

**House of Representatives  
Committee on Small Business**

**Protecting Small Businesses Against Emerging and Complex Cyber-Attacks**

**Thursday March 21, 2013**

**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 ..... 4
  - Cloud Infrastructure – Different Shapes of Clouds ..... 4
  - Cloud Service Models – Different Levels of Control ..... 6
- III. Why the Cloud? The Benefits to Businesses & Consumers ..... 8
  - The Cloud and the Coming Mobile Revolution ..... 9
- IV. Security in the Cloud & the Role of the Cloud Provider..... 10
  - A. Fundamentals of Cyber-Security in the Cloud ..... 11
    - The Increasing Complexity of Security and Compliance Requirements ..... 11
    - Compliance Standards and Regulatory Requirements ..... 12
  - B. Security Advantages of Utilizing a Cloud Service Provider ..... 13
    - The Responsible Service Provider as Security Partner..... 14
- V. Principles of Cyber-Security Policy..... 17
  - Responsible Information Sharing..... 17
  - A Light Touch – Flexible Approaches Which Reflect New Realities ..... 18
  - Respect for International Competitiveness & Consistent Regulatory Regimes..... 18


## 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. The broad experienced and developed enterprises which are members of the OpenStack foundation, along with the contributing members of the OpenStack community, help bring a new level of community support and dedication to the project which substantially enhances the security of the OpenStack platform and codebase.


## ABOUT RACKSPACE® HOSTING THE SERVICE LEADER IN CLOUD COMPUTING

**4,300+ RACKERS** 

**180,000+ CUSTOMERS**  
82,000+ SERVERS 

**GLOBAL FOOTPRINT**  
120+ COUNTRIES 

**8 WORLDWIDE DATA CENTERS** 

**PORTFOLIO OF HOSTED SOLUTIONS**  
Dedicated - Cloud - Hybrid 



WE SERVE **60%** OF THE FORTUNE® 100 **OVER \$1 BILLION \$\$\$\$** Annualized Revenue



"We consider our leadership position in Gartner's Magic Quadrant for Managed Hosting providers as further confirmation that Fanatical Support® provides great IT outcomes for our customers."

— Lanham Napier, CEO, Rackspace

**A LEADER**  
Gartner Magic Quadrant  
for Managed Hosting

RACKSPACE HOSTING | WWW.RACKSPACE.COM

Rev. 6.201

## 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 that with today’s cloud, where applications are consumed as remote resources, rather than software running on a local device.

The cloud commoditizes 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 intense “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).

### Cloud Infrastructure – Different Shapes of Clouds

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

**Dedicated Clouds** or **Private Clouds** are 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. Private clouds are a more traditional version of cloud infrastructure, they are not as scalable and involve greater initial costs, but they can be higher performing, are isolated to single users, and can be more appropriate for some types of data processing or applications. For example banks often operate all critical financial resources on a dedicated private cloud – but all their users are still logging in and sharing access to that cloud in a controlled fashion.



- Single-Tenant Infrastructures
- Consumed over private or dedicated networks
- Limited ability to scale
- High initial capital expenditure


**Public Clouds** are made up of shared servers whose resources have been virtually partitioned on a per-user basis. 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. Shared resources (compute and storage) are pooled for efficient pricing, utility billing, and are highly scalable. Public clouds can be provided as services (in the form of applications), or they can be provided as bare resources for businesses to turn on and off as needed.



- Multi-Tenant, shared infrastructure
- Consumed over the web
- “Infinite” Scale
- Utility Billing
- Resources On-Demand
- Zero Capital Expenditure

**Hybrid Clouds** are the latest and most flexible clouds. They allow an “anchor” deployment of a dedicate private cloud to meet higher security or foundation resource needs, and link the parts of the service that need to scale to a public cloud. So while a commerce website around the

holiday season might want a solid, stable, and powerful payment processing system, their needs for serving ads and displaying inventory in webpages are going to be dynamic and high – they might use a private cloud for payments, and a public cloud for web content like their catalog. A company may split its use of cloud systems 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 provide 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.

