

House of Representatives
Committee on the Judiciary

Cloud Computing: An Overview of the Technology, IP and Market Access Concerns
Facing American Innovators

Wednesday, July 25, 2012

Written Testimony of Justin Freeman, Corporate Counsel, Rackspace US, Inc.

Table of Contents

- I. About Rackspace – Fanatical Support and the Open Cloud..... 3
- II. An Overview of the Cloud..... 3
- III. Major Challenges Facing US Cloud Providers 5
 - A. Market Access & International Privacy Policies..... 5
 - B. Freedom to Innovate & Patent Litigation 6
- Appendix 1: Cloud Services..... 8
- Appendix 2: Cloud Security..... 9

I. About Rackspace – Fanatical Support and the Open Cloud

Founded in 1998 and headquartered in San Antonio Texas, Rackspace is the service leader in cloud computing — a fast-growing industry that helps businesses avoid the expense and hassle of owning and managing their own computer gear by providing computing resources to them over the Internet. Rackspace now serves more than 170,000 customers in 120 countries, including most of the global corporations in the Fortune 100. More than 4,300 engineers, software programmers, customer support representatives, and others provide famed Fanatical Support, the 24/7/365 customer service and support that has defined Rackspace.

One of Rackspace’s top priorities is focusing on the development and deployment of Open Cloud computing infrastructure, based on the OpenStack platform jointly developed with NASA. OpenStack is a set of open-source cloud computing technologies which are platform agnostic – meaning that a company utilizing OpenStack to run its cloud computing services is capable of migrating between a variety of hosting providers and platforms, instead of selecting only one provider and being stuck with that choice. These Open Cloud technologies represent a sea-change in cloud computing – by eliminating proprietary lock-in they help foster critical industry standards for cloud computing and create a robust ecosystem of services which span multiple cloud providers. Much like a cell phone that a user can take from carrier to carrier, applications built on an OpenStack infrastructure can easily be moved between hosting providers.

II. An Overview of the Cloud

At its heart, cloud computing is nothing radically new. “Cloud” essentially describes the use of remote computing resources, whether it be storing information remotely (such as by utilizing a web based email account to store emails in a providers cloud, rather than on a local laptop), or processing information remotely (which occurs when a user leverages the processing power of a remote computer to perform calculations – power which may not be available at a local laptop). These two fundamental computing resources, *storage* and *compute*, are the essence of modern information technology.

What is new is the ubiquitous availability of remote connectivity which drives the cloud revolution. During the first stages of the IT revolution, corporations deployed massive mainframes which handled all the storage and compute needs of users, who accessed these remote resources through terminals. Although few consider this cloud computing, because all the systems were local and required a physical link, the terminal-mainframe model informs modern cloud computing approaches.

As modern workstations increased their storage and processing capabilities, an increasing amount of work was done exclusively on a user’s local computer. Even in the early days of the internet, most storage was local, and local compute power was all that a user had access to. Contrast with today’s cloud, where applications are consumed as remote resources, rather than software running on a local device.

Along with the cloud we now see the commoditization of storage and compute resources, permitting companies to save substantial amounts of capital by paying for modern IT

costs on a utility basis, just like electricity consumption, rather than invest in large capital expense “homegrown” IT infrastructure. This utility model is the blessing of the modern cloud – it permits IT resources to be dynamically allocated as needed, and allows services to be delivered over the internet to almost any user on any device (whether a laptop, cell phone, or tablet). The enhanced user experience and savings drive modern innovation in virtually all sectors of the economy.

The flexibility of IT models has resulted in a lot of confusion regarding what constitutes a cloud. There is no concrete definition – “cloud computing” has become an expansive term encompassing types of infrastructure (dedicating servers to one company’s use, or sharing them to maximize cost savings) and types of services (such as remote email, or remote office applications like Microsoft’s Office 365).

