Big Data Cloud Computing

Industry and Technology Assessment

7 June, 2012

HTTP: Hunter Thomas Tony Project

Introduction

In the current age, information is treated as precious commodity. As society advances, the demand for cost efficient information storage and processing grows higher. Incentives for independent companies to optimize their IT management and minimize server maintenance costs become greater as usage demands prove to be increasingly unpredictable. Cloud computing offers a promising solution. Cloud computing is a computing model that relies on a large, centralized data center to store and process a great wealth of information. Computing power and storage space are provided on-demand to companies that outsource their IT management to the cloud service provider. The immediate advantage to this computing model is a lower infrastructure maintenance cost. Since companies that use cloud no longer require on-site servers, they eliminate the associated cost in IT management and electrical power. Furthermore, companies adopting cloud technology do not need to oversize their computing resources to compensate for unforeseen usage spikes, as cloud service providers automatically match supply with demand at all times. A centralized data center also complies with "green IT", a federal initiative to promote environmental friendliness in the IT industry. According to Forrester Research, an independent research firm, the global cloud computing market will reach $241 billion in 2020 – more than six times the size in 2010. To confirm this speculation, we conducted a number of interviews with professionals from both management and technical specialties working in cloud-oriented organizations to garner their thoughts on the technology. We also looked into administrative reports, market evaluations, and federal policies for additional perspectives. There is great potential for investments in this field over the next decade. In this report, we seek to present an overview of the current market state and formulate a sensible projection on the course of the industry over the next five to ten years.

Executive Summary

The concept of cloud computing originated in as early as 1960s, when John McCarthy, an American computer scientist, predicted the eventual convergence of computing infrastructure, which would allow a great degree of versatility in the distribution of IT resources in order to meet fluctuating and unpredictable business demands.

After the internet became popular, Amazon modernized their data centers and adopted the cloud architecture, allowing them to outsource computing power to external customers. It launched the service in 2006, and is still known as the Amazon Web Service. Following Amazon, numerous other companies took the initiative to create cloud-based infrastructures.

The main advantage that cloud computing offers is efficiency. Companies using cloud do not need to maintain large data centers to store and manage their information. Instead of leaving 90% of the servers idle in case of a usage spike, the company needs only to maintain enough IT resources to handle the task at hand. If a spike were indeed to occur, the company can immediately demand additional IT resources from its cloud service provider, and relinquish those resources once the spike is over. The company only pays for what it needs at any given time.

Cloud technology completely eliminates the need for companies to have an IT department. Unlike the conventional computing model that companies have historically resorted to, cloud computing revolutionizes the way information is handled. Several levels of cloud service are available for adoption. The most basic level is Infrastructure as a Service (IaaS). In this model, consumers are given full freedom to manage their data on the server. The service provider is only responsible for raw storage, computing power, networks, firewalls, and load balancers. This is often manifested as a virtual machine. The level above IaaS is Platform as a Service (IaaS). In the PaaS model, consumers are provided with an operating system, programming language execution environment, database, and web server. They do not have to be concerned with the cost and management in the hardware and software layers. The highest level is Software as a Service (SaaS). In this model, consumers are given access only to the application software, which can be run remotely from the data centers of the cloud service provider. The provider is responsible for the maintenance and support of the infrastructure and operating platforms. Since cloud service providers specialize in one area, they can provide reliable service at a fraction of the cost. For instance, Dr. Greg Gianforte, CEO of RightNow Technologies, mentioned in our interview that his company was able to achieve 99.99995% reliability – four orders of magnitude better than existing companies with their own self-maintained servers.

Several key technological advances have paved the way for cloud to expand rapidly. The most prominent one is the invention and popularization of the internet, which led to the public availability and exchange of information, online communication, a new venue of advertising and entertainment, and a greater incentive to start small businesses that rely on social networking and electronic transactions. Another great proponent is the invention of smart phones. Smart phones have high bandwidth demands, which resulted in the development of high speed secure data transport layers such as the 3G network. The mobility and versatility of smart phones limit their computing power and storage capabilities, but cloud computing allows all the information handling to occur off-site, pushing the technological barrier past hardware limitations onto network speed. Hardware improvements also made the concept of cloud possible today.

In our evaluations, we determined that cloud has great potential in becoming the next dominant computing architecture. As of now, cloud is still in the early stages of the S-curve. As the technology matures, it would become increasingly ubiquitous in the IT solutions of businesses at all scales, and it will enable data management at an efficiency and reliability never formerly achieved. In addition, cloud will enable other technologies that require a great deal of information to be processed in a short amount of time to grow; this technology has greater implications in science, medicine, military, social engineering, and civil surveillance than any other.

In this report, we will go over the current market state of for cloud computing, some existing applications for cloud, and how cloud handles big data. We will also formulate quantitative projections on both the technical and the business aspects of the industry over the next decade. As additional illustration, we will include a scenario demonstrating a way cloud may be used in our daily lives.

The Development of Cloud Computing

In the last decade, cloud computing has experienced enormous growth because of hardware and software improvements that have increased capabilities while reducing costs.

The major advancements on the hardware side are increased processing power, better and cheaper memory, improvements in connectivity, and superior caretaking of data centers. The raw processing power of CPUs has been growing exponentially in the last 10 years. Fig. 1 shows that a processor from 2002 could handle around one thousand MUCP (one million units of computing power). In 2011, Intel's Sandy Bridge processors were clocked at one million MUCP, a factor of a thousand increase in speed, or roughly doubling speed every year. Intel already has a new generation of processors, nicknamed Ivy Bridge that is even faster.

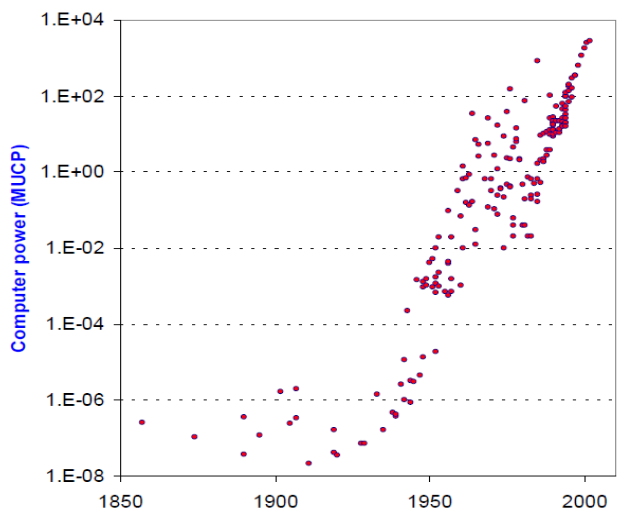

Figure 1

Intel's new Sandy Bridge processors, which were released in 2011, have a computing power of about 1E+06 MUCP, continuing the exponential growth trend shown historically here. (Nordhaus)

Big improvements in memory have also been made. In 2002, it cost around $50 for a 256 MB stick of memory that ran at a speed of 133 MHz. That's almost $1/5MB of memory. Today, it costs about $20 for a 4 GB stick that runs at 1600 MHz and has three times the memory channels (think of this as an additional factor of three increase in speed and efficiency). That's $1/200 MB of memory that is roughly 36 times faster than what was offered in 2002 (jcmit).