### Cloud Service Models – Different Levels of Control

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.

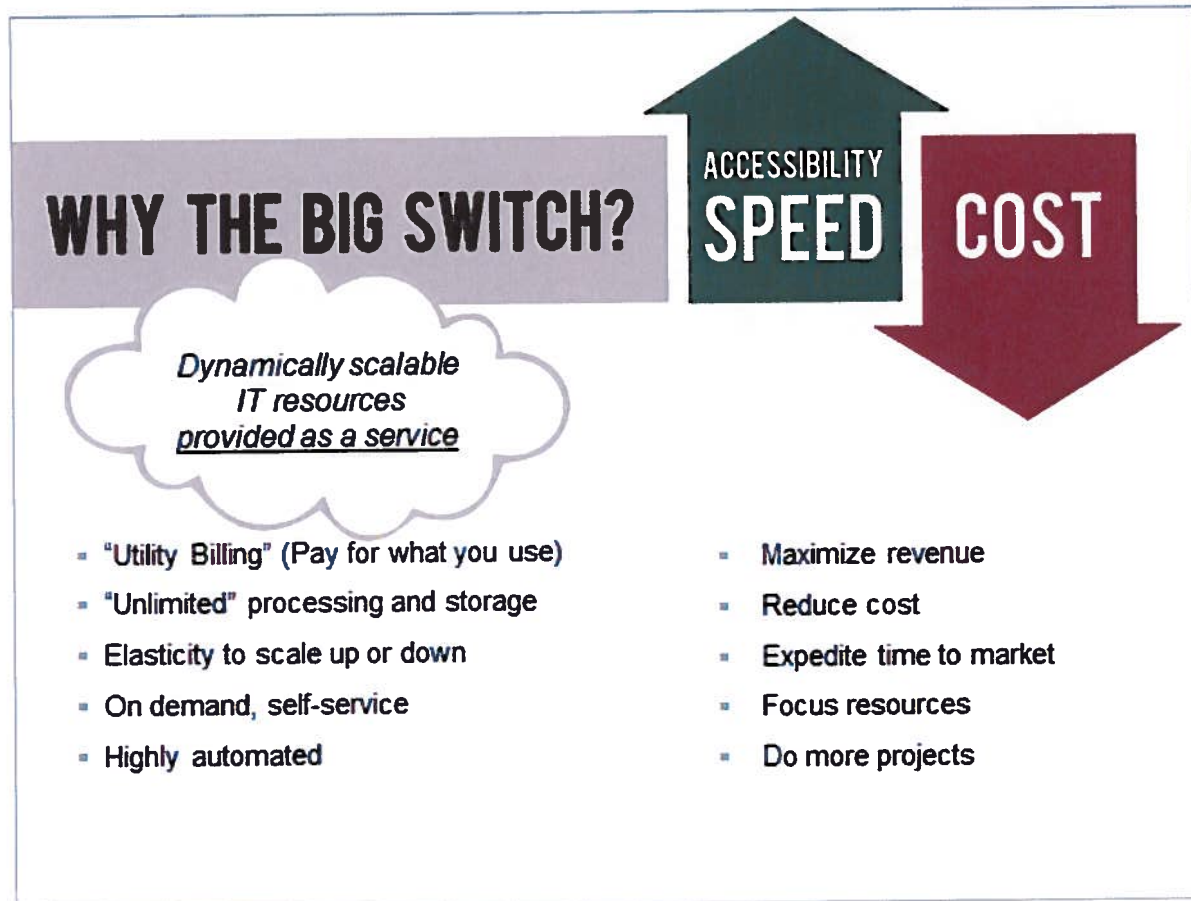
- **Platform as a Service (PaaS):** In the platform model, the provider controls the infrastructure (which of course may be subcontracted) and deliver systems ready to run user's applications. Users bring their applications and data and run them on a ready-to-go platform managed by the provider.
- **Software 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.

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. This paper will focus on cyber-security matters in the cloud, but the fundamental principles remain the same across all IT platforms.



