RJ10480 (A1103-003) March 7, 2011
Other

IBM Research Report

We've Looked at Clouds from Both Sides Now

Tobin J. Lehman
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099
USA

Saurabh Vajpayee
IBM Global Solutions
Charlotte, NC
USA

We've Looked at Clouds from Both Sides Now

Tobin J. Lehman
Almaden Services Research
IBM Almaden Research Center
San Jose, California, U.S.A.
toby@almaden.ibm.com

Saurabh Vajpayee
Global Solutions
ibm.com Sales and Distribution
Charlotte, North Carolina, U.S.A.
saurabhv@us.ibm.com [†]

***Abstract*— Cloud Computing is a versatile technology that can support a broad-spectrum of applications. The low cost of cloud computing and its dynamic scaling renders it an innovation driver for small companies, particularly in the developing world. Cloud deployed enterprise resource planning (ERP), supply chain management applications (SCM), customer relationship management (CRM) applications, medical applications and mobile applications have potential to reach millions of users. Cloud deployed applications that employ mobile devices as end-points are particularly exciting due to the high penetration of mobile devices in countries like China, South Africa and India.**

With the opportunities in cloud computing being greater than at any other time in history, we had to pause and reflect on our own experiences with cloud computing – both as producers and consumers of that technology. Our interests and attitudes toward cloud technology differ considerably for each side of the cloud-computing topic. As producers of cloud-like infrastructure, much of our interest was on the technology itself. We experimented with algorithms for managing remote program invocation, fault tolerance, dynamic load balancing, proactive resource management and meaningful distributed application monitoring. As consumers of cloud computing however, our focus switched from interesting technology to usability, simplicity, reliability and guaranteed rock solid data stability.

With an eye to the many cloud articles in the recent news, we have to ask, is cloud computing ready for prime time? After reviewing stories about current cloud deployments, we conclude that cloud computing is not yet ready for general use; many significant cloud service failures have been reported and several important issues remain unaddressed. Furthermore, besides the failures and gaps in the current cloud offerings, there is an inherent flaw in the model itself. Today, the cloud represents an opportunity for a client to outsource hardware/software function or program computing cycles. The missing piece is responsibility outsourcing – today something found only in IT Outsourcing contracts. This missing piece represents an essential component of a cloud offering. Without it, cloud consumers are left without any real reassurances that their data is safe from failures, catastrophe or court ordered search and seizure.

In this paper, we explore the different viewpoints of cloud computing. Leveraging our experiences on both sides of clouds, we examine clouds from a technology aspect, a service aspect and a responsibility aspect. We highlight some of the opportunities in cloud computing, underlining the importance of clouds and showing why that technology must succeed. Finally, we propose some usability changes for cloud computing that we feel are needed to make clouds ready for prime time.

Keywords-clouds; cloud computing; responsibility outsourcing; mobile computing;

INTRODUCTION

The term "Cloud Computing" has many connotations [1], [2]. For some it connotes grid computing, meaning, a mechanism for a person or business to acquire additional compute, storage or specialized hardware computing resources. For others it connotes software as a service, where it would provide a service that runs their own applications or provides access to third party software. For others it is everything; it is a complete computing infrastructure where the cloud provider manages and monitors all of a customer's computing activity. We believe any reasonable definition of a cloud should include self service, resource pooling, dynamic capacity, heterogeneous network access and measured metrics. [3]. Cloud computing has steadily grown in popularity, with many enthusiastic articles in the current literature [4], [5], [6]. At the same time, some less than satisfactory stories about cloud deployments have surfaced as well [7], [8], [9]. Cloud computing is a concept and technology that is growing and evolving. We believe there is much promise for "The Cloud", but it will have to undergo some changes in order to be the computing foundation for the real world that everyone seems to expect.

We have had experiences on both sides of "The Cloud". As producers of middleware, Service Oriented Architecture infrastructure and compute grid infrastructure, we focused on the science and technology aspect of the project. As potential consumers of a compute cloud, our focus shifted from technology to stability, assurance and ease of use. Interestingly, there's a fairly wide gap between those two viewpoints, even though it is regarding the same technology.