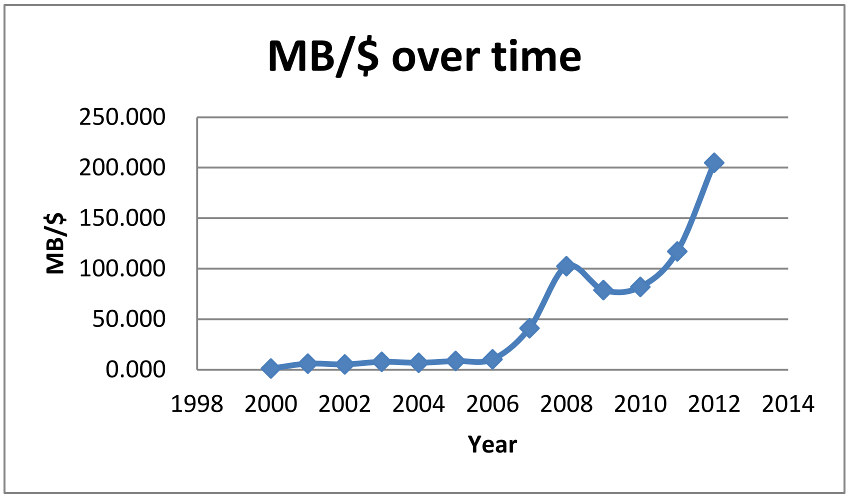

Figure 2

A graph showing the increase in memory per dollar spent since 2000. The first jump after 2006 was when outsourcing memory manufacturing started to affect the market. The decrease in 2009-2010 is due to a jump from DDR2 to DDR3 RAM, a major speed increase (Data from jcmit)

A similar trend exists in the hard drive market. 10 years ago, a 40 GB hard drive cost $150 or over $4 per GB of storage. In 2010, the price of a 1 TB hard drive was $70 or about 8 cents per GB (winchest). Most of these cost decreases can be attributed to the outsourcing of memory manufacturing from the US to various Southeast Asian countries – Thailand is now the world's leading exporter of hard drives.

Server bandwidth has also been an area of major progress. No matter how much faster or more efficient server hardware becomes, if the speed of data transfer does not keep up then it will become a bottlenecking factor. Advancements in fiber optics and the development of technologies such as PCI express (faster ports for data transfer) attributed to a factor of eighty increase in bandwidth speeds between 1997 and 2007 (Server bandwidth). Today's standard of 100 GB/s connections is more than ten times faster than what it was in 2007. Most large server farms run on connections that are multiple TB/s.

The proliferation of server farms (or data processing centers) can be partly attributed to the many hardware advancements that have minimized the cost portion of the business equation. Google has continuously invested into building up a literal army of server farms. It began in 1998 with its one data center at Stanford and now has over thirty significant (tens of thousands of servers in a significant center) data centers worldwide. It takes big bucks to power and cool a server farm. Companies such as Microsoft and Facebook spend millions of dollars per center per year to do both. Innovations, such as Google's plan to build their new centers out in the ocean to take advantage of sea wind power and cold currents, will drive costs down even further and make cloud computing an even more appealing business (justmeans).

Hardware improvements provide cloud companies with the tools they need to grow. Competition in memory and storage markets have also driven down the expenses of building data centers. As the hardware has evolved, so have the software components to the industry. Virtual machines have been one of the biggest software innovations for cloud computing. Before VMware, servers could only run one operating system/platform at a time, and as a result only cater to one customer per server. The advent of virtual machines allowed cloud providers to run multiple platform instances on a single server, thereby allowing each individual server to use its full capacity serving multiple clients. In the old system, the number of servers a data center had was often the limiting factor for the customers it could handle, so there was limited incentive to upgrade the hardware for higher capabilities. Nowadays, cloud service providers demand the highest spec systems so that they can maximize their workloads (Dr. Servan-Schreiber interview).

The proliferation of NoSQL database systems has also been a major advancement for cloud computing. NoSQL refers to a class of non-relational database systems that have been around in theory since the 1960s but only picked up steam in the early 2000s. Traditionally, programmers stored application data in tables and different data formats (such as ints and strings) would have their own column in said tables. The biggest disadvantage of this system is that it limits the development of applications. If the coder decides he needs to be able to handle a new type of data after he defines his data table, he must recode the entire table in order to add another column. In a NoSQL system, this is not a concern because of the way data is handled. Coders are free from worrying about how their programs handle new data. This concept is very important for cloud systems because cloud caters to a wide variety of customers and so they need a framework that supports handling many different types of data.

Major inroads in platform software have also helped cloud computing grow. The Hadoop project and Windows Azure are two standouts that have helped make server performance less reliant on having the most cutting edge hardware systems. These technologies allow large data problems to be broken down and split over many separate processors, then put back together to generate the final solution. By being able to distribute a problem in such a way, server farms can choose to become more cost efficient by having more midrange level hardware that give the same output as a center full of less cost effective high end systems. Hadoop and other offshoots of Google's MapReduce have become an integral part of any company that offers cloud solutions.