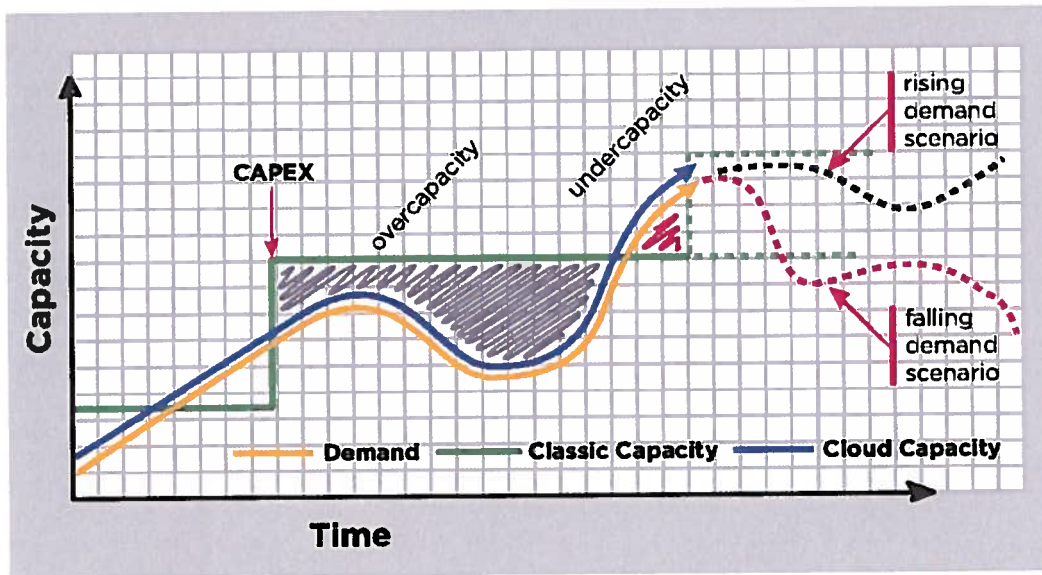
### III. Why the Cloud? The Benefits to Businesses & Consumers

The technological revolution that stems from cloud computing adoption well exceeds the level of change that we have come to expect in information technologies. "The cloud" is not just an incremental improvement in how fast or how heavy our laptops are, it is a massive change. One that represents an entirely new era in IT. Especially for consumers and small businesses who now have access to computing resources in ways never previously possible, breaking down barriers to entry and opening an international market of possibilities.



Cloud technologies shift IT to a public utility model. Centralized IT is like centralized power generation. Instead of generating all your power locally, with the cloud you're connected to a power grid, and can consume what you need from the grid. For businesses this results in a revolutionary approach to consuming IT resources – a push button approach to provisioning additional capacity, and the ability to turn capacity down, or even off. This utility approach is something that has never been possible before in the world of IT, and the savings it offers businesses which cannot afford large capital expenditure management cannot be overstated.





Before the cloud, in order to expand computing power, you had to make single, large investments (illustrated in the stair stepped green CapEx line). This meant that businesses were perpetually under-capacity (which resulting in poor user and customer experiences) or over capacity (which resulted in inefficient use of often limited capital resources). With the cloud in play, businesses can ride the demand curve for computing resources. On-demand, on-time, and on-budget is the promise of cloud for businesses.

Uniquely, businesses can utilize the cloud to export services without previously unapproachable overhead. Delivering services over the internet allows U.S. businesses of all sizes to readily offer their expertise and innovative solutions to consumers globally, opening up the ability to export services without locating a physical presence abroad, and thus supporting domestic jobs and economic growth.

### The Cloud and the Coming Mobile Revolution

The unprecedented access to readily available, highly scalable and utility priced computing resources is heralding a new age in mobile application development and utilization. To date, most mobile applications have focused on 'phone-sizing' a normal web browsing experience, with many applications simply reformatting a traditional web experience for a mobile device.