In this paper we review our experiences on both sides of The Cloud. We then review some of the current issues and problems with cloud computing, and then we use those lessons learned to explore what is still needed to make cloud computing truly integral in the computing world of the future. Finally, we paint a picture of the future, where trustworthy, well-managed clouds provide the computing backbone for an up and coming mobile computing deployment, where mobile endpoints become the majority end users for cloud-supported applications.

CLOUD PRODUCERS

Our experience with producing "enabling technology" is with a middleware package, called TSpaces, and a compute grid package called OptimalGrid. In each case, as producers of the technology, we were solely focused on the technical challenges of the technology. In hindsight, we realized that we gave little thought to usability or the actual needs of users. The reason we are emphasizing this point, is because we feel that many other producers of technology have the same mindset, which is technology first and practical usability second.

TSpaces

From 1998 to 2001 the TSpaces project [10] was actively underway at the IBM Almaden Research Center. Advertised as "Intelligent Connection-ware", TSpaces was designed to be network middleware for the new age of ubiquitous computing. It was a relatively simple concept. It was an asynchronous messaging infrastructure, connecting processes via Java objects (called "tuples") that employed a Java-based memory-resident database to manage all of the messages. Using the Tuplespace paradigm of asynchronous communication, first promoted by David Gelernter in the Linda System [11], it was truly a versatile connection service. It enabled communication between applications and devices in a network of heterogeneous computers and operating systems. With the addition of some service registration and location capabilities, it became a universal connection mechanism. With a relatively simple programming interface, users reported creating a simple distributed application in a matter of minutes [12].

TSpaces provides group communication services, database services, URL-based file transfer services, and event notification services. It is implemented in the Java programming language and thus it automatically possesses network ubiquity through platform independence, as well as a standard type representation for all data types.

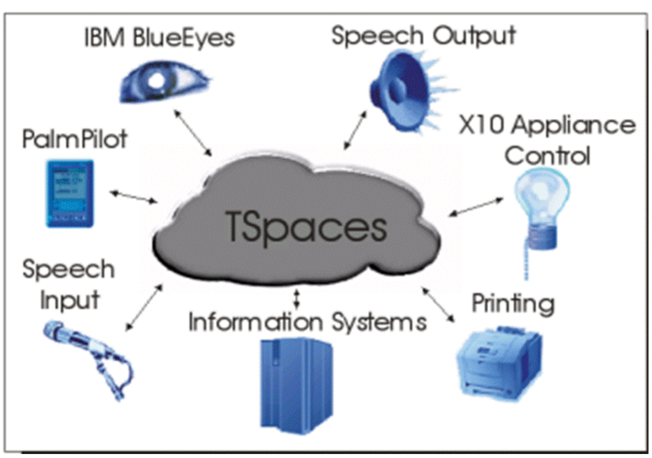

Fig 2.1: The TSpaces Connection-Ware Project

Like many research projects, TSpaces was partly a solution looking for a problem. Although it was created to solve two specific problems (asynchronous data delivery from web crawlers to client devices, and asynchronous communication between web crawlers) [12], it was clear that it could do so much more. As a result, we focused on possible feature extensions and additional applications. It turns out it was good at many things [13], but it did not specifically enable any new application or technology. It just made distributed communication between processes easier and, as a result, could enable rather large systems of services, along the lines of a trellis, as described by Gelernter [11].

Although TSpaces is embedded in several IBM products [14] and has been used in many university classrooms [15], [16] to study distributed computing, the fact that it is middleware practically ensures that it will never be visible to the public. Hindsight suggests that had we spent more time thinking about users and usability, we might have achieved much greater adoption and impact.

OptimalGrid

OptimalGrid [17], [18] was a follow-on project to TSpaces. Leveraging the TSpaces communication infrastructure, it was an autonomic, compute grid system that was primarily built to solve large connected problems that were too large to solve on a single computer. OptimalGrid was autonomic, which means it was designed to handle most types of internet errors, such as communication outage, compute node outage, or compute node slowdown. OptimalGrid's specialty was solving connected parallel problems – as opposed to disconnected parallel problems, such as SETI@Home [19]. Connected parallel problems comprise a special niche in parallel applications that are particularly interesting to computational scientists. From the outset we were focused on that special class of users that would be desperate to use our system, as opposed to users that might judge our system based on other metrics, such as usability.

