Elastic Compute Cloud (EC2) web service have both been struck by outages. The lack of control over the IT infrastructure puts off many executives from investing and trusting cloud computing.

By far the greatest concern of cloud computing is its security. Information is being sent out over the internet, out of the safe (relatively) confines of the business's firewall. A outside party is being trusted with your business's confidential information. Either concept alone is normally enough to give an IT staffer stomachaches, both together and many need antacid. Encryption, both in storage and in transit, is both the answer and the question. For online storage purposes, the problem is not so big; strong encryption will see it through. But for SaaS, encryption presents a conundrum. Data sent to the software program must be able to be understood by the program. Encrypted data is not understandable until it's been decrypted. So, either the user must send the data unencrypted, or the cloud service must decrypt the data that is sent for the application to be able to use it; either way, it takes the security and privacy of the

information involved out of the end user's hands. Encryption in the cloud is not the only problem that must be dealt with. For example, last year, three computer scientists from the University of California, San Diego and one from MIT found a flaw in Amazon's EC2 service. Hiring virtual machines (VMs) from Amazon to act as victims, they noted that if multiple VMs are bought in a small span of time, the VMs would have similar IP addresses, which indicated that they might be on the same server. Realizing this meant that they needed to hire the attack VMs at the same time as they victim would hire his VMs. One way they suggested to achieve this in a real world setting would be to bombard the intended victim's website with requests, encouraging them to hire additional VMs to handle the sudden spike. This is greatly eased if the hiring of such VMs for spikes is automated. When the victim hires the additional VMs, the attackers also hire them from the same service. The researchers found that the attack VMs ended up on the same server 40% of the time. The researchers then used the attack VMs that were on the same server as the victim VMs to monitor the victim VMs' use of computing resources. The researchers claim that though they did not steal any data, possibly for legal reasons, they could have. Amazon claims that the possibility of data theft was only theoretical, but they have since taken measures to prevent this kind of attack, though they did not state exactly how. This did illustrate how a fairly simple attack could conceivable have compromised a cloud computing service. (Ristenpart, Tromer, Shacham, Hovav, Savage) Although this attack has a fairly simple fix (seeing to it that VMs hired close together don't get assigned to servers in a predictable manner), it does conjure up questions of how many other simple attack vectors current cloud computing technology has opened up, and how soon and how well they can be closed.

Some of these problems already have solutions, though most are not ready for deployment yet. IBM researchers have come up with a way to scan virtual machines entering a

cloud to check for integrity and how it operates. New encryption technologies are being explored that would allow information to be sent and used by an online application while staying encrypted, such as Craig Gentry's "Ideal Lattice." (Talbot) It's ironic, but the recession of the last few years may have played a part in cloud computing's growth, as organizations look for ways to slash budget costs while sacrificing as little functionality as possible. That's why even the concerns listed above will not make institutions ignore the vast potential savings possible with cloud computing. Last October, the Los Angeles city council voted to outsource its employee email services to Google, Inc. (and Computer Sciences Corp.) for five years (Korzeniowski). In order to land the contract Google had to allay the city council's, LAPD's, and city attorney's security concerns. Google offered/were required to: (1) fingerprint all employees working on the project, (2) encrypt the data in transit, (3) shard the data once within Google's cloud (this means that in order to retrieve the information, the user needs the correct application and encryption key to make it readable again), (4) agree to store all the data within the U.S., and (5) limit access to the data to the employees for Google and CSC that met the city's security clearance requirements. Even with this, the contract was approved only once an amendment that would require Google to compensate the city in the event that the Google system was breached and city data exposed or stolen was added to the contract. All these measures were outside the normal range of assurances that Google offers, but they did not cost L.A. extra (L.A. actually got a 40% discount on the whole deal), possibly because of the marketing coup this meant for Google's cloud services. Why did L.A. require such additional measures? With this contract, L.A. became one of the first public concerns to buy into cloud computing; as such, they wanted higher assurance than would typically be offered to a private concern. With most companies, data security, storage, computer services, etc. are part of the cost

of, and part of the means to, the company purpose and products. For cloud computing services, the security, storage, and services *are* the companies' purpose and products. This is an important distinction. Companies invest in their primary purpose and products, and work to minimize any operating costs that they can (of which IT is usually a prime victim); but since IT infrastructure and data caretaking are cloud computing's purpose and their primary products, they will invest more, need to invest more, in security and privacy in order to have a viable product. Even some enterprises leery of cloud computing have gotten into the act, by investing in private clouds, a cloud computing service for their company alone. That way, they get some of the cost savings that cloud computing offers, while maintaining full control and possession over their data and operating within their security parameter.