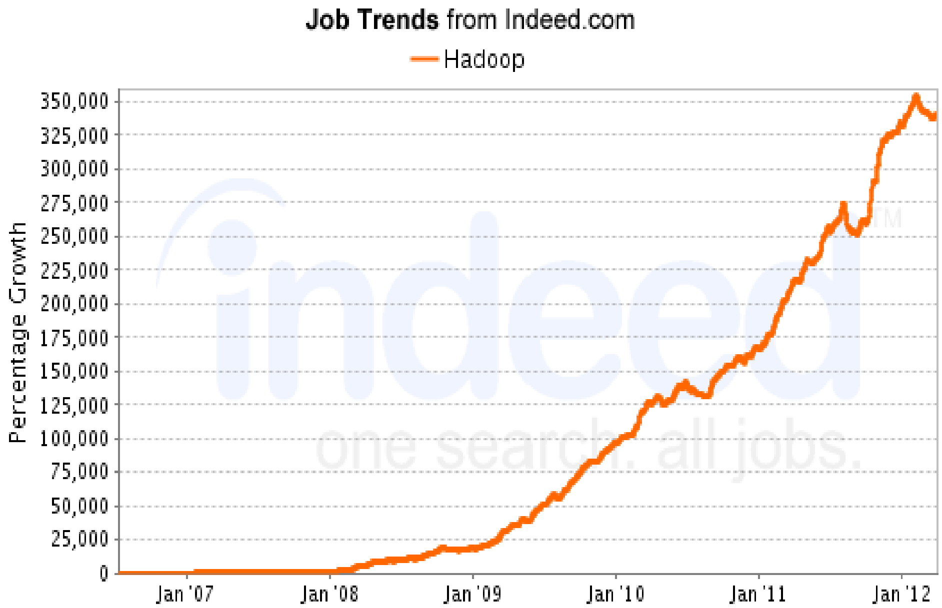

Figure 3

A graph of the percentage growth in job listings for Hadoop related IT positions. The shape of the graph matches the intro/growth portion of the S-curve. Data from indeed.com

The advancements in IT software have primarily helped lower the barriers of entry to the big data cloud computing market. The explosion of cloud computing startups over the last five years would not have been possible without these software innovations that have removed the need for a large capital investment to start a cloud computing business. As a result, the industry has become more competitive and faster developing than it would have, had it been dominated by a few giants such as Google and Amazon.

With the aforementioned technological innovations, big data cloud computing has rapidly increased its capabilities and sophistication. Consequently, it has been able to solve more complex problems over a wider range of applications than it could previously handle. Increased reliability and usefulness, in combination with more competition that drives down costs has made big data solutions an attractive alternative to many traditional methods as well as opening the door for new markets. Because of its powerful viability, big data has a bright future. Market projections show that its market will increase ten-fold in the next 5 years.

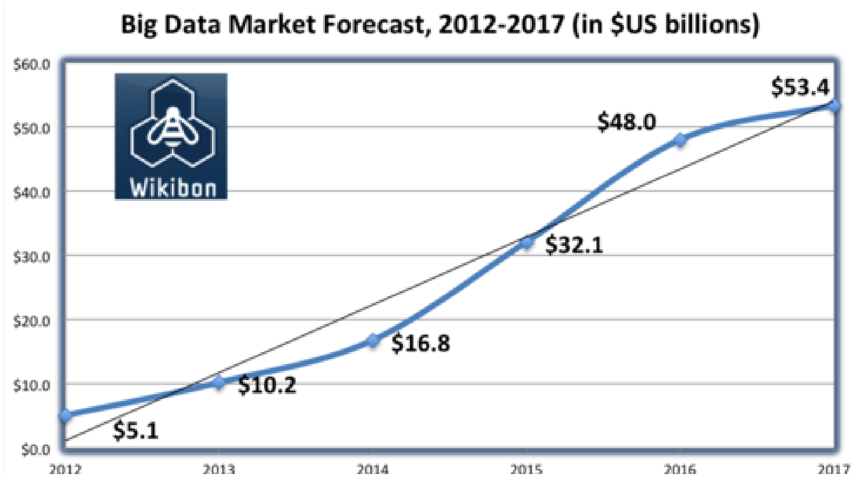

Figure 4

A graph of the projected growth of the big data industry. Graph provided courtesy of Wikibon.

Even more optimism can be drawn from the fact that tech CEOs and IT personnel all agree that big data is still a developing technology.

Currently, the cloud computing market has been dominated by larger companies with large cash reserves. Companies such as Dell and Google have been very aggressive in acquiring smaller companies for their cloud-based intellectual property. As a result, IBM currently owns a 20% market share in big data market (around one billion of the five billion total market revenue in 2011). However, the entrenchment of big companies has not inhibited the growth of cloud computing startups. In fact, the end goal of many current startups is to be acquired and receive a large payout. Low barriers to entry, such as being able to rent servers from Rackspace instead of investing capital in private servers, has been a key factor in helping startups compete with the Amazons and Googles. Smaller companies are able to concentrate their efforts on innovative ideas instead of having to worry over capital expenditures. Splunk is a recent success story that recently held its IPO. Splunk only offers software products and has very few physical assets yet its customer base includes businesses spread over twelve industries and multiple government agencies (BuffettBeater).

We see the big data industry to be on the early growth portion of the technology S-curve because the world continues to develop and the creation of data coupled with the demand for analyzing it will only grow in the foreseeable future.

Scenario: Real Time Route Planning with Big Data (and eventually Automation)

Smartphones and GPS devices not only offer driving directions, but also give travel time estimates. They base these estimates on information from various sensors and crowd-sourced data. The estimates lag behind real time, and sometimes this is disadvantageous.

Figure 5

Google directions from Caltech to the Staples Center. The estimated travel times were tested and found to be way off.

CA-110 S	**11.7 mi, 17 mins** In current traffic: 20 mins
CA-110 S and I-10 W	15.3 mi, 21 mins In current traffic: 23 mins
N Ave 64 and CA-110 S	12.0 mi, 22 mins In current traffic: 27 mins

