# The CTO Corner

## July 2010

## Cloud Computing

*Dan Schutzer, Chief Technology Officer, Roundtable Technology Group, Dan@fsround.org*

**What is the Cloud, why has it become important now?**

Cloud Computing applies the utility model to computational power and storage, making these components available on-demand and on a pay-as-you-go basis. Under this model, buyers can forego the capital expenses of computing equipment and just pay for it as an operational expense.

Cloud Computing is such a big and emerging topic that it has caused the National Institute of Standards and Technology (NIST) to provide a definition of Cloud Computing that, among other things, differentiates between the various types of Clouds[1]:

One part of the definition is the SPI (SaaS/PaaS/IaaS) model:

> **Infrastructure as a Service (IaaS).** When a provider allows a user to control the operating system for the hardware (e.g., Amazon's EC2).
>
> **Platform as a Service** (**PaaS).** When a user (application developer) has no control on the operating system, but has control on the application software, and use of application development services (e.g., the Google App Engine).
>
> **Software as a Service (SaaS).** When the end user has no control over the application software except for some configuration controls, like in the case of Salesforce.com.

But the line between these distinctions is gray. For example, Amazon's offerings are called IaaS, but as part of this service Amazon offers database services (SimpleDB & RDS) and a billing service (DevPay), which go beyond IaaS. Also, many refer to data clouds which could be viewed as a specialized SaaS where the service is the ability to store, secure and access data. Other ways of distinguishing between Clouds are to talk about Public Clouds (run and managed by a cloud service provider to any customer that wants to buy cycles); Private Clouds (run and managed by an enterprise for that enterprise); and Hybrids (some combination of Public and Private).

The ideas that cloud computing are based upon have been around for over 20 years. For example, until the 1990s, most software applications were written for mainframes and mini computers that often served as a computing resource to run many different applications in a time-shared mode. In the 90s, personal computing and client server applications began their dominance and many computers were dedicated to a single application. Today, the pendulum is swinging back to server-centric computing, much like the mainframe era, but adding the utility model.

---

[1] http://csrc.nist.gov/groups/SNS/cloud-computing/index.html.

What has changed to make this utility model possible? The main technical advances that allow cloud service providers to offer computational bandwidth on a utility basis is the ability to economically cluster low cost servers through virtualization to achieve the needed high levels of utilization. In addition, the availability of affordable, sufficiently reliable, high speed networks makes this processing and data streaming cost-effectively accessible remotely.[2]

Clouds are characterized by custom low-cost hardware, designed for efficient cooling, cross-time zone load-shifting made possible by the availability of affordable and reliable broadband connections, high CPU utilization, preferential network deals and high power usage efficiency (PUE). Today's public clouds (consisting of 50,000 or more servers) typically cost about .58 picocents[3] per CPU cycle[4], 1 bit storage/year about 5.3 -6 picocent, (which is about the same for small and medium size corporate data centers)[5]; and 1 bit network transfer at about 800 picocents.[6]

We can draw parallels from the power utilities industry and the telecommunications industries. For example, before electric utilities were reliable it was not uncommon for factories and enterprises to own electricity generators. We are now witnessing the cross pollination of ideas between different industry verticals and are in the midst of a convergence of the telecommunications and information technology industries.

**What are the business drivers for an enterprise going to the Cloud?**

The advantage for an enterprise moving to the Cloud is the flexibility it provides: an immediate access to unlimited capacity when you need it with lower operating costs. Clouds give the user the illusion of "unlimited" processing power, available on-demand with very short-term allocation, with just-in-time provisioning and de-provisioning of resources with no up-front commitment[7] and close to 100% transparency.

If you tried to satisfy these demands with an in-house data center, the data center would have to be sized to meet peak real-time processing demands, when most of the time this peak capacity will go unutilized. This is especially compelling for enterprises with short duration peak real-time demands for high performance computing, such as in finance and trading analytics.

Also contributing to the Cloud business case are the maintenance expenses associated with running a data center, which often exceeds the initial capital expenses. The complexity of hardware and software contributes to this trend. Maintaining computing resources by applying patches is cumbersome and time consuming. The cost benefits vary by the nature of the applications (e.g., are there short duration demands for high performance computing, such as experienced in the trading environment), and the size of the enterprise's processing needs. If a financial institution is large

---

[2] This has been the inevitable result of Moore's law – Transistor count doubling every two years; CPU speeds follow Moore ("Gigascale Integration-Challenges and Opportunities," Shekhar Borkar, Director Circuit Research, Intel Corp; and Jakob Nielsen's law that "high end connection speed grows 50% per year")

[3] 1 picocent = $10^{-14}$ USD

[4] This is in contrast with small and medium size corporation data centers which average 4.83 picocent/CPU cycle (5,000 server data center); 14.36 picocent/CPU Cycle (500 server data center) and 26.02 picocent/CPU (50 server data center).