In conclusion, despite the problems and suspicions, cloud computing will change the internet as we know it. The change may not be as all-encompassing or revolutionary as its advocates believe, but the lure of being able to use applications and computer capacities without having to invest in anything but browser capable computers, a reliable high speed internet connection, and the cloud computing service subscription or rate, will drive cloud computing from today's buzzword to tomorrow's mainstay. The important thing to remember is that cloud computing is a nascent technology, and as such will not be dismissed simply because of its immaturity. As it matures, problems will be solved and give way to new problems and then to new solution, but the overall technology points towards increased stability and security.

References

Knorr, E, & Gruman, G. (n.d.). What Cloud computing really means. *Infoworld*, Retrieved from http://www.infoworld.com/d/cloud-computing/what-cloud-computing-really-means-031?page=0,0

Iron Mountain. Online Backup. http://www.ironmountain.com/online-backup/online-backup.html

Sturdevant, C. (2010). Seeding security into the cloud. *eWeek*, *27*(6), 38. Retrieved from Academic Search Premier database.

Larkin, E. (2010). Will Cloud Computing Kill Privacy?. *PC World*, *28*(3), 44. Retrieved from Academic Search Premier database.

Rash, W. (2010). Cloud-based storage done right. *eWeek*, *27*(3), 14-16. Retrieved from Academic Search Premier database.

* Sansom, C. (2010). Up in a cloud?. *Nature Biotechnology*, *28*(1), 13-15. doi:10.1038/nbt0110-13.

Talbot, D. (2010). Security in the Ether. *Technology Review*, *113*(1), 36-42. Retrieved from Academic Search Premier database.

Hall, M. (2009). PIONEERS OF THE PRIVATE CLOUD. *Computerworld*, *43*(35), 14-19. Retrieved from Academic Search Premier database.

Hesseldahl, A. (2009). Forecast for 2010: The Coming Cloud 'Catastrophe'. *BusinessWeek Online*, 9. Retrieved from Academic Search Premier database.

Chura, H. (2009). Cloud Storage. *Time*, *174*(22), 102. Retrieved from Academic Search Premier database.

Korzeniowski, P., & Jander, M. (2009). Cloud Security. *InformationWeek*, (1247), HB4-HB14. Retrieved from Academic Search Premier database.

Villano, M. (2009). Cloud Cover. *Time*, *174*(14), 8. Retrieved from Academic Search Premier database.

Hawthorn, N. (2009). Finding security in the cloud. *Computer Fraud & Security*, *2009*(10), 19-20. doi:10.1016/S1361-3723(09)70131-9.

Rash, W. (2009). Is cloud computing secure? Prove it. *eWeek*, *26*(16), 8-10. Retrieved from Academic Search Premier database.

Buck, S. (2009). Libraries in the Cloud: Making a Case for Google and Amazon. *Computers in Libraries*, *29*(8), 6-10. Retrieved from Academic Search Premier database.

* Davis, R., & Kennedy, D. (2009). WORKING IN THE CLOUD. *ABA Journal*, *95*(8), 31-32. Retrieved from Academic Search Premier database.

(2009). Data in the cloud might be seized by government agencies without you knowing. *Computer Fraud & Security*, *2009*(8), 1. doi:10.1016/S1361-3723(09)70092-2.

* Ristenpart, Tomas, Tromer, Eran, Shacham, Hovav, Savage, Stefan. (2009). Hey, You, Get Off My Cloud: Exploring Information Leakage in Third-Party Compute Clouds. Retrieved on April 17th, 2010, from: http://cseweb.ucsd.edu/~hovav/dist/cloudsec.pdf