Figure 2.2 shows a typical deployment of OptimalGrid,. In this simple example application, a connected problem space has been divided into 36 segments. Each segments is allocated to a compute agent (CA), and in this configuration, each group of 9 compute agents is managed by a TSpaces communication node (aka whiteboard). The entire configuration is managed by an Autonomic Program Manager (APM), which uses TSpaces to direct the actions of the compute agents [20].

OptimalGrid offered a new programming model and a new application centric perspective on the use of grids. OptimalGrid provided autonomic functions to hide the complexity of creating and running parallel applications. Different autonomic functionality is required at the two stages of launching a grid application. At Stage 1, program load, a "problem builder" would partition the application and

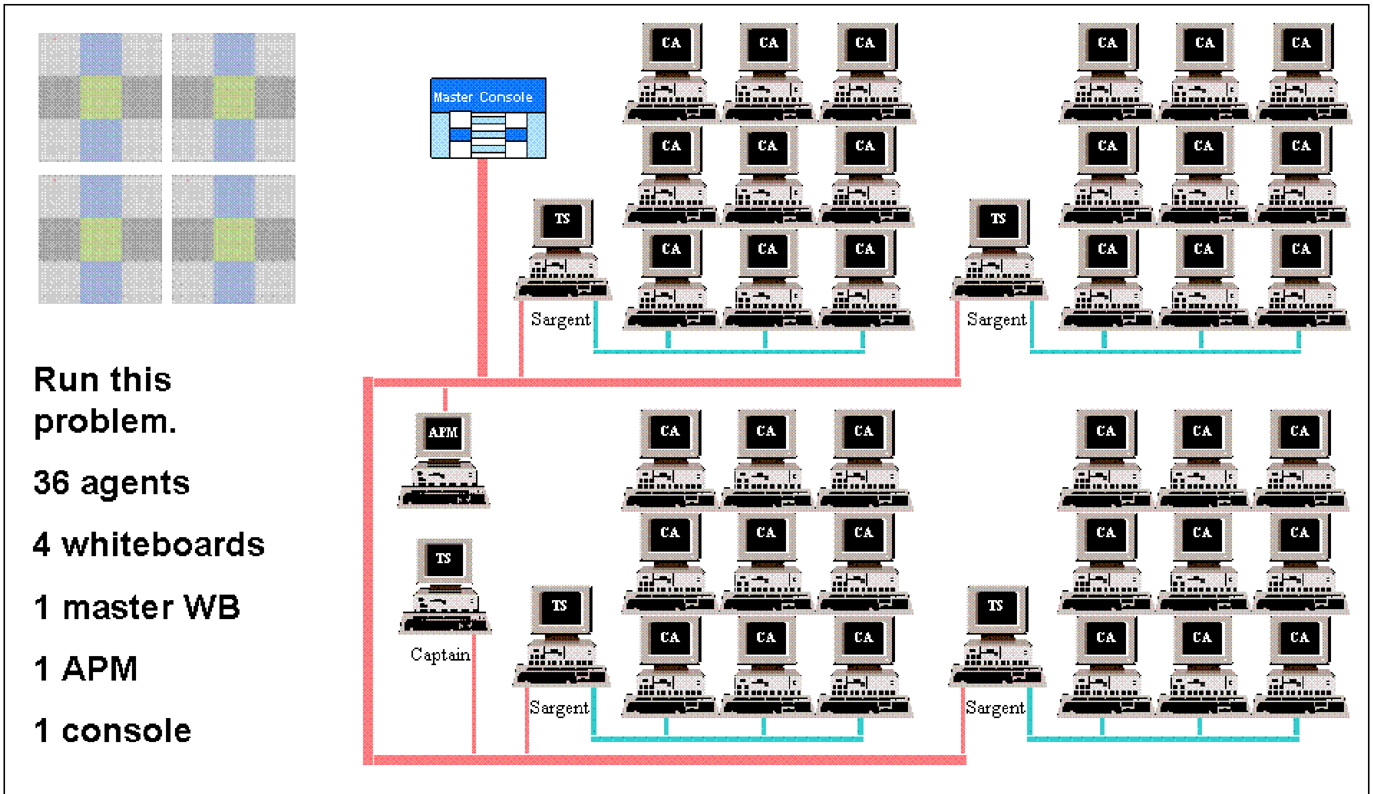

Figure 2.2: A typical OptimalGrid Deployment for solving a small connected problem.

