Technology Review
PUBLISHED BY MIT

January 2009

Security in the Ether

Information technology's next grand challenge will be to secure the cloud--and prove we can trust it.

By David Talbot

In 2006, when Amazon introduced the Elastic Compute Cloud (EC2), it was a watershed event in the quest to transform computing into a ubiquitous utility, like electricity. Suddenly, anyone could scroll through an online menu, whip out a credit card, and hire as much computational horsepower as necessary, paying for it at a fixed rate: initially, 10 cents per hour to use Linux (and, starting in 2008, 12.5 cents per hour to use Windows). Those systems would run on "virtual machines" that could be created and configured in an instant, disappearing just as fast when no longer needed. As their needs grew, clients could simply put more quarters into the meters. Amazon would take care of hassles like maintaining the data center and network. The virtual machines would, of course, run inside real ones: the thousands of humming, blinking servers clustered in Amazon's data centers around the world. The cloud computing service was efficient, cheap, and equally accessible to individuals, companies, research labs, and government agencies.

But it also posed a potential threat. EC2 brought to the masses something once confined mainly to corporate IT systems: engineering in which Oz-like programs called hypervisors create and control virtual processors, networks, and disk drives, many of which may operate on the same physical servers. Computer security researchers had previously shown that when two programs are running simultaneously on the same operating system, an attacker can steal data by using an eavesdropping program to analyze the way those programs share memory space. They posited that the same kinds of attacks might also work in clouds when different virtual machines run on the same server.

In the immensity of a cloud setting, the possibility that a hacker could even find the intended prey on a specific server seemed remote. This year, however, three computer scientists at the University of California, San Diego, and one at MIT went ahead and did it *(see "Snooping Inside Amazon's Cloud" in above image slideshow).* They hired some virtual machines to serve as targets and others to serve as attackers--and tried to get both groups hosted on the same servers at Amazon's data centers. In the end, they succeeded in placing malicious virtual machines on the same servers as targets 40 percent of the time, all for a few dollars. While they didn't actually steal data, the researchers said that such theft was theoretically possible. And they demonstrated how the very advantages of cloud computing--ease of access, affordability, centralization, and flexibility--could give rise to new kinds of insecurity. Amazon stressed that nobody has successfully attacked EC2 in this manner and that the company has now prevented that specific kind of assault (though, understandably, it wouldn't specify how). But what Amazon hasn't solved--what nobody has yet solved--is the security problem inherent in the size and structure of clouds.

Cloud computing--programs and services delivered over theInternet--is rapidly changing the way we use computers (*see Briefing, July/August 2009 (http://www.technologyreview.com/briefings/cloud/) , and "Clouds, Ascending" in above slideshow)*. Gmail, Twitter, and Facebook are all cloud applications, for example. Web-based infrastructure services like Amazon's--as well as versions from vendors such as Rackspace--have attracted legions of corporate and institutional customers drawn by their efficiency and low cost. The clientele for Amazon's cloud services now includes *the New York Times* and Pfizer. And Google's browser and forthcoming operating system (both named Chrome) mean to provide easy access to cloud applications.

Even slow-moving government agencies are getting into the act: the City of Los Angeles uses Google's Apps service for e-mail and other routine applications, and the White House recently launched www.apps.gov to encourage federal agencies to use cloud services. The airline, retail, and financial industries are examples of those that could benefit from cloud computing, says Dale Jorgenson, a Harvard economist and expert on the role of information technology in national productivity. "The focus of IT innovation has shifted from hardware to software applications," he says. "Many of these applications are going on at a blistering pace, and cloud computing is going to be a great facilitative technology for a lot of these people."

Of course, none of this can happen unless cloud services are kept secure. And they are not without risk. When thousands of different clients use the same hardware at large scale, which is the key to the efficiency that cloud computing provides, any breakdowns or hacks could prove devastating to many. "Today you have these huge, mammoth cloud providers with thousands and thousands of companies cohosted in them," says Radu Sion, a computer scientist at the State University of New York at Stony Brook. "If you don't have everybody using the cloud, you can't have a cheap service. But when you have everybody using the clouds, you have all these security issues that you have to solve suddenly."