No longer held back by local resources, and able to access diverse and powerful remote computing resources through the cloud, application developers are shifting their focus to innovative integrated application experiences. These new application approaches will integrate a contextual approach to mobile application utilization – instead of a user navigating a series of menus akin to our traditional application experience, this new generation will rely on integrating information about a user's location, past habits, configured preferences, capable of seamlessly interacting with all of a user's devices (be it a car, fridge, computer, or home monitoring and automation systems) to provide a "contextual experience" – one where achieving a task is intuitive and easy. Leveraging a complex web of cloud-based resources will provide the foundation for a coming revolution in mobile applications and services.

## IV. Security in the Cloud & the Role of the Cloud Provider

Using the cloud safely means using the cloud securely. This consists of four major areas which all cloud users, especially smaller businesses without large compliance and dedicated security teams, need to review. In order to build a secure cloud, it is essential to select the right cloud infrastructure, select the right provider, review the provider's 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. The best cloud configuration is a nimble cloud: a balance of cost and security which scales both based on your use case.
- **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 are 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).

## A. Fundamentals of Cyber-Security in the Cloud

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 layers 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).

## The Increasing Complexity of Security and Compliance Requirements

The most important reason that organizations struggle with security is that the security landscape is complex. There are three dimensions to this complexity:

- First, there are many **attack vectors and security disciplines**. These include identity, access management, encryption, data protection and trust management, denial of service attack mitigation, vulnerability assessment, end-point security and regulatory requirements, just to name a few.

- Second, there are **multiple layers in a given application’s architecture stack**, from the network, systems, operating systems and databases, to the application and its multiple access points. The modern web application is built on multi-tier service-oriented architectures with a myriad of moving parts, and deal with many types of access devices (mobile and desk based). Each one of these layers represents a subset of potential security weaknesses.
- The third is time: **security is an ongoing operation**, not a point-in-time individual project. There is a need for continuous management to maximize the investment one makes in security.

All of this complexity creates the opportunity for specific breakdowns of communication or of operational responsibility. But the most obvious change in the way application services are delivered is that companies and cloud providers have partial ownership over the whole computing stack of a given service. Sometimes, a cloud provider is responsible for the whole compute stack, as occurs with SaaS applications. More commonly, both the customer and the cloud provider take responsibility over a subset of the application stack. For example, it is possible that the customer manages the application-level components while the provider manages the network and the infrastructure (as is the case in many of our public cloud and dedicated offerings). It is this shared scope over the computing stack that complicates the processes and procedures that must be undertaken to deliver a service that reduces security weaknesses.

## Compliance Standards and Regulatory Requirements

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.

<b>Data Types</b>	<b>Examples</b>	<b>Example Regulations</b>
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

Cloud providers offer a host of services which enable customers to meet their regulatory requirements with regard to sensitive data types, although a discussion of these requirements is outside the scope of this comment, a key advantage of cloud services is the ability to implement cost-effective security and regulatory controls with a trusted and audited cloud service provider.

## B. Security Advantages of Utilizing a Cloud Service Provider

Both the cloud provider and the customer have to work together to ensure security best practices are followed and met. Given the role of the customer in the configuration and consumption of their cloud environment, it's imperative for the cloud provider and cloud customer to both put controls in place to manage the risks that multi-tenant environments can present. Both the cloud provider and cloud customer must accept responsibility for different aspects of the system and both must implement a range of controls in order to properly secure the service. Cloud security isn't a sprint – it's a marathon. It can also be thought of as a relay: strong cloud security comprises products and services from various players.

A recent Alert Logic whitepaper<sup>1</sup> (the twice-yearly data-driven analysis called the State of Cloud Security Report) addressed many of the common fears about “ceding control of data to the cloud” – the sense that a business gives up control and capability to deal with security challenges when it chooses a cloud services partner. Contrary to many concerns about security, the study found:

- When compared to traditional in-house managed IT environments, service provider environments show lower occurrence rates for every class of incident examined.
- Service provider customers experienced lower threat diversity (i.e., the number of unique incident classes experienced by a customer) than on-premise customers.
- On-premise environments were twelve times more likely than service provider environments to have common configuration issues, opening the door to compromise.
- While conventional wisdom suggests a higher rate of Web application attacks in the service provider environment, Alert Logic found a higher frequency of these incidents in on-premise environments.