application data. Conceptually, a large-scale connected parallel application is an application that won't "fit on" or run effectively on one machine. It must be divided so as to make use of multiple machines. This is, in many cases, the most intellectually challenging step. If one considers the parallel application as a graph where the nodes on the graph contain data, methods, and pointers to neighbors, the problem is to identify the optimal number and size of problem pieces and to partition the application accordingly [21].

The autonomic program manager (APM) could be launched in "headless mode" and run on a grid compute node, but typically a user would run the APM on their local machine so that they could view the various statistics, graphs and graphical output created by the APM.

Reflections on the Technology

As producers of the TSpaces and OptimalGrid technologies, our focus was on the interesting aspects of the technology itself. With TSpaces, it was the different types of devices and applications we could connect, or it was the added support for additional data types in the tuples, or it was adding replication and fault tolerance. With OptimalGrid, when connected parallel applications were not plentiful, we switched to other applications, such as parallel physics calculations, or neural net applications, or multi-player online game server management. In both cases, though, it was all about the technology and what interesting features, algorithms, or tricks we could add to the project. Ease of use, user friendliness or any overall usability was not a consideration. For both TSpaces and OptimalGrid, some amount of programming was required by the users, but since we were the initial testers and users, that seemed only natural. Then, by interesting coincidence, our initial user community also comprised technologists – who never thought to complain about the lack of user-centric design. It was only when we tried to take the next step to general adoption that we felt the first pushback from users who felt our system lacked usability. As is often the case in research and technology circles where the technical problems outshine the user community pain points, we realized after the fact that we had built a solution in search of a problem, rather than a solution for an existing problem.

CLOUD CONSUMERS

With our current project, we make the transition from technology infrastructure providers to infrastructure users. At the IBM Almaden Research Center, we've recently built a financial management application that models detailed costs for a financial system, such as an organizational budget or a detailed outsourcing contract. Some of the large contracts have on the order of 50,000 cost items in it. Figure 3.1 shows a breakdown of the types of calculations that are frequently done in these contracts. A calculation refresh of a full workbook involves around 60 billion double precision multiplications. Although such a calculation is not prohibitively expensive to do interactively, as one might use a spreadsheet, it is prohibitively expensive to do in batch mode, for operations involving simulations or "what if" analysis. A simulation could easily involve thousands of full workbook updates, which changes the 60 billion double precision multiplies to 60 trillion per simulation.

- 10 year deal, 120 months
- 4 different cost streams
- 50,000 cost items
- Average of 20 cost attributes per cost item
- 10 financial factors
- Average of 12 cost asset types

Figure 3.1: A common calculation breakdown

In order to do simulations or "What if" analysis, it would be necessary to run multiple instances, each instance varying different aspects of the deal: different global resources, different projected inflation rates, different cost items, different acquisition modes, etc. It would be relatively easy to fire off 1,000 different compute nodes, each of which could run 100 different configurations.

Running computations on a compute grid (a simple version of a cloud) is straightforward, but given the nature of some of these financial cases, it is important to guard against public disclosure. Therefore, it would be best to have a private cloud available to us for all of our calculations, but unfortunately, that's not the case. Like many small companies, our department budget is not sufficient for building and owning a private cloud. Our options are limited to a hybrid cloud or a public cloud - which is what started our investigation into the opportunities and problems with using public clouds.