Cloud Crises

Cloud computing actually poses several separate but related security risks. Not only could stored data be stolen by hackers or lost to breakdowns, but a cloud provider might mishandle data--or be forced to give it up in response to a subpoena. And it's clear enough that such security breaches are not just the stuff of academic experiments. In 2008, a single corrupted bit in messages between servers used by Amazon's Simple Storage Service (S3), which provides online data storage by the gigabyte, forced the system to shut down for several hours. In early 2009, a hacker who correctly guessed the answer to a Twitter employee's personal e-mail security question was able to grab all the documents in the Google Apps account the employee used. (The hacker gleefully sent some to the news media.) Then a bug compromised the sharing restrictions placed on some users' documents in Google Docs. Distinctions were erased; anyone with whom you shared document access could also see documents you shared with anyone else.

And in October, a million T-Mobile Sidekick smart phones lost data after a server failure at Danger, a subsidiary of Microsoft that provided the storage. (Much of the data was later recovered.) Especially with applications delivered through public clouds, "the surface area of attack is very, very high," says Peter Mell, leader of the cloud security team at the National Institute of Standards and Technology (NIST) in Gaithersburg, MD. "Every customer has access to every knob and widget in that application. If they have a single weakness, [an attacker may] have access to all the data."

To all this, the general response of the cloud industry is: clouds are more secure than whatever you're using now. Eran Feigenbaum, director of security for Google Apps, says cloud providers can keep ahead of security threatsmuch more effectively than millions of individuals and thousands of companies running their own computers and server rooms. For all the hype over the Google Docs glitch, he points out, it affected less than .05 percent of documents that Google hosted. "One of the benefits of the cloud was the ability to react in a rapid, uniform manner to these people that were affected," he says. "It was all corrected without users having to install any software, without any server maintenance."

Think about the ways security can be compromised in traditional settings, he adds: two-thirds of respondents to one survey admitted to having mislaid USB keys, many of them holding private company data; at least two million laptops were stolen in the United States in 2008; companies can take three to six months to install urgent security patches, often because of concern that the patches will trigger new glitches. "You can't get 100 percent security and still manage usability," he says. "If you want a perfectly secure system, take a computer, disconnect it from any external sources, don't put it on a network, keep it away from windows. Lock it up in a safe."

But not everyone is so sanguine. At a computer security conference last spring, John Chambers, the chairman of Cisco Systems, called cloud computing a "security nightmare" that "can't be handled in traditional ways." At the same event, Ron Rivest, the MIT computer scientist who coinvented the RSA public-key cryptography algorithm widely used in e-commerce, said that the very term *cloud computing* might better be replaced by *swamp computing*. He later explained that he meant consumers should scrutinize the cloud industry's breezy security claims: "My remark was not intended to say that cloud computing really is 'swamp computing' but, rather, that terminology has a way of affecting our perceptions and expectations. Thus, if we stop using the phrase *cloud computing* and started using *swamp computing* instead, we might find ourselves being much more inquisitive about the services and security guarantees that 'swamp computing providers' give us."

A similar viewpoint, if less colorfully expressed, animates a new effort by NIST to define just what cloud computing is and how its security can be assessed. "Everybody has confusion on this topic," says Peter Mell; NIST is on its 15th version of the document defining the term. "The typical cloud definition is vague enough that it encompasses all of existing modern IT," he says. "And trying to pull out unique security concerns is problematic." NIST hopes that identifying these concerns more clearly will help the industry forge some common standards that will keep data more secure. The agency also wants to make clouds interoperable so that users can more easily move their data from one to another, which could lead to even greater efficiencies.

Given the industry's rapid growth, the murkiness of its current security standards, and the anecdotal accounts of breakdowns, it's not surprising that many companies still look askance at the idea of putting sensitive data in clouds. Though security is currently fairly good, cloud providers will have to prove their reliability over the long term, says Larry Peterson, a computer scientist at Princeton University who directs an Internet test bed called the PlanetLab Consortium. "The cloud provider may have appropriate security mechanisms," Peterson says. "But can I trust not only that he will protect my data from a third party but that he's not going to exploit my data, and that the data will be there five years, or 10 years, from now? Yes, there are security issues that need attention. But technology itself is not enough. The technology here may be out ahead of the comfort and the trust."