Much of this difference can be explained by evaluating the relevant “IT surface area.” Service providers offer a discrete, centrally managed security environment with tightly controlled network access and enforced partitions between users. A large part of a service providers value stems from their operational expertise and dedication in managing large IT datacenter infrastructure, necessarily architected with security as primary consideration. In contrast, in-house IT deployments often grow as needed in spurts, without the benefit of

---

<sup>1</sup> Removing the Cloud of Insecurity, State of Cloud Security Report, Spring 2012. Alert Logic. Available at: [http://www.rackspace.com/knowledge\\_center/sites/default/files/whitepaper\\_pdf/Alert%20Logic%20Cloud%20Security%20Report%2C%20Spring%202012.pdf](http://www.rackspace.com/knowledge_center/sites/default/files/whitepaper_pdf/Alert%20Logic%20Cloud%20Security%20Report%2C%20Spring%202012.pdf)



advanced central planning which is essential to secure network environments, and are composed of a broad array of operating systems and applications. In-house IT deployments typically also suffer from a greater number of 'entry points' – routes that a malicious attacker can take to access systems and servers – because their networks include mobile devices, desktops, laptops, and other user-facing systems which co-exist in the same network infrastructure. This mixed and broad environment substantially increases the security challenges compared to a dedicated hosting environment. A service provider can also uniquely offer technologies which protect the entire datacenter network, but which are typically prohibitively expensive to single users. The cloud permits economies of scale both in security solutions and compute resources.

### **The Responsible Service Provider as Security Partner**

A well rounded service provider offers a robust approach to a customer's security in the cloud environment; there are four areas to focus on when evaluating a service provider's commitment to providing a secure environment to combat existing and emerging cyber-threats while helping customers meet regulatory requirements:

- **Security Oriented Product Services.** A service provider should offer both fundamental security solutions (such as firewalls, intrusion detection and prevention systems) as well as advanced compliance or security services (such as log management and review, and Denial of Service prevention systems). Often a robust service provider will partner with other trusted enterprises to deliver a broad range of security expertise to meet the myriad of niche compliance requirements facing customers. The provider should act as a trusted guide and advisor, helping a customer navigate the available security services in light of a customer's anticipate risk and compliance needs.
- **Clearly defined roles, access rights, and responsibilities.** A service provider and its customers should work together to identify areas of explicit control (such as the service providers exclusive control over physical security of a datacenter, or a customer's responsibility for securing user-access rights) and establish an on-going understanding regarding areas of shared control (such as management of firewall rules). Security controls documents should identify where a service providers existing controls end, and a customer's responsibility begins.
- **Protect data responsibly.** Service providers should commit to utilizing their customers data in a responsible fashion, including with regards to both civil and criminal data disclosure requests in accordance with applicable law. These commitments can be critical to enabling small businesses without an international presence to compete in an international environment flush with complex and sometimes contradictory privacy and information disclosure requirements.

- Provide audit documentation. Security and compliance is a shared commitment, and no service provider can offer a ready-made cloud compliance solution in a box – but they can lay the foundation for a customer’s use to be secure and meet compliance requirements (whether private or regulatory). A provider should have a commitment to annual third-party audits which it shares with its customers, and its datacenters and services should adhere to recognized certification frameworks. Still, a provider should have a deeper commitment than just having controls in place, it should actively and routinely revisit its security posture and systems controls in light of the evolving nature of cloud systems and threats.

At Rackspace, we believe that security is a partnership, a shared responsibility between our customers and our architects, security experts, and operations personnel. We believe that this is the only way to not only create a technical architecture that reduces the possibility of introducing vulnerabilities into your application, but also to create clarity and understanding about the processes and the proactive and reactive measures that must be put in place. More importantly, realizing that Security is a shared responsibility helps our customers and Rackspace focus on providing transparency on the roles and responsibilities of each party.