Legal Issues

The larger opportunity lies with the public cloud. That's where the majority of the smaller organizational users will turn for computing resources. But a search of the current literature showed that even though public cloud technology promises to deliver lower cost of ownership, significant legal challenges lie ahead before it can reach its prime. Some key legal issues are:

- Data tracking and auditing in a cloud – Due to the dynamic provisioning nature of cloud, data path and storage is not guaranteed upfront. This is a significant issue for financial industry which is required by law to always keep an audit trail of all financial data. The same can be said about privacy and data security. There are strict privacy guidelines by European Union and other countries. All enterprise level cloud-computing solutions will have to cater to data audit, privacy and security concerns.
- Managing Service Level Contracts – Even though metrics is a key billing concept in cloud-computing, legally enforced service level contracts are needed before a mission critical application can be deployed in a cloud environment. A service level agreement on cloud environment will not only enforce hardware resource like storage, network bandwidth, CPU and memory availability but also support concepts like business continuity and outage. For mission critical application, downtime due to new release deployment, hardware upgrade or a software bug will have to managed and agreed in an efficient way.
- Sub Contracting – A cloud provider will most likely outsource certain cloud operations to a more efficient smaller operator. This has been the business mantra – Keep core and out source rest. A business user accessing cloud-computing for a mission critical application will have to know any such sub contracts between a cloud provider and its sub contractors. This is against the very fundamental concept of cloud-computing where a user does not know the deployment and execution internals. We believe that for business adoption of cloud-computing computing, exposure of sub contracts will be required.
- Data Jurisdiction – This is a significant challenge as highlighted by the given examples. Since a cloud stored data can easily pass legal jurisdiction boundaries, it is a big concern especially of governmental agencies. The need of the hour is to form an international body to arbitrate data ownership and privacy disputes. Since such international cooperation is hard to achieve, we believe the first deployment of enterprise clouds will be geographical and it will be some time before an business acceptable global cloud can be acceptable.

Storm Clouds

Besides the issues with Cloud Technology that make the safety and security of data, well, somewhat cloudy, [Economist 2010] there are many more incidents of disaster and data loss that clients of cloud systems have experienced.

- Gmail fail casts dark cloud on cloud computing
- Microsoft (Danger) and T-Mobile
- iPad Data Loss Represents Increased Need for Cloud Security [7].
- Microsoft's Danger Sidekick data loss casts dark on cloud-computing [8].
- Palm Pre Suffering Cloud Data Loss...Could You Be Next? [22].

Global or local clouds?

Should cloud producers think of global clouds which are hosted in datacenters where there is cheap power and relatively low cost of hardware resources OR should cloud producers think of local clouds? Considering the data jurisdiction, privacy and lack of international cooperation on data ownership disputes, cloud producers are likely to move towards a more local approach to clouds. Following examples highlight the risk of jurisdiction disputes: RIM (providers of Black Berry mobile phones) was asked by India and UAE to allow decryption of emails on servers located in country before forwarding encrypted emails to Black Berry data center in Canada. See the articles in The New York Times [23] and The Wall Street Journal [24]. Google is required by Chinese law to filter out certain search results [25].

What Exactly is a Cloud?

Achieving a clear understanding cloud technology is a challenge. Different providers offering different services with different levels of support all use the term "cloud" to describe their products and services. Hence, the public perception of cloud technology is, well, cloudy. [2] [26]. In simplest terms, cloud computing is an offering of a variety of computing services, including hardware, software applications, computing environments or even complete platforms. [1] [2]

In general, cloud technology is extremely valuable. The ability to pay for ***only*** the needed computing resources is a huge benefit to many organizations. An interesting point along these lines is that IT organizations themselves may benefit greatly from cloud services, as it gives them access to immediate resources, rather than forcing them to suffer delays in procuring machines to satisfy short or long term demand [27].

It is our thesis, though, that there is an important distinction between temporarily offloading some compute cycles or non-essential data files to come remote machine and offloading the responsibility of preserving sensitive corporate data to a third party. Just as we learned our lesson about being focused on interesting technology and not about the real needs of users, we claim that many cloud providers are giving the impression of taking full responsibility for user data when in fact they are only temporarily hosting the data, not taking responsibility for it. Therefore, as the need for increased responsibility by the cloud providers grows, the costs of cloud computing will rise with it. Others have shared this similar view [28].

CLOUDS WITH SILVER LININGS

Hindsight tells us that producers of technology are using an entirely different thought process than the consumers of that same technology. This observation was brought into focus during a recent discussion of cloud technology and the level of responsibility that cloud providers actually accept [29]. That then begs the question, "How can we change the way cloud providers offer their services in order to accept more responsibility for user data?" We feel that the answer lies in the mechanism a user would employ to configure their cloud service.