In a nondescript data center in Somerville, MA, just outside Boston, lies a tangible

reminder of the distrust that Petersonis talking about. The center is owned by a small company called 2N+1, which offers companies chilled floor space, security, electricity, and connectivity. On the first floor is a collection of a dozen black cabinets full of servers. Vincent Bono, a cofounder of 2N+1, explains these are the property of his first client, a national bank. It chose to keep its own servers rather than hire a cloud. And for security, the bank chose the tangible kind: a steel fence.

Encrypting the Cloud

Cloud providers don't yet have a virtual steel fence to sell you. But at a minimum, they can promise to keep your data on servers in, say, the United States or the European Union, for regulatory compliance or other reasons. And they are working on virtual walls: in August, Amazon announced plans to offer a "private cloud" service that ensures more secure passage of data from a corporate network to Amazon's servers. (The company said this move was not a response to the research by the San Diego and MIT group. According to Adam Selipsky, vice president of Amazon Web Services, the issue was simply that "there is a set of customers and class of applications asking for even more enhanced levels of security than our existing services provided.")

Meanwhile, new security technologies are emerging. A group from Microsoft, for example, has proposed a way to prevent users of one virtual machine on a server from gleaning information by monitoring the use of shared cache memory by another virtual machine on the same server, something that the San Diego and MIT researchers suggested was possible. And researchers at IBM have proposed a new kind of security mechanism that would, in essence, frisk new virtual machines as they entered the cloud. Software would monitor each one to see how it operates and ensure its integrity, in part by exploring its code. Such technologies could be ready for market within two or three years.

But fully ensuring the security of cloud computing will inevitably fall to the field of cryptography. Of course, cloud users can already encrypt data to protect it from being leaked, stolen, or--perhaps above all--released by a cloud provider facing a subpoena. This approach can be problematic, though. Encrypted documents stored in a cloud can't easily be searched or retrieved, and it's hard to perform calculations on encrypted data. Right now, users can get around these problems by leaving their information in the cloud unencrypted ("in the clear") or pulling the encrypted material back out to the safety of their own secure computers and decrypting it when they want to work with it. As a practical matter, this limits the usefulness of clouds. "If you have to actually download everything and move it back to its original place before you can use that data, that is unacceptable at the scale we face today," says Kristin Lauter, who heads the cryptography research group at Microsoft Research.

Emerging encryption technologies, however, could protect data in clouds even as users search it, retrieve it, and perform calculations on it. And this could make cloud computing far more attractive to industries such as banking and health care, which need security for sensitive client and patient data. For starters, several research groups have developed ways of using hierarchical encryption to provide different levels of access to encrypted cloud data.

A patient, for example, could hold a master key to his or her own electronic medical records; physicians, insurers, and others could be granted subkeys providing access to certain parts of that information.

Ideally, we'd make it more practical to work with sensitive data that needs to be encrypted, such as medical records, so that unintended viewers couldn't see it if it were exposed by a hack or a glitch at the cloud provider. "The general theme of cloud computing is that you want to be able to outsource all kinds of functionality but you don't want to give away your privacy--and you need very versatile cryptography to do that," says Craig Gentry, a cryptography researcher at IBM's Watson Research Center in Yorktown, NY. "It will involve cryptography that is more complicated than we use today."

To find and retrieve encrypted documents, groups at Carnegie Mellon University, the University of California, Berkeley, and elsewhere are working on new search strategies that start by tagging encrypted cloud-based files with encrypted metadata. To perform a search, the user encrypts search strings using mathematical functions that enable strings to find matches in the encrypted metadata. No one in the cloud can see the document or even the search term that was used. Microsoft Research recently introduced a theoretical architecture that would stitch together several crytographic technologies to make the encrypted cloud more searchable.

