Presented by Selena Deckelmann, Sr Director of Firefox Browser Engineering
Oct 1, 2019
View Source Conference, Amsterdam

Speaker notes (and... these aren't edited— just cut and paste from the slide deck. I'll edit and condense next week)

Hi! Thanks to the View Source staff for having me here. There are a lot of overlapping themes to our talks — I also reference the Simpsons, and have an anecdote about Mike Taylor — and like many of the other presenters, I love the web. I saw my first webpage in 1994 and it was for some electronic media mashup artist that was super weird and it blew my mind. The web is a wild and creative place, as well as a business place. It's a place where people from all over the world connect. It's an incredible political and social experiment.

The web enabled me to see this news story. This story had everything — a crime, a jailed bird and a stale piece of bread. And let me just zoom in a bit.

If you missed this — the bird is in a jail cell because its owner was caught in a robbery, and the police didn't have a bird cage, so they put it in a cell and gave it water and a bit of bread. It's a Lovebird (originally misidentified I think as a parakeet) — but it wasn't misidentified because of the bar across it's eyes. Privacy is a complicated thing, and bless GDPR or the ePrivacy Directive for contributing to create this image. I just want to let you know if you want to protect a bird's privacy, you're going to need to disguise the feather patterns, perching style, length of tail and size and kind of beak. While we're here, I feel compelled to say that a bar across the eyes will almost never protect a bird's privacy.

In the words of Bruce Schneier, I'm a short term pessimist and a long term optimist. I feel incredibly lucky to get to work on Firefox. The threats and problems I'll talk about are faced by everyone, and many of the solutions I'll talk about are things every browser does. Some of the work Mozilla does, however, is uniquely from Mozilla and springs from our mission. The backdrop here is a web increasingly influenced by big data and psychographic-driven marketing. There's some fundamental questions to be debated. I'm just going to scratch the surface, but I do think that browsers are uniquely situated to have a technical and policy impact on massive data collection, and we all have a stake in what comes next.

Before I can talk about the web and privacy, I feel like I should explain how I think about privacy. Privacy is the how we form our sense of self, become individuals and form relationships. It's also how we form critical thoughts, and then decide whether we share them. Sometimes we invent technology that helps with this — as a manager, I spend a lot of time in meetings with people all over the world over video. As a person with a lifelong problem with controlling my face and reactions, being able to face mute in meetings is a critical feature for me.

Mozilla's founders believed privacy to be such an important concept that they enshrined it in our manifesto here. Mozilla Privacy Principle #4 is Individuals' security and privacy on the internet are fundamental and must not be treated as optional. This principle guides much of the privacy technical work we do, balanced against social and political realities. We're working to develop products that demonstrate our world view, but also make effective, global policy change that puts people's rights at the center.

But privacy is more than a personal right to choose. In order to exercise privacy, we need to be free from surveillance. There's this great paper called "What Privacy is for" by Julie E Cohen. And she makes the argument that privacy is not just an individual right, or simply the right to choose to be private, but it's a collective right that requires nuanced and sometimes unprincipled thinking to defend.

She also describes the purpose of privacy in a way that I fundamentally agree with - which is that privacy ensures that — "the development of subjectivity and and development of communal values don't proceed in lockstep". This is the source of the creative tension behind change in culture and making a better world. And it's incredibly personal to me as a child of a Chinese man and a white American woman — without privacy and freedom of thought, it's likely that I literally would not exist. It's also true that privacy and freedom are needed in order for companies like Mozilla to exist.

One of the critical definitions Julie Cohen shares is that surveillance is attention that is purposeful, routine, systemic and focused. The purpose of modern surveillance — the purpose that I feel like I'm fighting the most today —is to reduce volatility and ensure sustainable profits. Government overreach is certainly a problem, but the existential threat we have today comes from systems that nudge, to try to smooth out profits and make people's behavior more predictable. She calls this process "Modulation" — which is an extremely powerful tool to shape people's behavior. I watched a video of Alexander Nix from Cambridge Analytica on my flight here, and in it he repeated the claim that people are willingingly — as individuals — choosing to share their personal information as inputs into this system. I think he's right to a point — but where he's not right is that most people do not at all understand the conclusions being drawn, or the feedback loops that are in place that make use of their personal information and the surveillance that goes along with it. So, let's have a look — what kinds of things are pervasive surveillance used for?

I still am an instagram user and have a few friends with kids who swap photos. The system has decided that I need some new shoes, so I've seen the same ad for some shoes now about 100 times. I know a lot about this stuff, and STILL one day I found myself on the shoe site starting to give them my credit card to buy a pair of shoes that I really didn't need. They did move me, despite my personal rule not to buy shoes that haven't actually already been on my feet.