One of the attractions of cloud technology today is the price. When one pays for only those services it uses, the cost of computing can be very affordable – almost cheap. It is much like renting versus buying a car. It doesn't make sense for someone to buy a car that they drive only once a year. However, it also (usually) doesn't make sense to rent a car for four years. The style of payment should match the use.

Also, the quality or robustness of the service affects the price. The rental price of a new Hyundai would logically be much less than the rental price of a new Ferrari [27]. And, a rental car dealership would likely not do well if its only vehicle type was a Ferrari – it would not be competitive. A rental agency would need a balanced product family – from low line rentals to high profile rentals. Customers would select the product to rent based on need. That is exactly the argument we want to apply to cloud computing.

One of the reasons that the companies offering cloud computing services can't provide iron-clad service contracts is that they would not look competitive with other providers. There must be an obvious distinction in the quality of service. What is needed is a direct correlation between the quality of the cloud service and the price. To go with that correlation, there must be user controls that allow a user to set the cloud service level – from cheap commodity level, to highly specialized and personalized. The more effort a cloud company puts into a cloud service, the more money it can charge for that service.

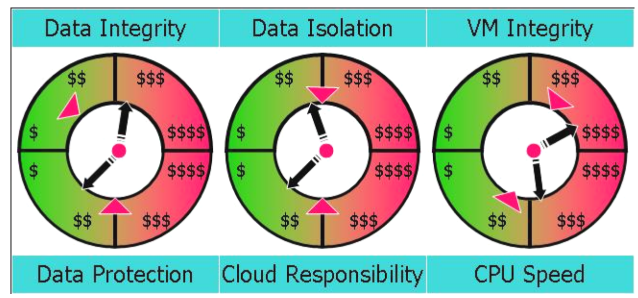

Fig 4.1: Cloud Configuration Settings

Consider the possibility of having user-oriented dials that specify a level of function in the cloud for a specified cost. Figure 4.1 shows an example of user controls for configuring a cloud. The Cloud Configuration Settings (CCS) – would be similar to Service Level Agreements (SLA) – for all of the major system issues. The default settings would be clearly marked (the red triangles in this example), but the user would have the option of selecting stronger CCS values, with the appropriate change in price.

One of the issues often talked about in the news is the problem with data isolation, security and integrity. One of the reasons cloud computing is cheap is the economy of scale, the server consolidation, and the virtualization of the offered cloud machines [27]. For the standard settings, the "regular cloud" configuration would be used. However, for stronger data isolation, which would include physical, logical and legal isolation, specific measures could be taken by the cloud owner to create explicit barriers for certain types of protected data. In fact, a configuration with maximum security, isolation and responsibility settings would be the equivalent of an IT Outsourcing contract.

CLOUDS ON THE HORIZON

Developing economies like India and China are full of opportunities. The number of technology startups in India can be guessed from a simple Google search of "Number of technology startup in India" which yields 1,930,000 results. A similar search for China yields 2,240,000 results. We believe the true opportunity for cloud lie in developing economies of India and China.

The low cost of cloud computing and its dynamic scaling renders it an innovation driver for small companies, particularly in developing world. Cost is a major inhibitor for innovations to move from concept phase to production. The cost of hosting a standard instance of Windows on Amazon's EC2 cloud in APAC region is $0.12 per hour. The cost of 1GB of out bandwidth is as low as $0.19 and storage is $0.15 per GB per month. . A similar instance of dedicated server will cost more and it will be costly for a small organization to provision more servers for scaling up. A bigger cost is associated with under utilization of hardware resources during non peak time. A cloud instance can scale up and down based on the demand, giving a higher utilization of hardware resources. Developing economies are teaming with opportunities. Traditional high IT cost has prevented benefits of applications like enterprise resource planning, supply chain management and customer relationship management to reach small businesses. A cloud deployment of web accessed applications can give small businesses the benefit of information technology.