The problem of how to manipulate encrypted data without decrypting it, meanwhile, stumped researchers for decades until Gentry made a breakthrough early in 2009. While the underlying math is a bit thick, Gentry's technique involves performing calculations on the encrypted data with the aid of a mathematical object called an "ideal lattice." In his scheme, any type of calculation can be performed on data that's securely encrypted inside the cloud. The cloud then releases the computed answers--in encrypted form, of course--for users to decode outside the cloud. The downside: the process eats up huge amounts of computational power, making it impractical for clouds right now. "I think one has to recognize it for what it is," says Josyula Rao,

senior manager for security at IBM Research. "It's like the first flight that the Wright Brothers demonstrated." But, Rao says, groups at IBM and elsewhere are working to make Gentry's new algorithms more efficient.

Risks and Benefits

If cloud computing does become secure enough to be used to its full potential, new and troubling issues may arise. For one thing, even clouds that are safe from ordinary hackers could become central points of Internet control, warns Jonathan Zittrain, the cofounder of Harvard's Berkman Center for Internet and Society and the author of *The Future of the Internet--and How to Stop It*. Regulators, courts, or overreaching government officials might see them as convenient places to regulate and censor, he says.

What's more, cloud providers themselves could crack down on clients if, say, copyright holders apply pressure to stop the use of file-sharing software. "For me," Zittrain says, "the biggest issue in cloud security is not the Sidekick situation where Microsoft loses your data." More worrisome to him are "the increased ability for the government to get your stuff, and fewer constitutional protections against it; the increased ability for government to censor; and increased ability for a vendor or government to control innovation and squash truly disruptive things."

Zittrain also fears that if clouds dominate our use of IT, they may turn into the kinds of "walled gardens" that characterized the Internet in the mid-1990s, when companies such as Compuserve, Prodigy, and AOL provided limited menus of online novelties such as news, e-commerce, and e-mail to the hoi polloi. Once people pick a cloud and applications they like, he says--Google Apps, for example--they may find they have limited access to great apps in other clouds, much as Facebook users can't network with people on MySpace.

But such concerns aren't stopping the ascendance of the cloud. And if cloud security is achieved, the benefits could be staggering. "There is a horrendous amount of computing and database management where cloud computing is clearly relevant," says Harvard's Dale Jorgenson. Imagine if today's emerging online repositories for personal health data, such as Google Health and Microsoft HealthVault, could link up with the growing number of electronic records systems at hospitals in a way that keeps private data protected at all times. The resulting medical megacloud could spread existing applications cheaply and efficiently to all corners of the medical profession. Doctors could easily compare patients' MRI scans, for example, with those of other patients around the country, and delve into vast databases to analyze the efficacy of treatments and prevention measures *(see "Prescription: Networking (http://www.technologyreview.com/computing/23545/)," November 2009)*. "The potential there is enormous, because there are a couple of transformations that may occur in medicine in the near future from vast collections of medical records," says Ian Foster, a computer scientist who leads the Computation Institute at Argonne National Laboratory and the University of Chicago. Today, he points out, individuals are demanding access to their own medical information while medical institutions seek new sources of genomic and other data. "The two of those, together, can be powered by large-scalesharing of information," he says. "And maybe you can do it in the cloud. But it has particularly challenging security problems."

This isn't the first time a new information technology has offered profound benefits while raising potentially intolerable security risks. The advent of radio posed similar issues a century ago, says Whitfield Diffie, one of the pioneers of public-key cryptography, who is now a visiting professor at Royal Holloway College at the University of London. Radio was so much more flexible and powerful than what it replaced--the telegraph--that you *had* to adopt it to survive in business or war. The catch was that radio can be picked up by anyone. In radio's case, fast, automated encryption and decryption technologies replaced slow human encoders, making it secure enough to realize its promise. Clouds will experience a similar evolution. "Clouds are systems," says NIST's Peter Mell. "And with systems, you have to think hard and know how to deal with issues in that environment. The scale is so much bigger, and you don't have the physical control. But we think people should be optimistic about what we can do here. If we are clever about deploying cloud computing with a clear-eyed notion of what the risk models are, maybe we *can* actually save the economy through technology."

David Talbot is Technology Review's chief correspondent.