

Testimony of Edward W. Felten
Professor of Computer Science and Public Affairs, Princeton University

United States House of Representatives, Committee on the Judiciary
Subcommittee on the Constitution, Civil Rights, and Civil Liberties
Hearing on
ECPA Reform and the Revolution in Cloud Computing
September 23, 2010

Chairman Nadler, Ranking Member Sensenbrenner, and members of the committee, I thank you for the opportunity to testify about cloud computing and the Electronic Communications Privacy Act (ECPA).

My name is Edward W. Felten. I am a Professor of Computer Science and Public Affairs at Princeton University. I also serve as the founding Director of the Center for Information Technology Policy, an interdisciplinary research and teaching center at Princeton that focuses on public policy issues relating to computing, communications, and the Internet. My primary background is in computer science, and my main subfields in computer science include security and privacy, and Internet technologies. I have served as an advisor or consultant to the U.S. Departments of Defense, Homeland Security, and Justice, and the Federal Trade Commission. I have testified several times previously before House hearings and once before a Senate hearing. I am a Fellow of ACM, the leading professional society for computer scientists, and I serve as Vice-Chair of USACM, which is ACM's U.S. Public Policy Council.

I have been asked to testify about cloud computing technology and its impact on the security and privacy issues implicated by ECPA.

Changes Since 1986

In 1986, when ECPA was passed, the Internet consisted of a few thousand computers. The network was run by the U.S. government for research and education purposes, and commercial activity was forbidden. There were no web pages, because the web had not been invented. Google would not be founded for another decade. Twitter would not be founded for another two decades. Mark Zuckerberg, who would grow up to start Facebook, was two years old.

In talking about advances in computing, people often focus on the equipment. Certainly the advances in computing equipment since 1986 have been spectacular. Compared to the high-end supercomputers of 1986, today's mobile phones have more memory, more computing horsepower, and a better network connection—not to mention a vastly lower price.

But even more important than changes in the equipment have been the changes in how people use the Internet and the role it plays in their everyday lives. From wedding photos to financial records, from love letters to business plans, from grocery lists to boarding passes, the records of our lives are increasing created and archived online. Even those few who don't use the Internet and don't have mobile phones will leave extensive electronic trails online, including their health records and financial transactions.

Cloud Computing for End Users

One of the biggest changes, driven by the vast improvements in storage and networks, has been the shift to cloud computing. At the most basic level, cloud computing means that your data is stored on somebody else's computer. Rather than keeping the only copy of your data on your own computer, you rent computing resources from a service provider, and that provider keeps the primary copy of your data and manages its use. You access the data across the Internet, using your own computer(s).

For example, I use a cloud service to manage my calendar. Keeping my calendar in the cloud has several advantages:

- *professionally managed*: The computers that store the master copy of my calendar are managed by the service provider's employees. They take care of backing up data, maintaining security, and updating software.
- *accessible anywhere*: I can access my calendar from my desktop computer, my laptop, or my mobile phone. If I make a change to my calendar on one of these devices, the change is visible instantly on the others. I don't have to worry about keeping the copies "in sync" because that happens automatically. I can even work while disconnected from the Internet (on an airplane, or in the subway); when I reconnect, any changes made while I was disconnected will be reconciled.
- *easily sharable*: I can give others access to my calendar. My wife can see everything; my students can see when I'm scheduled to be in my office. If I schedule a new appointment, this is visible to all of them right away.
- *supports collaboration*: People can share a calendar, to schedule access to a shared resource such as a conference room or piece of equipment.

Technologically, this kind of cloud service is implemented by cooperation between a server and a set of end-user devices. The server computer sits in the service provider's data center, and acts as the "boss," coordinating the participating devices and storing the archival, master copy of my calendar. People interact with the system on end-user devices such as my desktop computer, my laptop computer, and my mobile phone. The end-user devices get information as necessary from the server, display the information to the user, and tell the server about any modifications that the user makes to the information, such as any new appointments added to the calendar. Together, the server and devices provide the user with the illusion that all of the end-user devices have a view onto a single, shared calendar.

From the standpoint of user experience, a cloud service might look a lot like a traditional local service: information is viewed and manipulated using the local device's display and keyboard, and the information is always available just as if it were stored locally. A less sophisticated user might not even realize that he is using a cloud service, unless he notices that the same data is also available on another device and he stops to think about how this is accomplished.

Many types of services are provided in the cloud. Common examples include email (Yahoo Mail, Microsoft's Hotmail, Google's Gmail), document management (Google Docs, Microsoft Office Live, Zoho Office), investment tracking (Mint, Wesabe), photo-sharing (Flickr, Picasa), project management

(Basecamp, Goplan), hard-drive backup (Mozy, Dropbox), and many more. Any computer-based activity that will benefit from the advantages listed above can be, and probably is, offered via the cloud.