Mobile phone penetration in emerging economies is huge. The total wireless phone subscribers in India has crossed 670 million [31] and total wireless subscribers in China has crossed 830 million people [32]. These emerging economies have started investment in faster cellular networks like 3G. Smart phones with highly capable mobile handset represent the fastest growing sector in mobile handsets [33]. Mobile phone is fast becoming a ubiquitous device to access the internet. The future of banking, social network, sales and marketing, advertisements, etc. is on a mobile platform. Cloud computing is well suited for middleware deployment of mobile applications. The depth of mobile phone penetration, local language support and known user interface renders cloud deployed mobile applications or what we call mobile clouds, the next big technical revolution. In a way, we can think of Apple's AppStore [34] as a mobile cloud. Mobile cloud deployed applications are well suited for government welfare and crisis management too. Weather prediction and warning, agricultural information, census, disease outbreak detection are good candidates for mobile end point applications deployed on a cloud.

We believe low cost, cloud deployed “pay as you go” billing model applications in the space of supply chain management, customer relationship management, accounting, education and e-commerce will boom in emerging markets. In all likelihood these applications will be developed by local start-up companies to cater local business needs. Emerging markets also place low demand on legal issues, especially on privacy. Opportunities offered by cloud computing are vital for kick starting the innovation wheel in India and China. Cloud computing offers the low cost option to transition from concept to prototype to production needed by inventors.

Conclusion

We’ve looked at clouds from both sides now. From each vantage point, there are entirely different interests and concerns. From a technology point of view, there are interesting technical problems to solve. From a service or consumer point of view, there are essential usability, stability, and reliability problems to solve.

We are at a crossroads with cloud technology. On the one hand, there are many stories of problems with clouds, from data loss, to service interruption, to compromised sensitive data. To stay relevant, to remain meaningful, to grow in the service space, the cloud providers must step up their game and produce robust cloud implementations.

On the other hand, the world is poised to explode with a billion new devices that will be desperate for the very technology that clouds almost offer today. It is possible that the wave of users, applications and demand will just wash over the cloud landscape, regardless of how robust they are.

Regardless of the level of cloud technology adoption, whether clouds are just successful or wildly successful, we feel that cloud offerings must change to offer user-oriented controls, especially those that directly affect sensitive user data. Users need control over a price/performance choice for their data. If the cloud providers are too slow to provide safe, secure, reliable data storage and application services, they may miss one of the greatest opportunities of this century.

Miscellanious

[†] The address and email for Saurabh Vajpayee changed after paper submission. It is now:

G-22 Shanti Nagar, Kanpur-Cantt, Kanpur-208004, India.

References

[1] D. Linthicum, “Selecting the right cloud,” book excerpt, InfoWorld Cloud Computing Deep Dive, InfoWorld, Sept 2009

[2] E. Knorr, G. Gruman, “What cloud computing really is,”, InfoWorld Cloud Computing Deep Dive, InfoWorld, Sept 2009

[3] S. D. Levi, K. C. Riedel, “Cloud computing: understanding the business and legal Issues”, White Paper, http://us.practicallaw.com/8-501-5479

[4] J. Harris, S. Nunn, “Cloud computing’s great promise,” Forbes, June 30, 2010 http://www.forbes.com/2010/06/30/pharmaceuticals-mobile-salesforcecom-technology-cloud-computing.html

[5] E. Knorr, G. Gruman, “What cloud computing really means,” InfoWorld, http://www.infoworld.com/d/cloud-computing/what-cloud-computing-really-means-031

[6] S. M. Fulton, “Here comes the cloud manifesto, to be signed by the usual suspects,” Beta News, March 27, 2009, http://www.betanews.com/article/Here-comes-the-Cloud-Manifesto-to-be-signed-by-the-usual-suspects/1238167857

[7] Kristin, “iPad data loss represents increased need for cloud security”, Backupify, June 11, 2010 “http://blog.backupify.com/2010/06/11/ipad-data-loss-represents-increase-need-for-cloud-security/

[8] D. F. Diglger, “Microsoft's danger sidekick data loss casts dark on cloud computing,” Apple Insider, October 11, 2009 http://www.appleinsider.com/articles/09/10/11/microsofts_danger_sidekick_data_loss_casts_dark_on_cloud_computing.html

[9] “Cloudy with a chance of rain,” Tech.View, The Economist, March 5, 2010, http://www.economist.com/node/15640793

[10] TSpaces Homepage: http://www.almaden.ibm.com/cs/TSpaces/

[11] D. Gelernter, “Mirror worlds, or the day software puts the universe in a shoebox… how it will happen and what it will mean,” Oxford University Press, USA (January 28, 1993)

[12] S.H. Teng, Q. Lu, M. Eichstaedt, D. Ford, T. Lehman, "Collaborative Web Crawling: information gathering/processing over internet," HICSS '99 - Proc of the Thirty-second Annual Hawaii International Conference on System Sciences, Vol 5, 1999.