[5] Clouds don't really have a strong benefit in getting cheaper storage; most storage cost is in the energy, for which the Cloud does not have an advantage.

[6] Radu Sion's BITS/FSTC R&D Special Interest Group talk, May 26, 2010.

[7] Often known as "pay as you go."

 

enough to maintain 5,000 or more servers at a high average utilization rate[8], then that enterprise might be able to realize the same economies as the large public clouds, such as Amazon EC2, Google Apps, MS Azure, and IBM's Cloud.

Some of the advantages of Public Clouds with respect to cost are:

- Their large size allows for special demands of their suppliers; e.g., Google mandated Intel hardware to be especially designed for them to run at three degrees higher with the same MTBF (mean time between failures). This greatly reduced the cost of cooling.
- Similar buying power advantages allow them to negotiate lower real estate and preferential networking deals. It also allows them to realize high Power Usage Efficiencies. These costs, plus the personnel to run the data center, can represent as much as 60 to 80% of the cost.
- They offer their services across time zones to effect time zone and load balancing shifting.
- They can locate their data centers in areas where there is a lower cost for cooling their systems; e.g.. next to water dams.
- They can experience efficiencies in staffs in that they can spread the cost of their staff and subject matter experts across a very large number of users.
- The can deliver the services independent of the hardware used, with standard Service Level Agreements (SLAs) and contracts. Once contracts and agreements are customized, the cloud service provider loses a significant part of the cost advantages.

Based upon these benefits and economies, when computation is routine and not a differentiator for an enterprise, the enterprise must consider the options of outsourcing it or purchasing it from a utility provider.

**Concerns, Issues and Barriers**

Some of the main concerns and challenges include:

- Privacy and data confidentiality, including a concern over adequate controls by third-party and side-channel effects[9]
- Shift in liability. Liability issues are extremely important, but it is also important to understand that regardless of cloud deployment[10] and delivery modes in place[11], the risk management responsibility always remains that of the cloud user and business that owns the data.
- Lack of interoperability and standards across public clouds. Concerns about the availability of back-up and vendor lock-in[12] are being addressed by organizations such as the Cloud Computing Interoperability Forum[13], NIST, ITU Cloud Computing Focus group and IEEE Cloud Computing Standards Study Group.

---

[8] Over 60% as contrasted with the typical 15% experienced by many in-house data centers.

[9] http://www.wordiq.com/definition/Side_channel_attack, in cryptography, a **side channel attack** is any attack based on information gained from the physical *implementation* of a cryptosystem, rather than weaknesses in the mathematical algorithms (compare cryptanalysis). For example, timing information, power consumption or even sound provide an extra "channel of information," which can be exploited to break the system.

[10] Public, private, hybrid, or community.

[11] IaaS, PaaS, SaaS.

[12] By vendor lock-in, we mean the inability of the user to change cloud services providers as their systems could not be easily transferred from one cloud provider to another due to lack of standards and interoperability. Once one ports and develops their systems for one cloud, one cannot easily port or transfer them to another cloud provider without significant development.

[13] http://www.cloudforum.org

 

- Regulatory compliance. Regulated entities such as financial institutions still need to ensure they meet compliance and audit expectations. However, this adds to the cost of the service and could represent a customization of the standard SLAs (Service Level Agreements) and contracts to account for much of the cost-savings.
- Transparent infrastructure scalability. A system of records resides on physical premises under someone else's control which could result in trust and transparency issues.
- Portability of legacy applications to the cloud.
- Economic modeling of new markets for a better understanding of the underlying economics and the various pricing models. This also includes issues of chargeback by IT to business users of Cloud services.[14]
- Security. Clouds have become a much larger and more attractive target[15] for hackers and fraudsters, raising concerns about insecure interfaces and API[16], malicious insiders, shared technology vulnerabilities, intentional an unintentional data loss/leakage, account, service and traffic hijacking, supply chain risks, and unknown risks.
- New energy efficient designs.[17]
- Potential for the growth of a Cloud Computing Oligopoly.[18] Because the cost of entry to become a viable Cloud Computing service provider is great, we may see the emergence of a few dominant Cloud service providers and the potential for loss in innovation, restrictive practices and rise in prices.[19]