Clouds come in various types and shapes, the configuration of the underlying servers and devices constitutes the infrastructure of the cloud. While the potential recombination is substantial, there are fundamentally three different types of cloud infrastructure:

- **Dedicated Clouds** comprised of physical infrastructure dedicated to one company’s use. That company controls the servers and storage devices exclusively. Also known as private clouds, these are the “single family homes” of the cloud. Dedicated clouds can be located anywhere – at a company’s corporate headquarters or at hosting providers data center.
- **Public Clouds** made up of shared servers whose resources have been virtually partitioned by user. These are the “apartments” of the cloud – all users rely on the same set of underlying devices, and a provider typically manages the segregation of those resources by user. These are the most cost effective types of cloud infrastructure, as the overall capital costs are shared amongst the users, who typically pay only for what they use. Because of their shared nature, public clouds are almost always maintained by a hosting provider at premises that it operates.
- **Hybrid Clouds** come in two flavors, and represent the majority of modern IT usages. A company may split its user of cloud resources between resources dedicated to its use (a dedicated cloud) and resources it shares (a public cloud) in order to balance the need for control provided by dedicated clouds with the cost savings of public clouds. A company may also make use of some computing resources which it runs at its own offices, and some which it outsources to a hosting provider. This balancing act often results as a trade-off between security, control, and cost.

These “different types of clouds” reflect different configurations of computing resources, which are then used to provider different types of services in a ‘pay as you go’ approach. Cloud service models often scale control with cost, and reflect different methods of delivering services through the cloud in a utility pricing model. For a more detailed review of the types of cloud services and their impact on control, please see Appendix 1.

Although the types of resources used by the cloud are not novel, the combination of choice and the ability to hand-off control of IT resources at various levels is. Ultimately,

securing the cloud requires you to know who is in charge of what layer of security, and what they are doing about it (how are they protecting your data?). The fundamentals of IT security are quite similar in the cloud, the focus of a responsible cloud user should be on ensuring that at each layer of cloud security, appropriate controls are in place. Ultimately, the party which controls the data has the most fundamental level of security responsibility – they can encrypt sensitive data and thereby truly protect it from malicious or unauthorized access. For an introduction into the fundamentals of cloud security, a discussion of appropriate cloud security controls, and examples of data types and applicable regulations, please see Appendix 2.

III. Major Challenges Facing U.S. Cloud Providers

The United States is home to the most innovative IT sector in the world, and is especially vibrant when it comes to adopting and innovating in the Cloud. Unfortunately, market barriers resulting from globally inconsistent data protection standards threaten the ability of U.S. companies to compete internationally. Moreover, patent trolls (also known as non-practicing entities or NPEs) are attempting to monetize questionable patents in an all out legal assault directed at the cloud computing industry. It is impossible to overstate how critical market access and an innovative environment are to the ongoing success of the U.S. cloud services industry.

A. Market Access & International Privacy Policies

Many U.S. cloud technology companies are attempting to compete overseas. Much of the time these services are provided out of a U.S. based datacenter to remote users – a position which is increasingly met with opposition from foreign countries concerned about friction between their domestic privacy principles and U.S. law. U.S. cloud providers and technology companies are facing a growing threat to their ability to compete internationally in the form of uncertainty and misrepresentation about their ability to protect and secure data.

EU countries are required to adhere to the principles (implemented differently in each member state) of the EU Data Privacy Directive, a set of requirements intended to protect the rights and privacy of citizens of the EU member countries. EU law currently mandates specific requirements regarding the treatment of data regarding citizens of the member states, including required notifications if the data is shared with third parties. Unease about the U.S. Patriot Act, which requires U.S. companies to comply with U.S. government data requests is driving EU business and regulatory concerns about doing business with U.S. companies.

Cloud sales in Europe trail those in the U.S. by almost 2 years, in part because of these concerns.¹ At Rackspace we routinely face concerns by potential customers based in the EU

¹ Kevin J. O'Brien. "New European Guidelines to Address Cloud Computing." *The New York Times*, July 1, 2012, Technology Section. Available at: https://www.nytimes.com/2012/07/02/technology/new-eu-guidelines-to-address-cloud-computing.html?_r=1&pagewanted=all.

that their mere utilization of cloud services by Rackspace (even in a European data center) would place them in violation of applicable EU regulations. The same uncertainty is appearing in the Indian market, as recent privacy reforms there aligned closely with the EU Data Privacy Directive. The lack of a consistent international privacy regime has resulted in uncertainty that is crippling the ability of U.S. cloud companies to access and compete in international markets.²

EU regulatory authorities are increasingly moving in the direction of denying U.S. cloud providers access to EU markets as a result of this uncertainty. Privacy concerns all too often poison competitiveness as they become the foundation of protectionist measures. U.S. healthcare IT companies have already seen this occur in the form of Canada's FOIPPA healthcare privacy law, which prohibits Canadian healthcare providers from storing patient data on systems located in the U.S. Distrust of U.S. privacy standards by EU regulatory authorities is often general in nature, without a specific legal reasoning or regulatory provision to blame. It is critical that the U.S. government take steps to allay business unease with the unclear regulatory environment and to quash protectionist impulses in the cloud computing market.