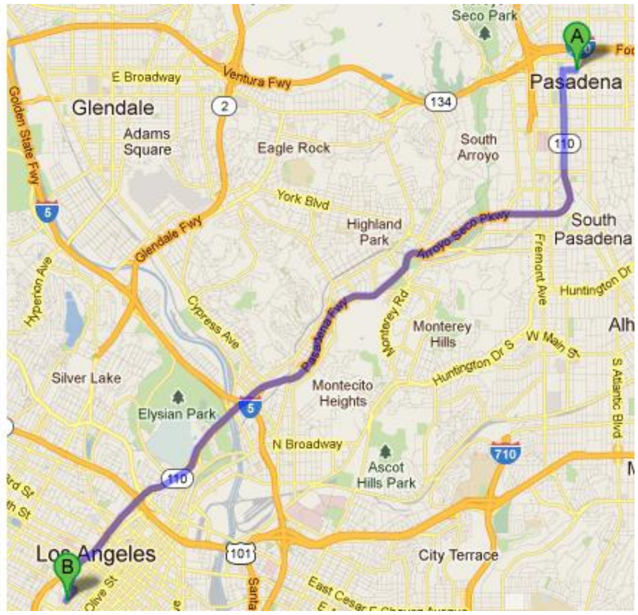

We believe that Big Data can be the missing link that gives route planning and updates in real time. The technology is already in place for individual cars to send and receive data. GPS systems have been built into cars for years, and have become very affordable (the new Camry model even offers a full-fledged mobile OS system). All that's left to do is come up with a way to process data on order of 10^5 variables at a sufficiently fast rate. With the ability to break down large data problems over many data centers and tens of thousands of computing cores, companies such as Amazon Cloud Services and IBM are already equipped with the raw processing power to handle such a task. In the future, Big Data can help link together all the travelers in a city so that traffic is spread out over more roads and GPSs can automatically pick a route based on up-to-the second conditions and even traffic light schedules.

Even if it is not possible to simultaneously handle the traffic calculations for every major city in US at once, we're well on our way to getting there. Processing power is growing exponentially (see Fig. 1) and there is ample research in both faster analysis algorithms and swarm theory as it applies to traffic (Garcia-Nieto). At this rate, the capabilities will become very well developed within the next 5-10 years.

The technology aspects are not the chief barriers inhibiting this goal; rather it is the fact that we currently lack widespread social acceptance for this new technology. Many people would balk at the idea of sharing their location and travel data in a public mainframe. But this is an older way of thinking that is gradually being phased out. The younger generations who are growing up in the age of social media live much more publicly than their parents or grandparents do. Improvements in cyber security can also help mitigate these reservations. In order to implement this system, there would need to be a massive overhaul of what is currently in place and a great deal of public investment. In order to make it worth it, we have to ask if this is all worth the trouble. What does the public gain from this big data-driven system? Hypothetically speaking, if the new system can save, on average, 10-20% of travel times, this equates to roughly 10-20% less time cars spend on the road. This would indirectly decrease both CO_2 emissions and gas consumption. Clearly, there are many benefits to using an improved system.

Taking the concept one step further, using big data to plan driving routes can be integrated into autonomous cars as well. There has already been a lot of interest in automated vehicles. Google's self-driving car has already been licensed in Nevada (PCworld) and Caltech had a team participate in a contest for autonomous highway navigation back in 2003 (Caltech). As mentioned above, scientists are also researching how swarm theory can apply to help regulate traffic. There is active research in replacing every car on the road with auto-piloted vehicles capable of higher safety and efficiency than human drivers. Big data cloud computing can provide the real time processing power necessary to link every individual car on a street to streamline traffic. We believe that developments in this area will begin to show within the next couple decades.

Applications of Cloud Computing

Cloud computing has traditionally been the domain of business applications. The bulk of the money being spent on cloud software development is by large corporations for their own internal systems. Companies such as RightNow Technologies and NetSuite, led by CEOs Greg Gianforte and Zach Nelson respectively, have been in the cloud industry for some time and provide business solutions for companies. RightNow Technologies provides customer relationship management software-as-a-service for well-known companies such as Motorola, Electronic Arts, British Airways, Nikon, and many others (Dr. Gianforte Interview). NetSuite also offers software-as-a-service business management software for accounting, order management, inventory, service automation, and E-commerce (Dr. Nelson Interview).

RightNow Technologies and NetSuite are two model cloud computing companies that were founded before 2000 and have been in the industry since its beginning. They demonstrate the past and current market for cloud computing in business applications. However, as cloud IT solutions are becoming more widespread and accepted the potential markets for cloud are also expanding rapidly. Big data cloud alone is not a business solution but an IT tool. Historically companies have learned how to outsource certain elements to streamline their processes. For instance, companies nowadays typically hire legal consulting firms rather than having their own internal law department because outsourcing lawyers is much more efficient than maintaining their own lawyers. The same can be seen in other departments such as security, transportation, etc. Cloud computing is the next step that allows for outsourcing of IT. Now instead of maintaining their own IT department with physical servers and technical specialists companies can hire a cloud service company to provide all its IT needs.

Also, cloud computing allows computing to be treated as a commodity. In the past if a company needed computing power and storage it would have to purchase its own processors and servers and maintain them. If the company was not using those IT resources at any given time they would be going to waste. However, with cloud computing a company can purchase exactly the computing power it needs. If maximum computing power is only needed occasionally the company does not have to pay for those IT resources to just stay idle. This fundamental change of computing to a usage based model makes cloud computing a much cheaper alternative to the current physical IT structure (Michael Olson Interview).

The outsourcing inherent in cloud computing reduces the cost of maintaining IT and that reduction in cost is the bottom line benefit. That benefit is not limited to companies needing business applications. Cloud computing has a variety of potential applications in other fields that can potentially be explored. One particularly exciting new avenue for cloud applications is science and medicine.

Turbulent air flow is one currently unsolved problem in science that cloud computing could have a major impact on. Turbulence is a phenomenon in fluid flow characterized by chaotic and irregular fluctuations. As a result turbulence is very hard to simulate analytically. However, turbulence is also an essential topic because it is a universal feature found in almost all aspects of life. Air flow over objects ranging from baseballs to airplanes is affected by turbulence on the air boundary. The mixing of air and water at different temperatures that occurs in weather is another application of turbulence. Fluid flow in engines and other machines is also frequently turbulent. Turbulent flow is an integral part of nearly every fluid application. Nobel Laureate Richard Feynman was quoted saying that turbulence is “the most important unsolved problem in classical physics.”

The two computational approaches to handling turbulent flow are modeling and simulation. Turbulence models have existed since the 19^{th} century but so far every model has had significant limitations. The pinnacle of turbulent flow research would be a single closed model that applies universally but no such model has since been created.

Turbulence simulation aims to calculate all the details and variations in flow. Turbulent flow is an amalgamation of constant changing and interacting eddies and vortices. These features cover a wide range of sizes and time scales and evolve, deform, and adapt through time and space. These features are responsible for the chaotic nature and turbulent flow and make simulation an incredibly difficult endeavor. Resolving all the details and scales of the turbulent eddies and vortices often requires 10^8 to 10^{10} computational nodes and up to millions of CPU hours.

