

Giving Web a Memory Cost Its Users Privacy

By John Schwartz

Sept. 4, 2001

See the article in its original context from September 4, 2001, Section A, Page 1 [Buy Reprints](#)

New York Times subscribers* enjoy full access to TimesMachine—view over 150 years of New York Times journalism, as it originally appeared.

[SUBSCRIBE](#)

*Does not include Crossword-only or Cooking-only subscribers.

One day in June 1994, Lou Montulli sat down at his keyboard to fix one of the biggest problems facing the fledgling World Wide Web -- and, as so often happens in the world of technology, he created another one.

At that moment in Web history, every visit to a site was like the first, with no automatic way to record that a visitor had dropped by before. Any commercial transaction would have to be handled from start to finish in one visit, and visitors would have to work their way through the same clicks again and again; it was like visiting a store where the shopkeeper had amnesia.

At 24, Mr. Montulli was the ninth employee hired by what would come to be known as Netscape Communications, and was already known as a programmer of exceptional skill. So he quickly came up with an ingenious idea to address the problem and hammered out a five-page document describing the technology that he and co-workers would design to give the Web a memory.

The solution called for each Web site's computer to place a small file on each visitor's machine that would track what the visitor's computer did at that site. Mr. Montulli called his new technology a "persistent client state object," but he had a catchier name in mind, one from earlier days of computing. When machines passed little bits of code back and forth for such purposes as identification, early programmers called the exchanged data "magic cookies." Mr. Montulli would call his invention, a direct descendant, a "cookie."

It was a turning point in the history of computing: at a stroke, cookies changed the Web from a place of discontinuous visits into a rich environment in which to shop, to play -- even, for some people, to live. Cookies fundamentally altered the nature of surfing the Web from being a relatively anonymous activity, like wandering the streets of a large city, to the kind of environment where records of one's transactions, movements and even desires could be stored, sorted, mined and sold.

Since then, cookies have become nearly ubiquitous -- and that has many people upset. A recent survey by Public Opinion Strategies, a Republican polling organization, found that 67 percent of Americans identify online privacy as a big concern -- far more than those who identify fighting crime (55 percent) or building an antimissile shield (22 percent).

Yet while public anger has grown over invasions of privacy both real and imagined, momentum in Washington to restrict the use of cookies and other high-technology tools for monitoring Internet users' activities has slowed.

In Washington, at least 50 privacy-related bills are awaiting consideration, though the current leadership in the House has focused its attention on privacy invasions by government, not by private business. President Bush's recently appointed chairman of the Federal Trade Commission, Timothy J. Muris, is just preparing his first statement on the commission's direction on privacy, to be delivered next month.

Whether willingly, begrudgingly or unknowingly, however, most Web users have already traded a slice of their privacy for the convenience that cookies bring to the Web. Most people accumulate cookies unknowingly; a search on the average

Internet user's machine will turn up dozens, or even hundreds, of the small files.

Thanks to cookies, a customer shopping at a site who walks away from the shopping cart before buying can come back later to have the site ask if he wants to complete the order. Cookies also allow sites to show advertisements tied directly to the parts of the site a visitor has seen, so that someone visiting a health-oriented site who reads information about diabetes drugs might see an advertisement for a newly approved medication for the condition.

All these functions can be performed without knowing the name of the visitor because the anonymous, unique identifier included in the cookie is enough. But if a Web site owner can combine that identifier with personal information, say from having visitors register with the site, then the cookie becomes a powerful mechanism for personal tracking.

"Before cookies, the Web was essentially private," said Lawrence Lessig, a professor at Stanford Law School who studies the ways that software code and public policy collide. "After cookies, the Web becomes a space capable of extraordinary monitoring."

Most business Web sites now use cookies (including the sites of The New York Times Company) and most use them responsibly, privacy experts say. But many in business fear that privacy concerns could put a further drag on the hobbled high-technology economy. "The danger to the digital economy's longevity is not from the bursting of the dot-com bubble," said Richard H. Brown, chief executive of the technology giant EDS, in a recent speech.

He cited examples like Toysmart, a company that offered to sell its customer records as part of its bankruptcy settlement -- potentially including children's names and addresses. "Those effects are minuscule compared with those inflicted by breaches of trust," Mr. Brown added.

Still, cookies are not going away, said Koen Holtman, a Dutch computer scientist and privacy advocate who has fought to limit the expanding abilities of cookies.

Web users "can't really live with cookies because of user-tracking issues," he said, "but also can't live without them because that would lose them some important functionality or reliability."

Mr. Montulli's first description of cookies can still be found on Netscape's Web site. The document describes how a relatively few bits of text can perform tasks like identifying a visitor, tracking the items he is preparing to buy and setting a date for the cookie to be destroyed. In a whimsical example drawn from Saturday morning cartoons, Mr. Montulli displayed a cookie that might be set on a customer's computer by the fictional Acme Corporation:

```
Cookie: CUSTOMER=WILE_E_COYOTE; PART_NUMBER=ROCKET_LAUNCHER_0001
```

The document was technically thorough. But one word appears nowhere within it: privacy.