[13] T. J. Lehman, et al, "Hitting the distributed computing sweet spot with TSpaces," Computer Networks, Vol 35, Issue 4, pp. 457-472,

[14] IBM Products containing TSpaces: Lotus Domino Instant Host, IBM InfoPrint Manager, IBM Update Connector http://www.almaden.ibm.com/cs/TSpaces/licensing.html

[15] E. Reinhard, "Concepts of parallel and distributed processing", Computer Science Dept, Univ Central Florida, Fall 2002 http://www.cs.ucf.edu/~reinhard/classes/cop4520/cop4520.html

[16] R. Van Engelen, "CIS5930: Internet Computing," Computer Science Dept, Florida State University, Fall 2006, http://www.cs.fsu.edu/~engelen/courses/CIS5930/

[17] The OptimalGrid Project page on IBM Alphaworks: http://www.alphaworks.ibm.com/tech/optimalgrid

[18] The OptimalGrid Project Homepage: http://www.almaden.ibm.com/cs/projects/optimalgrid/

[19] SETI@Home, http://setiathome.ssl.berkeley.edu/

[20] T. J. Lehman, R. G. Deen, J. Kaufman, "An autonomic operating environment for large-scale distributed applications," Autonomous Computing, Vol 13, No 1 (January 2006), pp. 81-99

[21] J. Kaufman, T. J. Lehman, J. Thomas, "Grid computing made simple," The Industrial Physicist, http://www.aip.org/tip/INPHFA/vol-9/iss-4/p31.html

[22] R. Evans, "Palm pre suffering cloud data loss...could you be next?" I4U News, Nov 26, 2009 http://www.i4u.com/32111/palm-pre-suffering-cloud-data-losscould-you-be-next

[23] V. Bajaj, "In India, blackberry gets 2-month reprieve on threat of ban," The New York Times, August 30, 2010, http://www.nytimes.com/2010/08/31/technology/31rim.html

[24] M. Coker, T. Falconer, P. Dvorak, "U.A.E. puts the squeeze on BlackBerry," The Wall Street Journal, Technology, August 2, 2010, http://online.wsj.com/article/SB10001424052748704702304575402493300698912.html

[25] Editorial, "Google vs china," The Washington Post, January 13, 2010, http://www.washingtonpost.com/wp-dyn/content/article/2010/01/13/AR2010011302908.html

[26] R. Scheier, "The 9 myths of cloud computing," InfoWorld Cloud Computing Deep Dive, InfoWorld, Sept 2009

[27] M. Beckman, "Cloud options that really help IT," Cloud Computing Deep Dive, InfoWorld, Sept 2009

[28] D. Durkee, "Why cloud computing will never be free," Communications of the ACM Vol. 53 No. 5, Pages 62-69 10.1145/1735223.1735242, http://cacm.acm.org/magazines/2010/5/87259-why-cloud-computing-will-never-be-free/fulltext

[29] T. J. Lehman, "I've looked at clouds from both sides now," presentation at the High Performance Transaction Systems (HPTS) workshop, Asilomar 2009, http://www.hpts.ws/agenda.html

[30] About.com, "Least expensive cars in 2011," http://cars.about.com/od/helpforcarbuyers/tp/Cheapest_09.htm

[31] "Communications in India," Wikipeadia http://en.wikipedia.org/wiki/Communications_in_India

[32] List of countries by number of mobile phones in use, Wikipedia, http://en.wikipedia.org/wiki/List_of_countries_by_number_of_mobile_phones_in_use

[33] "Gartner says worldwide mobile phone sales grew 17 percent in first quarter 2010," http://www.gartner.com/it/page.jsp?id=1372013

[34] Apple iPhone Application Store: http://www.apple.com/iphone/apps-for-iphone/#heroOverview