Services provided via the cloud often substitute for traditional packaged software. Rather than buying a software product, installing it on my computer, and using it to manage data locally on my computer, I might subscribe to a cloud service that provides similar functionality via the cloud. This latter approach is sometimes called “Software as a Service” (abbreviated “Saas”). Alternatively, a user might choose to use a cloud service in conjunction with traditional desktop software.

Cloud Computing for Businesses

End users are not the only ones who can benefit from outsourcing their information management. A business can put its back-office (i.e., payroll, sales, inventory, etc.) and customer-facing computing infrastructures “in the cloud” by contracting with a service provider to lease access to resources in the provider’s data center. In the limiting case, the business would not run its own data center at all, but would build all of its computing functionality (other than employees’ desktop computers) to operate in the service provider’s data center. The business would essentially be putting its computing infrastructure into the cloud, in the same way I put my calendar into the cloud. This approach is sometimes called “Infrastructure as a Service” or “IaaS.”

Doing this poses obvious trade-offs for the business. On the one hand, it loses some control over its computing infrastructure and becomes vulnerable to failures by the service provider. On the other hand, it benefits from the economies of scale inherent in the service provider’s larger data center operation, which can lead to lower cost, higher reliability, better energy efficiency, and more professional security management.

Even high-tech start-up companies often turn to the cloud for their computing resources. Although a start-up might have the engineering expertise needed to build and run a data center, it might prefer to outsource that function and have its engineers work on product development instead. An additional benefit is scaling: if the start-up’s business grows rapidly and it needs to expand its computing capacity dramatically to handle a flood of new customers, this is easily done in the cloud, by simply increasing the number of servers the start-up is renting from the provider. Some cloud providers allow their customers to change their resource allocations in nearly real time, something that would not be possible for a customer that operated its own data center.

The point of this discussion is not that everyone should be using the cloud, but that cloud computing offers important advantages which, despite its disadvantages, will make the cloud the right choice for many businesses.

An Example: My Own Use of the Cloud

To illustrate the usefulness of cloud computing, and the variety of cloud services available, let me describe how I use the cloud, as an end user and as the manager of a research group.

As an end user, I keep essentially all of my personal and business data in the cloud. (The only exceptions are a few documents and emails which colleagues or clients ask me to store locally.) Because of the nature of my job I am often on the move, either across the campus or across the country. The ability to access all of my data wherever I am, without having to remember to copy data from here to there or to do an explicit “sync” operation before switching devices, is key to working efficiently while on the move.

By keeping my information in the cloud I can have access to it anywhere, and as a bonus I don’t have to worry about backing up the data because I know there is always a safe copy on the cloud provider’s system. If my laptop computer were destroyed today, I could carry on without loss of data. I would simply get a new laptop, log it in to my accounts on various cloud services, and get back to work, confident that my data would be streamed across the Internet to my new laptop as needed.

I use commercial cloud services to manage my email and calendar. My calendar is shared with my family, colleagues, and students (with appropriate levels of access for each person) as described above. I often use a commercial cloud service to store and edit documents. Any software code and technical reports on which I’m working are stored in the cloud via a version-management tool called “git.” I manage other miscellaneous files using a cloud backup/synchronization service which automatically copies new or modified files to a cloud server and then onto my other computers. My wife and I store our grocery shopping list in the cloud; if my wife adds milk to the shopping list, this will show up on my mobile phone when I get to the store; when I check the box to say I have bought milk, this will be visible on her phone so she won’t accidentally buy milk too.

I wrote this testimony using a cloud editing and document management service. I did some of the writing on my desktop computer and some on my laptop; the editing windows on the two devices were always in sync. When I wanted to ask a few colleagues for feedback on a draft, I shared the cloud document with them, and they wrote comments in the (virtual) margins. I could see their comments in real time, and my colleagues could see me editing to address their comments.

Beyond my personal use of the cloud, I also use the cloud in my research and administration role. My group at Princeton builds various public web sites and tools, which are typically run or hosted using cloud resources that we rent. Although we could run our own servers, this is not cost-effective for us—we can save money and be more agile by renting cloud space for many research purposes.

The Cloud and ECPA

I am not a lawyer, so I will not offer testimony on how ECPA should, or does, apply to data kept in the cloud. However, I understand that the legal treatment of information under ECPA can depend on whether that data is held directly by the user or business, or by a third-party.

In some cases it may be difficult for a user to tell whether or not his data is stored in the cloud, because cloud services can offer nearly the same user experience as local services. This could cause legal uncertainty if the legal status of data depends on whether or not it is stored in the cloud. Even if a user or business knows that data is stored in the cloud, it might not be clear exactly where the data is stored.

From the standpoint of technical and economic efficiency, the cloud offers substantial advantages for at least some users and businesses. Ideally, customers would choose to use the cloud or not by comparing the inherent technical and economic advantages of the cloud approach against its inherent disadvantages, rather than making the decision in order to stay on one side or the other of a legal distinction. To the extent that ECPA considerations dictate decisions to use the cloud or not, this makes computing less efficient and impedes progress toward better technology.