**What are BITS and FSTC doing with respect to Clouds?[20]**

BITS is exploring Cloud Computing from the aspects of its four program areas: Security, Fraud, Vendor Management and Regulation. Some of the activities being initiated in this regard are:

1. Executive Briefing Paper describing changes to be incorporated in service provider risk assessments. This paper will outline the security risks associated with the three key cloud computing areas of data in the cloud, software in the cloud (i.e., software as a service) and infrastructure in the cloud. Privacy issues may also be a consideration for this paper.

2. Regulatory Framework Briefing Paper will detail current regulations and regulatory guidance related to all aspects of cloud computing. This includes an attempt to identify and document anticipated regulations or regulatory guidance.

---

[14] http://itknowledgeexchange.techtarget.com/total-cio/it-chargeback-a-political-hot-potato-is-tossed-up-by-cloud-computing/?track=NL-964&ad=774399&asrc=EM_NLN_119, IT chargeback: A political hot potato is tossed up by cloud computing, SearchCIO.com blog posted by Laura Smith.

[15] Cloud Security Alliance (CSA) published in February the top ten security concerns for Cloud, "Security Guidance for Critical Areas of Focus in Cloud Computing Version 2.1" at http://cloudsecurityalliance.org/csaguide.pdf). The paper focuses on data and application issues, and is quite comprehensive.

[16] API = Application Programming Interface, http://en.wikipedia.org/wiki/Application_programming_interface

[17] This is of growing importance as energy costs become an ever larger percentage of the data center cost.

[18] http://en.wikipedia.org/wiki/Oligopoly

[19] The growth of standards and interoperability could help prevent these outcomes, lower the barriers to entry, and keep the Cloud marketplace competitive.

[20] Roundtable/BITS/FSTC members interested in participating in these projects should contact Paul Smocer, PaulS@fsround.org.

 

3. Community Cloud Development Considerations Paper. Using the knowledge gained developing the aforementioned papers, participants will continue to research and develop a paper outlining the potential advantages and disadvantages of developing a private or shared community cloud environment for use solely by the financial services industry (in particular Roundtable financial institution members). The initial focus will be on a cloud for data storage

4. Contractual Considerations Best Practices Paper. Develop a set of recommendations for inclusion in service provider contracts specific to cloud services.

**Questions and Opportunities**

As businesses outsource more of their processing to the Cloud, they would expect to have all the associated services also move to the Cloud to maximize the realized savings. It is not a stretch to envision the emergence of specialized platforms targeted for industry specific applications. For instance, Authorize.net provides a set of services for credit card payment accessed over the Internet or "in the Cloud." Some envision a Cloud-based Identity Ecosystem where services related to identifying and authenticating individuals and organizations are handled by identity service providers collaborating and delivering these services to the Cloud-based services[21].

Because of the security and regulatory concerns, we might see the emergence of special "high assurance clouds" designed to provide a higher level of assurance with respect to security, privacy, legal and regulatory compliance, and identity verification. The specialized levels could lead to community-based clouds where communities with special security and regulatory needs, such as the financial services community, establishes its own utility that can economically meet the unique or more demanding needs of its community.

The economics of the Cloud Computing paradigm are transformational because users no longer need to purchase resources for peak demand. As the capital expenses for computation decrease, it creates a level playing field for small businesses and provides them with high performance computing, which was previously out of reach for them. With Cloud Services, an independent chip designer can now potentially rent expensive design tools for short projects and financial service firms can run compute intensive Monte Carlo simulations[22] (to model market movements and credit defaults) and risk analysis workloads without purchasing expensive hardware and software equipment.

Although there are many concerns today regarding the use of Cloud Computing (especially for high risk processing where security, control and resiliency needs are paramount), the underlying economics suggest that these issues will ultimately be satisfactorily addressed and Cloud Computing with likely over time become a dominant computing model. Accordingly, it is important to pilot and keep on top of trends in this area, and understand how the Cloud Computing movement could impact your business and to explore ways to take advantage of these trends.

[21] Cloud Identity Summit, July 20-23, 2010, Keystone, CO.
[22] http://en.wikipedia.org/wiki/Monte_Carlo_method

 