The sophistication level of the next generation of cloud-based applications is increasing, and with it, the security landscape is turning more complex. Businesses must pay attention to the details of the relationship as they consider any hosting or cloud provider and dive into the details of the technology, processes and policies to provide an appropriate security level.

A sample of Rackspace’s security controls follows this page, identifying our approach to several security areas.<sup>2</sup>

---

<sup>2</sup> Available at: <https://www.rackspace.com/security/>



## Physical Security

Physical Security includes locking down and logging all physical access to our data centers.

- Data center access is limited to only authorized personnel
- Badges and biometric scanning for controlled data center access
- Security camera monitoring at all data center locations
- Access and video surveillance log retention
- 24x7 onsite staff provides additional protection against unauthorized entry
- Unmarked facilities to help maintain low profile
- Physical security audited by independent firms annually

## Operations Security

Operational Security involves creating business processes and policies that follow security best practices to limit access to confidential information and maintain tight security over time.

- ISO 27001/2 based policies, reviewed at least annually
- Documented infrastructure change management procedures
- Secure document and media destruction
- Incident management function
- Business continuity plan focused on availability of infrastructure
- Independent reviews performed by third parties
- Continuous monitoring and improvement of security program

## Network Infrastructure

Network Infrastructure provides the availability guarantees backed by aggressive SLAs.

- High-performance bandwidth provided by multiple network providers
- Elimination of single points of failure throughout shared network infrastructure
- Cables properly trunked and secured
- Proactive network management methodology monitors network route efficiency
- Real-time topology and configuration improvements to adjust for anomalies
- Network uptime backed by Service Level Agreements
- Network management performed by authorized personnel only

## Environmental Controls

Environmental Controls implemented to help mitigate the risk of service interruption caused by fires, floods, and other forms of natural disasters.

- Dual power paths into facilities
- Uninterruptable power supplies (minimum N+1)
- Diesel generators (minimum N+1)
- Service agreements with fuel suppliers
- HVAC (minimum N+1)
- VESDA / fire suppression
- Flood detection
- Continuous facility monitoring

## Human Resources

Human Resources provides Rackspace employees with an education curriculum to help ensure that they understand their roles and responsibilities as they relate to information security.

- Background screening performed on employees with access to customer accounts
- Employees are required to sign non-disclosure and confidentiality agreements
- Employees undergo mandatory security awareness training upon employment and annually thereafter

## Security Organization

Security Organization includes establishing a Global Security Services team tasked with managing operational risk, by executing an information management framework based on the internationally recognized ISO 27001 Standard.

- Security management responsibilities assigned to Global Security Services
- Chief Security Officer with oversight of Security Operations and Governance, Risk, and Compliance activities
- Direct involvement with Incident Management, Change Management, and Business Continuity

## **V. Principles of Cyber-Security Policy**

Federal regulations have several key roles to play in helping provide for a safe, secure internet which supports innovation, protects consumers, and permits U.S. businesses to export services and remain globally competitive. As in all complex environments, there is no magic bullet, but a coherent and informed set of sector specific regulations, security policies and practices, privacy protections, and collaborative commitments can provide a strong foundation on which U.S. businesses and providers can continue to operate safely and securely.

Sector specific regulations recognize that blanket approaches fail to provide sufficient nuance in implementation and appropriate focus, stifling innovation and fostering inefficiency. Limiting regulations to high-risk areas avoids unnecessary bureaucracy and permits sufficient enforcement resources to be dedicated. High-risk areas and systems should be carefully identified and appropriate standards established to guide, rather than force, the implementation of a stable set of fundamental security practices. No regulatory guidance or set of requirements can themselves establish a network or system which is impervious to attack, but they can provide solid guidance and expertise to inform users and businesses. I urge policy makers to consider the following three principles when considering regulatory or policy action.

### **Responsible Information Sharing**

With appropriate safeguards, increased cyber security collaboration can help prevent denial-of-service attacks, thefts of sensitive information and attacks on the nation's infrastructure. Private industry currently engages in widespread, but often very inefficient, dissemination of security threat information, often limited to a couple of parties or through the security professional community.