It is essential to move towards a consistent international privacy and data transfer framework, while simultaneously providing clear interpretations of U.S. laws which may impact the obligations of U.S. companies serving international customers.

B. Freedom to Innovate & Patent Litigation

Even relatively established cloud computing companies rely on rapid innovation for their success, and the freedom to innovate is critical for nascent cloud technology and service providers. The U.S. patent system is increasingly abused by patent trolls which gather complex software and business practice patents with the sole intent of extracting payments from truly entrepreneurial companies.

The direct costs of this abuse of the patent system are staggering: reaching approximately \$29 billion in in 2011 alone.³ That number is exclusive of the related and often crippling business impact these actions impose on innovative companies such as resource diversion, product delays, and losses of market share.⁴ These costs are a pure social loss – not the result of a repayment to an inventor or a small company whose innovations were unjustly exploited. Instead these lawsuits routinely target large and small operating companies similarly; the only net beneficiaries are the aggressive non-practicing entities which originate these lawsuits. In fact, a substantial number of small and medium businesses are targeted - they comprise about 90% of the companies sued, and their portion of these costs is near 40% of the total.⁵

² Patrick Baillie. "Can European Firms Legally Use U.S. Clouds To Store Data?" *Forbes*, January 2, 2012. Accessed July 23, 2012. Available at: <http://www.forbes.com/sites/ciocentral/2012/01/02/can-european-firms-legally-use-u-s-clouds-to-store-data/>.

³ James E. Bessen & Michael J. Meurer. "The Direct Costs from NPE Disputes." Boston Univ. School of Law, Law and Economics Research Paper No. 12-34. Available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2091210.

⁴ *Id.*

⁵ *Id.*

The cloud industry is under siege. Computing services companies are routinely one of the industries most impacted by patent troll litigation, and the high-tech sector consistently accounts for more than half of all such suits filed.⁶

While recent efforts to reform the patent system have addressed many long-standing problems, the patent trolls have continued their predatory litigation, and further reform is necessary. Focusing on the behavior of the entity, rather than its status as simply a non-practicing entity is a promising way forward.⁷ Fee shifting to favor defendants in cases brought by non-practicing entities, strict limitations on the applicability of notoriously difficult to interpret software and business method patents, and alignment of awards with the value of the underlying patent are potential approaches to this problem.⁸ Absent reform, it is clear that aggressive patent litigation will continue to constrict the resources of well established companies, while exerting a potentially decimating impact on the innovative small and medium businesses the patent system is intended to protect.

⁶ Patent Freedom. "Exposure by Industry." Data captured as of July 13, 2012. Available at: <https://www.patentfreedom.com/about-npes/industry/>

⁷ James E. Bessen & Michael J. Meurer. "The Direct Costs from NPE Disputes." Boston Univ. School of Law, Law and Economics Research Paper No. 12-34.

⁸ *Id.*

Appendix 1: Cloud Services

Types of Services

The different types of clouds (configurations of computing resources) are used to deliver different types of service models. These service models scale control with cost, and are different methods of delivering services in a cost effective utility model. As a user moves from consuming IT resources in the form of dedicated devices (such as servers in a company data center) to consuming IT resources as a service they gradually cede control to providers and third parties.

- **Infrastructure as a Service (IaaS):** In this most fundamental type of IT service, providers control the datacenter, the network, and physical access to servers and storage devices. Users control the rest, and are often responsible for their administration of the IT resources. Most IaaS providers will not permit their customers physical access to devices – all their users share the same physical location, although many of the actual devices are dedicated to particular users rather than shared.
- **Platform as a Service (PaaS):** In the platform model, the provider controls the infrastructure (which of course may be subcontracted) and deliver systems ready to run user's applications. Users bring their applications and data and run them on a ready-to-go platform managed by the provider.
- **Software as a Service (SaaS):** In a SaaS model the underlying IT resources are obfuscated from the user, and the provider delivers a ready-to-use application, maintaining responsibility for the underlying platform and infrastructure. This is the most common type of cloud service for consumers (gmail & Office 365 are great examples – the user consumes and email or office application, without having the software installed locally), and is increasingly relied on by businesses looking for customized off-the-shelf applications, without having to make substantial investments in new computing infrastructure.