Currently only turbulent flow at low Reynolds number, such as in the image below, can be simulated in this way (Imperial). Turbulent flow at higher Reynolds number would be even more complex and convoluted. The limit of feasible Reynolds numbers is significantly less than practically relevant Reynolds numbers. Because of the limitations in modeling and simulating turbulent flow most research relies on experimental studies.

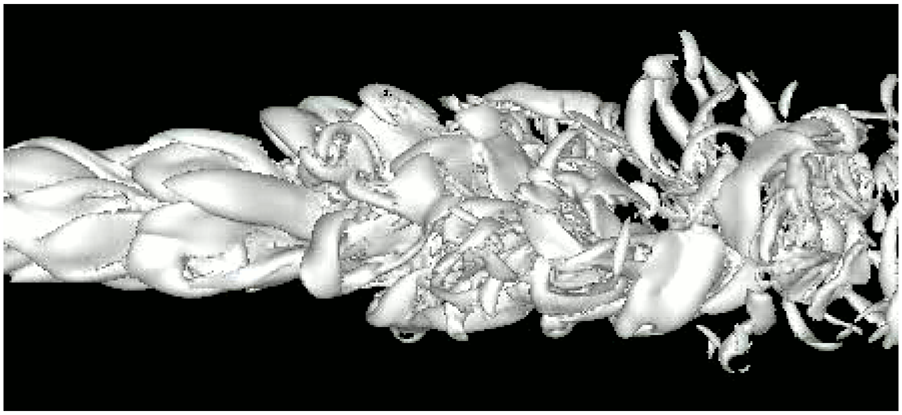

Cloud computing offers a way to harness increased computing power at a low cost. Renting time and processing power through services such as Amazon's cloud server farm is a lot cheaper than purchasing and maintaining super computers and computing clusters, which most research groups currently do. Creating a model for turbulent flow still remains an unsolved scientific problem but big data cloud can help solve the problems posed by turbulent flow in a much less elegant but simpler way. Computational power is the barrier that holds back direct numerical simulation of turbulent flow and cloud processing offers improved computing power at lower prices, expanding the feasible range of simulation. This can push the boundary of what is currently achievable in turbulent flow modeling.

Another potential application for cloud computing is computational biology. One of the problems in simulating experiments is simplification because there are often times too many factors. Introducing every variable at the same time creates an incredibly complex computational problem that is beyond the scope of most computing clusters. For example, in studying cancer treatments researchers simulate millions of artificial compounds to test if they are viable drugs. Simulating these synthetic drugs is the fastest way to identify new potential pharmaceuticals but it goes beyond the limits of the computational resources available at research institutions. Typically researchers introduce simplifications to reduce the complexity of their models in an attempt to make their calculations more feasible. However, this is not ideal because those simplifications can also lead to false results that undermine the entire purpose of the experiment.

Researchers at Schrodinger, a New York firm that creates simulation software for use in pharmaceutical and biotechnology research, have started utilizing the capabilities of cloud computing to improve modeling in computational biology (Arstechnica). Cloud computing clusters make much more processing power available to researchers, who no longer have to cut as many corners as they did in the past to make feasible simulations. Schrodinger was attempting to model 21 million potential cancer drugs and testing their behavior. Running this full test on in-house resources would have taken on the order of years of computing time. The only practical way to run this test on in-house resources was to make simplifying assumptions to reduce the complexity of the calculations.

However, Schrodinger instead took their test to Cycle Computing, which adapts user's software and designs them to utilize the computing power of Amazon's cluster. The Schrodinger test was run in a mere three hours at a cost of $4,828.85 per hour on up to 51,132 cores on Amazon's virtual supercomputer. The processing power was tapped from data centers across the world. Schrodinger found that, compared to the pared down test run on their in-house resources, the full test run on Amazon's computing cluster revealed many false negatives and false positives. Developing an equivalently powerful supercomputer would have cost Schrodinger upwards of $20 million, so reducing the test expense the approximately $15,000 was monumental savings. There was an additional unknown cost required in this process for the subscription service to Cycle Computing, but the savings were still massive. 3D rendering of some of the potential drugs uncovered in Schrodinger's experiment are shown below.

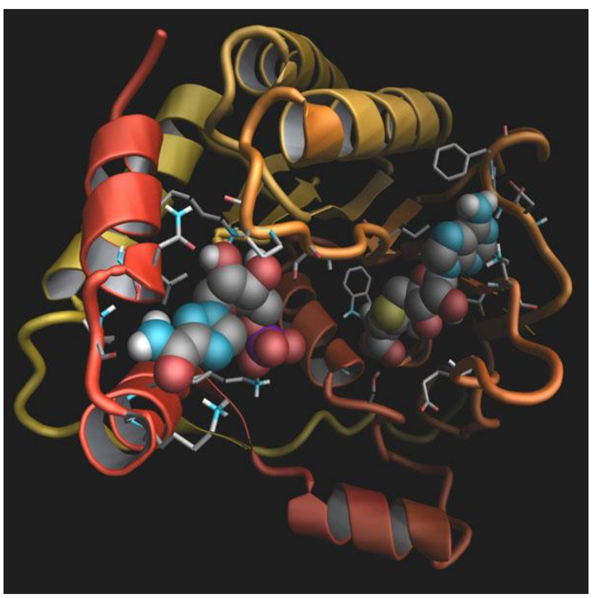
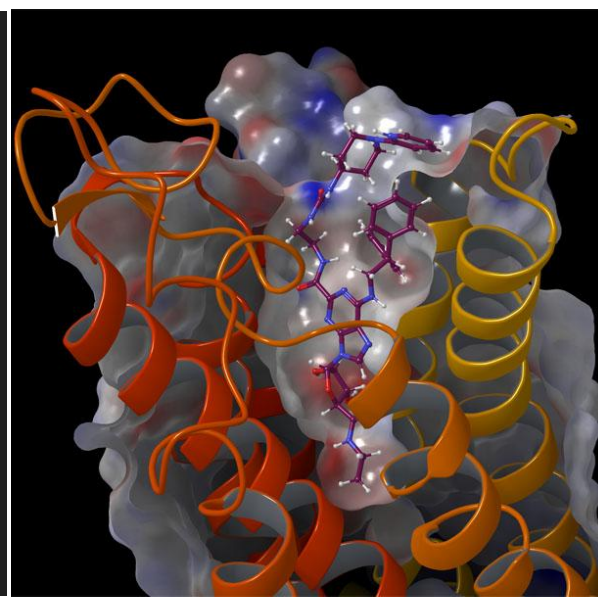

Some supercomputing centers do already exist in universities and research institutions, but this new type of cloud service virtual supercomputer threatens to take over that market. Supercomputing clusters sell time to researchers but buying time on them can be a long and drawn out process. Amazon's cloud and systems like it can be utilized dynamically without any physical interaction and is therefore a much cheaper alternative. Cycle Computing already reported in 2011 that they had developed a similar virtual cluster for another unnamed client. Amazon also markets this service heavily to computational biology and bioinformatics groups across the country. Cloud processing offers a tremendous amount of computing power for much less than it has cost in the past, opening a lot of doors for researchers in computational biology to ask difficult questions and have the computing ability to answer them.