Microsoft Takes Notice

The engineers did build in a few privacy precautions, however. Cookies did not identify the user by name. Instead, each site issues a unique ID number to each visitor's computer. Mr. Montulli said that he also considered and rejected an idea for creating a single ID number that a person's browser would use in all Web explorations; while convenient, it would be, he knew, a privacy nightmare. "We didn't want cookies to be used as a general tracking mechanism," he recalled.

But, Mr. Montulli said, he had also planned for cookies to be a flexible tool -- like all Netscape creations. "We were designing the next-generation communications system," he said, and the designers of revolutions don't think small.

"We wanted people to be able to use it for other uses" besides shopping carts, Mr. Montulli said, including "things we hadn't thought about."

By 1995, as Netscape's browser introduced millions of people to the wonders of the Web, another company had taken notice of its success and wanted in on the game. Microsoft aimed at the market for Internet browsers and servers and began a

concerted effort that became the focus of the federal antitrust suit against Microsoft.

But when it came to keeping track of online shopping carts, Microsoft decided not to reinvent the wheel, said Michael Wallent, the head of the company's browser efforts. The company's entry in the browser wars, Internet Explorer, largely incorporated Netscape's cookie system as a "no brainer," Mr. Wallent said.

"I don't think anyone ever thought that cookies were anything that could be excluded in the browser and have that browser become a success in the marketplace," he said.

Like Netscape, Microsoft kept its cookies under the table: cookies were designed to be exchanged silently, without alerting the user. With other Web browser functions, like encrypted communication, an icon appears on the computer screen when the technology is in use. Mr. Wallent explained that privacy was not, at the time, a central consideration because the Web "was a very different place."

"While privacy was an issue, it was much less of an issue than you see today," he said.

Although they were not obvious to the average computer user, cookies were quickly noticed within the technology community. Members of the Internet Engineering Task Force, a group that evolved from the time of the Internet's predecessor, the Arpanet, to become the standards-setting body for the ever-evolving worldwide computer network, started in April 1995 to discuss cookies.

Despite Mr. Montulli's prowess, the technology was less than robust. Simon St. Laurent, the author of "Cookies," a technical work, said of Mr. Montulli's original version: "It kind of works, but it's definitely concocted overnight." Discussions began among Internet experts about the kinds of things that Internet engineers fret over, like ways to make the system more secure and reliable. Within the discussion, some were pressing for consideration of privacy issues.

And so, in 1995, a group was formed to come up with proposed standards for cookies and their uses; it was led by David M. Kristol, a scientist at Bell Laboratories whose outside interests included the intricate interplay of chamber music. He estimated that the job would take a few months.

He worked on it for nearly six years.

Like all such groups, the work was public and carried out largely through online postings and e-mail. Mr. Montulli was an active participant -- at least at the beginning. "I remember saying that it was very important that if we made any changes at all to the way things work, that it needed to be a more forward-compatible kind of thing: the old stuff should still work, and people's general idea of cookies will stay the same."

The members of the working group agreed: although they wanted to improve on cookies technology, they realized that whatever recommendations they came up with should work a lot like the current cookies, or the effort would be wasted.

Increasingly, the group became concerned about the ways that cookies might be used to violate consumer privacy. Mr. Holtman, the Dutch computer scientist, issued a warning to the group in December 1995 that would turn out to be prophetic.

Although cookies can only be read by the site that created them or a related site -- another of Mr. Montulli's early privacy measures -- Mr. Holtman realized that companies could, by agreement, place cookies across a network of related sites, and that those cookies could be used to track users.

"Someone is bound to try this trick," he wrote, "and it will, when discovered, generate a lot of bad publicity for the whole Web."

What Mr. Holtman did not know was that companies were already planning to exploit this wrinkle of the Web. Before long, large Internet advertising companies like DoubleClick and Engage were displaying ads across thousands of sites, using a

common cookie across the network that allowed the company to recognize a visitor wherever he wandered on the Web. The innovation allowed these companies to rotate the ads the user sees from site to site.

DoubleClick's Web site says that it "allows marketers to deliver the right message, to the right person, at the right time." The concern of privacy advocates, however, was that these "third-party cookies" could also be used to build a detailed profile of a Web user's habits.

If a Web surfer visited a large number of sites about AIDS treatment, for example, and if that data were tied to information that identified him -- say, registration at one of the sites -- an insurance company could, conceivably, collect the cookie data from an ad network and use it in a quiet decision to decline an application for a policy. (Advertising networks insist that they do not sell data for such purposes.)

Third-party cookies were precisely the kind of tracking mechanism Mr. Montulli had tried to prevent through his privacy measures. He describes it today as a surprise -- and something of an embarrassment. "That's the one 'gotcha' we had," he recalls with chagrin.