As each entity is capable of identifying unique trends and attacks relating to its specific infrastructure, collaboration most often results when one entity's infrastructure is used (after being compromised by a malicious party) to attack another company. Similarly, federal law enforcement agencies track complex international cyber-criminal organizations, but often share that information only with victims, meaning, after the fact. It is critical to establish a clear flow of cyber-threat information from all stakeholders and ensure its widespread accessibility.

But any cyber security bill that is designed to encourage and protect information sharing must recognize the complexities created by shared computing infrastructure environments which are at the heart of cloud computing. Selective information sharing, in the name of stopping cyber threats, can have the unintended consequence of magnifying a security risk. A user of public cloud computing services should not be permitted to disclose information about any vulnerability in that cloud without first disclosing the vulnerability to the cloud provider. This approach would give the provider an opportunity to resolve the vulnerability as quickly and simply as is possible — or to limit or block disclosure in cases where the disclosure would pose a threat to other users served by that provider.

A careful balance must be struck to promote responsible disclosure of vulnerability and threat information without creating additional risk to users of shared computing environments,

and useable, actionable information must be timely disseminated to infrastructure operators and stakeholders.

### **A Light Touch – Flexible Approaches Which Reflect New Realities**

Cyber-security legislation should reflect the new realities of how IT works in the cloud. Any new law should be tailored to fit the way that individuals and companies today use shared computing resources, which are available on demand, across the globe.

Regulations should avoid a reliance on particular or specific technological solutions. These can only stifle innovation by mandating a reliance on a particular technology. Security requirements should be addressable through a reasonable mix of administrative, physical, and technical controls which are evaluated in light of a party's utilization of computing resources.

Likewise while it is critical to allow companies the flexibility to implement to address control requirements, those requirements must avoid excessive vagueness and especially avoid retrospective analysis or results-determined compliance. While regulations should seek to provide controls in appropriate contexts which aim to ensure the security of sensitive data, simply mandating that entities commit to implement controls to 'ensure' security is ineffectual. What controls are sufficient to ensure security? One can only identify gaps after a breach – meaning the ones that were not in place when a compromise happened. The introduction of vagueness or uncertainty into security or privacy regulations can undermine the effectiveness of otherwise well designed approaches, and risk heightening liability for business which are otherwise acting in good faith but find themselves the victim of a malicious third party.

It is also critical to compliment, not contradict, the internet's enabling technologies. Manipulation of fundamental and shared communication systems, such as the Domain Naming System (DNS), are fraught with dangerous consequences and should be avoided without a clear and broad consensus from all stakeholders that such approaches are net-positive for the internet as a whole. Damaging the underlying technologies which enable the internet to be a self-healing, globally effective communications system will only stifle innovation and develop new, unforeseen security risks.

### **Respect for International Competitiveness & Consistent Regulatory Regimes**

Especially when it comes to data privacy and disclosure policies, it is essential at all levels to recognize the impact that data-privacy laws of other countries, particularly in Europe, have on U.S. technology companies — and on U.S. competitiveness. Some European officials, along with European companies that compete against U.S. Internet firms, are actively spreading misinformation about U.S. law.

We continually hear the canard that it is "unsafe" to host data with U.S.-based cloud-computing companies in their European data centers because the U.S. government can access customers' data easily, without due process. While that's simply not true (there is no U.S. law that overrides the data-privacy laws of European nations and The Patriot Act does not allow U.S. law enforcement agencies to access data in overseas data centers) policies which permit information disclosure in the name of law enforcement or security concerns routinely re-ignite

this issue, to the detriment of U.S. companies and the competitiveness of one of America's fastest growing industries.

Congress should provide specific assurances that the data privacy laws of our trading partners will be respected. It is essential that our security and privacy policy regime 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.



THE  
FUTURE OF THE  
CLOUD  
IS  
OPEN

 **rackspace**  
the open cloud company