A third potential application of cloud technologies is health care and biomedical informatics. Current clinical records are the domain of hospital systems and most of these systems are independent. Therefore clinical records are typically only shared on a patient by patient basis and not as a whole. This is a severe limitation because clinical data cannot be mined on a large scale. If all clinical data in the United States could be organized and processed as a single system the possibilities would be vast. Clinical practice typically responds slowly to adapt because of the reluctance to accept significant changes in health care. Hospital systems have only switched predominantly from paper based records to electronic records within the last decade and very few currently use cloud based electronic health records (Dr. Abernethy Interview).

Having a centralized clinical database would allow for substantial new clinical research and rapid application to actual health care practice. Mining for trends in clinical data could shed light on patient behavior to specific drugs and how those drugs may act differently on different individuals. For instance, certain drugs may have different effects on specific groups within the population. Without access to a centralized clinical database noticing that aberrant effect would take much longer. Data mining on clinical records could shed light on obscure trends that are very difficult to observe on the small scale of an individual hospital system.

A centralized clinical database would also unify how health care is administered. Right now there is a lot of variation between clinical perspectives and treatment in various hospitals and having a centralized database to give the most complete clinical data available would help improve the delivery of health care. Also, having a centralized system would allow clinical research results to be disseminated more rapidly. Communication between hospital systems would become much easier.

However, there are barriers that prevent cloud technologies from being implemented for health care in this way. Surprisingly, cost is a barrier. The reason is because hospital systems are very independent and each one maintains its own records. As a result the economy of scale that is typically one of cloud's biggest strengths simply does not apply. In addition, as previously mentioned most hospital systems recently upgraded their infrastructure to electronic health records. This upgrade has yet to go through its complete lifespan. Upgrading yet again so soon would be very expensive because the additional upfront cost of upgrading compounds with the cost of the previous system that has not fully be used up. This is a situational problem to cloud adoption in the health care industry. Also, patient care data is extremely sensitive and storing it through the cloud and giving access to hospital systems across the country poses a very difficult security situation. Right now there are no adequate cloud security measures available commercially to make hospital systems willing to adopt cloud databases. There is no denying potential benefits to cloud computing in health care but the current situation coupled with cloud's current limitations prevent cloud solutions from taking over the health care industry at this moment.

A fourth application of cloud technologies is in seismology. Earthquake data is collected across the world and has to be brought together and processed to identify global events. Right now Professor Chandy of Caltech manages a system that uses cloud software to analyze global seismic activity. Local sensors are positioned across the globe to constantly collect seismic data. When an anomaly is detected that data is sent to a central cloud server. That server can then use its own criteria and algorithms to determine, based on anomalous data from many sensors, if a global event is occurring. Cloud services allow this data integration to happen rapidly and in real time (Michael Olson Intervew).

Also, earthquakes are infrequent occurrences and the processing power required for seismic analysis fluctuates significantly. The usage based pricing model of cloud allows this system to be much cheaper that if a private server had to be maintained solely for this purpose. Also, that private server would itself be susceptible to earthquakes. With this system if an earthquake were to hit Caltech a physical server could be disabled, ironically by the exact event it was meant to detect, but with the virtual server offered by a cloud model the processing would not be affected.

Turbulent flow, computational biology, clinical informatics, and seismology are all either current or potential applications for cloud technology. The diversity of these applications speaks to the universal appeal of cloud and the capabilities that it offers. There are still some roadblocks to additional cloud computing adoption that need to be overcome before cloud software becomes truly widespread. The foremost technological concern is data security. So far scientific applications such as research in turbulence, computational biology, and seismology are relatively unconcerned with security because they do not deal with particularly sensitive data. However applications in domains such as health care, banking, of government require much more cyber security than is currently commercially available. The United States government has a vested interest in improving cyber security especially in the current global climate and hopefully military research in this field will drive discovery and lead to improved cyber security in all aspects for cloud computing in the future.

Another major roadblock for cloud computing are geopolitical barriers. Cloud data storage is amorphous and for a cloud company that accesses servers internationally there is no way of knowing where a particular piece of data is stored. Right now storing sensitive data across international borders is a legal problem in many countries. For example, in Europe there are many small countries that cannot warrant having their own data centers based on the principle of economy of scale, but their laws prevent the storage of citizens' data internationally. A cloud company would not be able to operate in a country such as that. Since cloud computing was not on the map until recently many of these legal barriers were not designed with cloud data storage in mind. It is not certain what countries will do to address this problem as cloud computing becomes more universal. However, it is certain that for cloud to operate across international boundaries, which it inherently should do to take advantage of economy of scale, these current legal limitations will have to change or be overcome.

There are other minor concerns with cloud computing that will have to be addressed. One of these is the proliferation of end user interfaces. Current consumers demand that their software not only be conveniently available from a workstation but also from their mobile devices, tablets, automotive computer assistants, etc. A big challenge for cloud software will not just be the software but also the user experience. Also, there are many problems that cloud simply has not yet be applied to solve. One of those applications is credit scoring. As cloud is adapted to other business problems it can create a new market due to its fundamental nature as IT outsourcing and usage based commodity computing.