A Hot Media Topic

By 1996, the existence of cookies and third-party cookies was becoming a hot topic in the news media and in online forums; Mr. Montulli and Netscape altered the company's browsers to distinguish cookies coming directly from the site being viewed from third-party cookies and to give consumers some control over them, allowing them to turn off all cookies or just the third-party variety. Microsoft, too, implemented some cookie control tools over time. But by default, browsers were set (and are still set) to accept such cookies automatically unless the user told the software not to -- which meant that a great majority of people ended up accepting cookies unknowingly from nearly every site they had visited.

The Internet Engineering Task Force was pursuing a different tack, however, recommending in 1997 that browsers be set to block any cookie that did not come directly from the site being visited.

Mr. Kristol said that the response from the advertising companies, which were by then well established, was: "This is terrible. This will destroy our business." Each argument caused further delay -- time in which the advertising companies became more powerful and the market crystallized around the two leading browsers.

Mr. Kristol was not surprised, then, that neither Netscape nor Microsoft took to heart the recommendation that browsers block cookies unless instructed not to. He acknowledged that there was little he could do to persuade companies to adopt the voluntary standards. "There's no Internet police going around knocking on doors and saying, 'Excuse me -- the software you're using doesn't follow I.E.T.F. standards.' "

By then, Mr. Montulli said he had drifted away from the process, saying that the working group had, in fact, called for the kinds of technical changes that companies would not comply with. "I was hoping we'd get some kind of incremental improvement" out of the working group, he said -- ideas like the cookie control mechanisms he was working into new versions of the browser.

"But what the new standard required," he said, "was that you start over."

To Mr. Montulli, the conflict came down to the differences between pure researchers like Mr. Kristol and commercial engineers like himself. "The cold reality of the software business is you have to ship something that's good enough and get it out there," he said. "That's the way you ship software, and hopefully make money. If you wait forever trying to make something perfect, you may never ship."

In an article that Mr. Kristol prepared for Communications of the Association for Computing Machinery, the journal of the leading computer science professional organization, he said several factors kept him on his somewhat quixotic task. On one level, "I simply wanted to see the effort through to an appropriate completion," he said. But in his paper, Mr. Kristol -- who recently retired from Bell Laboratories -- writes, "Feeling I was being bullied" by the industry "made me more determined to persist, and I didn't like to see an attempt to bully the I.E.T.F., either."

If nothing else, the effort raised the visibility of the issues underlying cookies, Mr. Kristol said. Thanks in part to his group's work, he said, companies can't violate consumer privacy, or even appear to, without attracting unwelcome attention.

He cited the controversy that arose when DoubleClick announced in 1999 that it had bought Abacus Direct, a company that maintained a database of the buying habits of 88 million catalog shoppers, and planned to match and merge some of the data that it was collecting online with the offline data from Abacus. The resulting data trove would portray millions of consumers' habits at a level of detail unparalleled in its intimacy.

A Public Outcry

Public outcry over the plan was fierce, and the Federal Trade Commission began an inquiry into the company's practices. DoubleClick abandoned the plan, and the Federal Trade Commission dropped its inquiry. DoubleClick's chief privacy officer, Jules Polonetsky, said, "Companies are learning from the missteps of the past year, and are obligated to bake privacy into the infrastructure of their new products lest they face the wrath of the critics."

Mr. Montulli, now 30, has since gained a measure of fame -- not just as the inventor of the cookie, but also as one of People magazine's runners-up for "sexiest man alive" in 1999. He says that he has dialed back from the 120-hour work weeks at Netscape -- a punishing life that contributed to the breakup of his marriage to the daughter of Netscape's founder, Jim Clark, in 1997.

He left Netscape in 1998, a millionaire many times over thanks to the company's high-flying stock. He helped to create epinions.com, a site for comparison shopping, but has since left that company as well.

Ask about his latest achievement, and he talks about climbing Mt. Shasta with his girlfriend, Ashley Dearruigunaga -- and, at the summit, asking her to marry him. ("At 14,162 feet, I figured she couldn't say no," he said.)

When it comes to cookies, he says that he is satisfied with the way things have worked out. Even though he does not favor the use of third-party cookies, he calls the existence of third-party cookies "the best possible error," because "the only way it could be exploited is by someone who is extremely public, who is extremely large and who has a very long reach" -- a company, in other words, that cannot afford a public relations fiasco, he said.

Over time, the views on cookies from privacy advocates have evolved. Richard M. Smith, the chief technology officer for the Privacy Foundation, a think tank in Denver, said that he now believed that most cookies were benign.

"My first reaction was, 'Oh they're terrible!' Over the last year and a half as I've looked at the Internet and how it works, it would be very difficult to have the Internet without them."

About the Series

Over the next two days, articles in this series will examine public opinion about online privacy and how companies have tried to profit from it, and also how the privacy issue is being dealt with in Washington.

A version of this article appears in print on , Section A, Page 1 of the National edition with the headline: Giving Web a Memory Cost Its Users Privacy