Appendix 2: Cloud Security

Fundamentals of Cloud Security

Ultimately, securing the cloud requires you to know who is responsible for each aspect of the cloud resources, and how each layer of security is being addressed. There are three fundamental levels of security in the cloud:

- **Physical Security:** This most fundamental layer relates to having physical access to the IT appliances. If the servers running a cloud are not physically secured from unauthorized access then there is little else that can be done. A malicious party with physical access to a server can readily engage in obvious sabotage such as data theft (even as simple as removing the physical hardware) and physical damage causing data loss, as well as more complicated security risks, such as injecting malicious code or viruses through a thumb drive.
- **Network Security:** It is critical to secure networked systems both from local threats (other users on the same network, including other employees in the same office for example) and remote threats (malicious attacks over the internet). Network security in the cloud is often split amongst multiple parties, so it is especially important for a security conscious user to understand who is responsible for what portion of the network. Insecure networks can permit unauthorized access, the injection of malicious code and viruses, to the more common denial of service attack – where a third party shuts down the ability of servers to function by overwhelming their network capabilities, without necessarily engaging in theft.
- **Logical Security:** The broadest layer of security, logical security relates to controlling user permissions and securing applications from vulnerabilities. Controlling who can get to what based on their access credentials is a fundamental requirement for a secure system. Role based access restrictions are a mechanism of getting users access to the data they need (like quarterly financial statements) while keeping them out of data they don't (like HR records). It also relates to the security of the applications users run – the most common security gap occurs when a user fails to update their operating systems (such as with Microsoft's routine patches) or their anti-virus definitions (without constant updates, anti-virus programs can easily become obsolete).

Selecting the Right Cloud Provider

In order to build a secure cloud, it is essential to select the right cloud infrastructure, select the right provider, review the providers security and operational controls, and to ensure sensitive data is always encrypted.

- **Selecting the right cloud infrastructure:** while an infrastructure dedicated to a single user is typically the most secure, public clouds formed of shared resources can be just as

secure. The key element is identifying how data is secured, regardless of the type of cloud it resides in.

- **Selecting the right provider:** It is critical that cloud users have a clear understanding of the security practices undertaken by their providers, and that the providers are willing to demonstrate compliance with their controls. There is an incredible number of potential combinations of security responsibilities, so users must make sure they choose a combination that meets their needs and capabilities.
- **Reviewing security controls:** It is increasingly common to require a third-party audit of a provider’s security controls, achieving both confidence in the provider and often in order to meet regulatory requirements. Three common audit and control reports are the SSAE16 (a third party review of a company’s ability to meet its stated operational controls), a PCI-DSS audit (commonly utilized in the payment card transaction industry), and a Safe Harbor Self-Certification (especially critical in business ventures between U.S. and EU businesses).
- **Encryption, encryption, encryption:** Regardless of who is responsible for the layers of security, there is only one fundamental method of securing data: encryption. Encryption ensures that even when a system is breached (which increasingly seems like an inevitability even with the best security practices in place) the attacker is unable to utilize any data stolen, mitigating the risks to privacy (in the case of personal information), competitiveness (in the case of proprietary business information), and national security or defense (in the case of military information).

Examples of Regulated Data Types

Cloud users should be especially sensitive to regulatory requirements (whether industry or governmentally based) regarding the types of data they store in the cloud. Below are some examples of the different types of data commonly stored in the cloud and applicable U.S. regulations.

Data Types	Examples	Example Regulations
Personally Identifiable Information (PII)	Credit Card Processing Information	PCI-DSS, Gramm-Leach-Bliley
Protected Health Information (PHI)	Health Records	HIPAA/HITECH
Sensitive Corporate Governance Data	Corporate Audit & Financial Reports	Sarbanes-Oxley
Sensitive Business Information	Forecasts, Development Plans, Strategic Proposals	None – High Economic Value
Generally Public / Non-Sensitive Information	Marketing Collateral, Miscellaneous Documentation	None – Low Economic Value