Six Forces Analysis

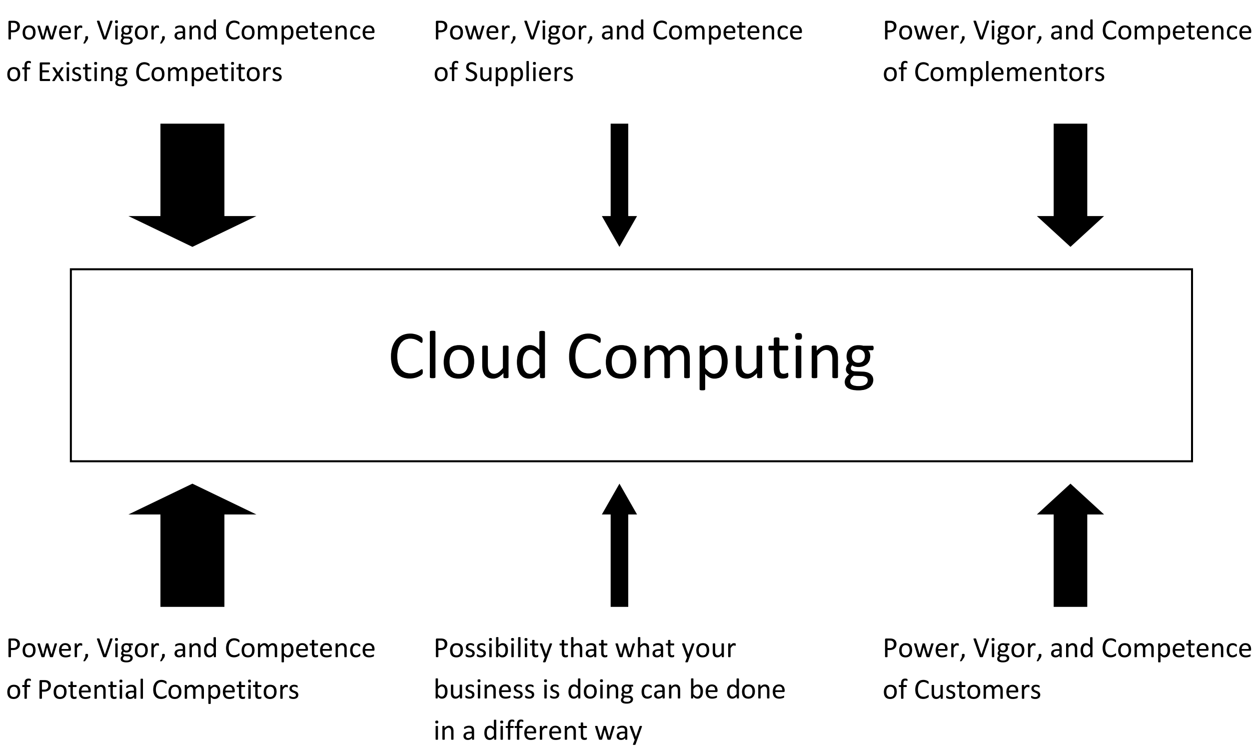

The two biggest forces on the cloud computing market are the existing and potential competitors. There is very little barrier to entry right now in the cloud computing market because physical resources required for a startup are minimal. Existing competitors are some of the biggest names in computing such as Amazon, Oracle, Dell, HP, and others and they are actively driving research, pursuing new industries, and acquiring startups. There is also vigorous investment and the market prospects for cloud computing are high, adding further incentive for potential competitors to join the market.

Two somewhat milder forces on the cloud computing market are the complementors and customers. Complementors are companies that offer software and computing services that may switch to a cloud system in the future. That would open up a new market for cloud to occupy. This is not a strong force because the cloud industry is already healthy as it is and can enter some of those markets independently without the existing market leaders turning to cloud. However, having the existing companies in that market adopt cloud would make the implementation of cloud software as a whole in that market significantly easier. Also, customers are a mild force on the cloud computing industry. Customer support for cloud computing is high because it has a lot of applications and offers many advantages. The customer base is also very large and therefore stable so while some customer demand can drive cloud development that demand is not needed for cloud to succeed.

Finally suppliers and potentially different computing models are not a threat to the cloud industry. Suppliers for cloud computing are simply the companies providing servers and there are a whole host of companies that can do that. Providing a server farm for cloud computing is a necessity but in no way in high demand because there are many roughly equal companies doing so already and providing servers is fairly easy to do. Also right now there is no fear that a different computing model could replace cloud computing. Cloud is that different computing model that is driving out the old style of IT. Zach Nelson, the CEO of NetSuite, called cloud computing the last computing architecture and the final computing paradigm. As of now cloud computing is the pinnacle of computing infrastructures so there is no possibility that another business can be offering computing services in a superior way.

Conclusion

Drawing from our analysis and interviews, we can conclude that the current state of cloud technology still places it in the early stages of the S-curve. The amount of investments into the industry is growing exponentially according to recent market evaluations. Countless start-ups are popping into existence every year, and many recent college graduates are pursuing jobs with these start-ups rather than large and established IT oriented companies such as Google and

Microsoft. The outlook for cloud industry is extremely optimistic, and we expect to see the cloud model adopted by most, if not all companies that currently have self-maintained IT departments.

There are several notable roadblocks that may impede the growth of this technology. The first, which is universally agreed upon by all the professionals we consulted with, is data security. Having a centralized data handling unit means that a great deal of security measure must be put in place before it can be used for any legally or militarily sensitive purposes. A compromise in such a data center can result in information leak that have potentially devastating consequences. A second roadblock that many agreed on is the restriction imposed by geopolitical barriers. Since certain countries currently enforce regulations that require data to be stored locally, much policy reformations needs to occur internationally before cloud can reach its full potential.

As the industry grows, we also expect to see more start-ups in the near future as the hardware and software required for market entry become cheaper over time. In the far future, however, the initial capital investment needed to enter the market will become higher with the existence of dominant cloud service providers that have set the industry standards for data processing and storage rates.

Hypotheses

By 2022, the majority of data storage and maintenance will be conducted off-site using cloud computing technologies.

There is an increasing demand and growing incentive for data storage and maintenance to be moved off-site for many private companies. Currently, the majority of companies in every consumer market have an IT department. By outsourcing their IT management to cloud service providers, they will save a significant amount of money, achieve a greater degree of reliability, and have access to consistent and professional technical support.

By 2020, all television, video games, and other multimedia will be delivered and serviced exclusively through cloud technology.

Currently, a significant portion of PC multimedia is already being distributed through the cloud. For example, Steam, Origins, Hulu, Youtube (acquired by Google), and OnLive are just a few of the most prominent players in this field. Console gaming is also expected to move onto cloud-based distribution. In fact, the Wii console already supports online purchases for minor games. On-demand movie rental services such as Netflix also support online movie streaming using a pay-per-view basis. These companies are all expanding, so there is no reason to refute the hypothesis.

We have rejected the following hypothesis:

By 2025, 90% of day to day economic transactions will be conducted electronically through the cloud.

Economic transactions require highly secure transport channels. Since data security is a primary concern for cloud, this may not be feasible until the general public has grown comfortable with the cloud concept, and has become trusting of handling their personal wealth through electronic means. Given the current concerns with internet security, public perception may not change for at least another two decades.

Teamwork Evaluation

This team worked spectacularly well over the course of this class. Each one of us was capable of doing any part of any assignment, but we each had things we are better at. Hunter, for instance, took care of all the elevator speeches. Tommy was extremely efficient at transcribing the interviews. Tony was an effective motivator. He was also especially apt at providing business perspectives (both Hunter and Tommy are highly technical oriented).

In terms of organization, we made our private Facebook group. Each member took the initiative to suggest meetings as needed, and each responded in a timely fashion. We coordinated as well as we can to conduct interviews whenever all of us can be present, although this is not always possible given the schedule of the interviewee.

There were no significant conflicts. Working together has proven to be an enjoyable experience.

Appendix A: References

Nordhaus, William D. The Progress of Computing. Yale University and the NBER. http://nordhaus.econ.yale.edu/prog_030402_all.pdf

Jcmit: Memory Prices (1957-2012). http://www.jcmit.com/memoryprice.htm

Winchest: Cost of Hard Drive Storage Space. http://ns1758.ca/winch/winchest.html

Server Bandwith: The Need for Server Bandwidth Benchmarking. http://serverbandwidth.blogspot.com/2007/09/need-for-server-bandwidth-benchmarking.html

justmeans: Server Farms: How Biggest Resource Users Drive Sustainable Innovation. http://www.justmeans.com/Server-Farms-How-Biggest-Resource-Users-Drive-Sustainable-Innovation/48302.html

BuffettBeater: Splunk: The Google of Big Data. http://beta.fool.com/buffettbeater/2012/04/23/splunk-google-big-data/3833/

PCworld: Google's Self-Driving Car Licensed to Hit Nevada Streets. http://www.pcworld.com/article/255204/googles_selfdriving_car_licensed_to_hit_nevada_streets.html

Caltech: Caltech Car: No CD Player, No Seats, No Driver. http://media.caltech.edu/press_releases/12416

Imperial College: Research: Turbulent Flow Modeling and Simulation. http://www3.imperial.ac.uk/tfms/turbulentflowmodelling

Arstechnica: $4829-per hour supercomputer built on Amazon cloud to fuel cancer research. http://arstechnica.com/business/2012/04/4829-per-hour-supercomputer-built-on-amazon-cloud-to-fuel-cancer-research/

Garcia-Nieto, J., E. Alba, A. Carolina Olivera. Swarm Intelligence for Traffic Light Scheduling: Applications to Real Urban Areas. Engineering Applications of Artificial Intelligence. http://www.sciencedirect.com/science/article/pii/S0952197611000777

Executive Office of the President. Big Data Across the Federal Government.
http://www.whitehouse.gov/sites/default/files/microsites/ostp/big_data_fact_sheet_final_1.pdf

Appendix B: Interview Summaries

1. Dr. Edouard Servan-Schreiber, Director of Solutions Architecture at 10gen

   Dr. Servan-Schreiber has a PhD in computer science from Berkeley and has been working with databases for his entire career. He explained how VM was one of the major breakthroughs for cloud computing. He also talked about the state of database technology and how NoSQL really changed the way things were done. We asked him about the future of the cloud and he explained his perspective of how cloud will take over IT management for businesses top to bottom in the future. Finally, he addressed some concerns with cyber security and emphasized how it was more of a social acceptance problem.

2. Dr. Greg Gianforte, Former CEO of RightNow that was sold to Oracle

   Dr. Gianforte has been around cloud application companies since its advent in 1997. He explained to us the business models related to cloud computing and stressed how it was a geared to be a business solution. He talked about the advantages of cloud such as improved reliability at reduced costs. He also went into detail about the geopolitical barriers to widespread cloud adoption. Finally, he mentioned that cloud is not currently viable for some businesses because it does not make financial sense for them to switch to cloud. This will change as the number of problems big data can solve increases as it develops.

3. Dr. Zach Nelson, CEO of NetSuite

   Dr. Nelson graduated from Stanford right around the time the first Mac computers came out. He's been around the IT industry for a long time. He called cloud the fusion of telecommunications and computation. He stressed business advantages of using the cloud that were very similar to what Dr. Gianforte said. He sees cloud as the final computing architecture and big data as a solution for the tons of data we create as we integrate the internet and the cloud into businesses. He believes the biggest challenge for cloud computing is simultaneously supporting the various types of user interfaces such as tablets and smart phones beyond traditional computers.

4. Dr. Amy Abernethy, Director of the Duke University School of Medicine

   Dr. Abernethy works in health care policy and improving the delivery of care through IT-enabled systems. She described a possible centralized clinical data center that would facilitate clinical research. It would also improve communication at the intersection of clinical research and clinical practice to more rapidly convert laboratory discoveries to health care benefits. She thinks cloud services will eventually be applied to this purpose but that right now hospital systems are not willing to adopt cloud. This unwillingness is because of high costs due to the private systems of health records in place today and need to replace the entire current infrastructure with a new cloud infrastructure. Security of patient data is also a concern that will have to be addressed before cloud is adopted universally for health care data centers.

5. Dr. Mia Levy, Cancer Clinical Informatics Officer at Vanderbilt

   Dr. Levy works on the continuum of cancer care from computational biology in cancer research to medical informatics for clinical treatment of cancer. She described computational biology and clinical informatics are two potential applications for cloud computing. Computational biology would utilize cloud processing for accessible and affordable usage based computing power. This allows simulations of drug models at unparalleled speed and accuracy compared to traditional computing clusters or in-house supercomputers. However clinical informatics has more difficulty implementing cloud purely because of security concerns. Clinical informatics data is patient data and therefore a lot more sensitive than drug models. As a result clinical informatics will not go to cloud solutions until further advancements have been made in cyber security for cloud computing.

6. Michael Olson, cloud computing expert for Professor Kanianthra Mani Chandy at Caltech

   Olson works on utilizing cloud computing for global earthquake sensor networks. Local sensors transmit anomalous data to a central cloud server that can utilize global information to decide whether an event is an earthquake or not. Cloud computing allows for real time analysis and usage based systems. Commodity computing allows users to only pay for the processing power they need, which helps save a lot of money. Cloud computing in this model has the potential to replace the current market of computing clusters and supercomputers that are run by many universities and research groups.