So here's another example of how the ad surveillance network is functioning. We all hear about creepy ads where people swear their device must have been listening to them in order to have

served up the ads for something incredibly specific to them in the moment - Mike Taylor told me about driving and discussing a place to have a meal, and when he opened his phone to look up the address, he found that his maps app had already helpfully located the place he'd been discussing and dropped a pin on the location. A friend's husband recently suffered a serious injury to his hand and he got this ad for a poster about safety the next day… with homer and an injured hand? So clearly some system in Facebook learned something from his posts, probably that he had an injury - maybe specifically to the hand? and someone at Grangier had the good fortune to have bought ads maybe for "interest in hand injuries" for this safety poster. But could we get even more specific?

If you haven't seen this site, It's called the Spinner. I and people I know are pretty sure it's a parody based on some silly interviews the site operator has given. And the idea is for $49 you can buy ad time that targets a specific person. As you'll see from some of the profiles we know are created for targeted ad buying that I'm going to show later, it's not out of the question that you could figure out how to buy an ad whose parameters could be used to ensure that a specific person, within an increasingly smaller group, would see it in the moment you wanted them to.

So these are all kinda goofy examples, that are real world, and with the exception of the Spinner — which I think is a fake site — really don't do much harm to people. I'm sure some of you have opinions about ad targeting that may be more like hard line than mine, but I wanted to take a moment just to get out there a couple issues that give me pause, and really motivate my work.

I'm a woman and those shoes I was seeing ads for are for women, so makes sense that they'd want to target women for their ads. Harmless enough.. But it turns out that this kind of audience segmentation can be used in ways that are harmful.

This is a pretty obvious abuse of the profiling, and rose to the level of getting covered by this great organization called ProPublica that specializes in investigations of abuse of data. And on September 24, they learned that the Equal Employment Opportunity Commission (a regulatory agency in the US) investigated the claims made by ProPublica and was forcing Facebook to restrict the use of job, housing and a few other categories to reduce the obvious discrimination harm.

Another reason why I care about ad targeting has to do with this thing called profiles. Here's a list of just a few of the possible categories - with a link to the full list. Some of these are problematic alone, and in combination they can be particularly concerning. These are the conclusions drawn by algorithms about people, turned into a terse label for advertising. They can be eerily accurate about extremely personal things. Some profiles being generated and used for real time ad buying systems are extremely problematic. We could assume that all advertisers have our best interests in mind, and only would be showing us ads for things that would help us in the conditions described, but it turns out that we have adversarial users of our ad buying systems, and on social media for that matter. On these systems the people buying ads aren't trying to get you to buy shoes anymore, they're lying about who they are, faking

information with the intent to influence political views and actions, and influencing elections. Stopping the data collection that leads to these profiles being useful is definitely part of the work of my team.

Given all that — how is it that Firefox helps? And I want to say - some of the things I'll talk about are either already part of every browser, or they could be. I think about this as a combination of ensuring basic security of the browser, fighting surveillance and finally creating better defaults paired with education about that.

The first thing we have to ensure is that we maintain security for our users, despite the fact that the browser is this place where we run untrusted code from people we don't know, and all just hope for the best. Hope may be underselling it a bit, as there's a ton of work that goes into this.

Here's some of the big areas of work today: sandboxing to restrict access to APIs that could be abused, process-level isolation of sites (called site isolation in chrome, and our project to do this is called Fission), we fix a lot of security bugs, and we use the Safe Browsing API to prevent users from accidentally visiting malware and phishing sites. We also write things in Rust, or sometimes rewrite. For example, we're implementing a new protocol (HTTP/3 or QUIC) in Rust that's just starting to be integrated now. We're also thinking about interesting uses of WebAssembly in the browser, where we could get lighter-weight sandboxing of things like 3rd party libraries. And honestly there's a lot more we do here — both in terms of standards involvement, implementation, our build process and collaboration with Tor.. and most other browsers do most of those things as well at this point because we all take security extremely seriously.

Once we secure the browser, we have privacy features that we implement. And those features don't exist in a vacuum. They're often accompanied by policies, open process and maybe even a community to support them.

Tracking Protection is one of the ways Firefox fights surveillance in the browser. Safari has ITP, which is a local machine learning approach to blocking that we think is cool. What Tracking Protection does in it's most strict form is simply not load trackers, blocking them from loading at the network level, and in it's least strict form it blocks trackers from having read access to local data and may strip cookies of sites we find trying to circumvent the system. We operate this system by working with a third party - Disconnect - that creates and maintains the lists, and established an published criteria for adding and removing sites. This is part of our tech combined with a transparent and accountable policy approach that underpins everything we do.

Another thing my team works on is the Certification Authority (CA) Root Program and the Common CA Database. This project is run as a radically open forum that manages which organizations are trusted by Firefox to provide SSL/TLS certificates. Our program is considered the gold standard for deliberative and open process in this industry, and there are many organizations that simply follow our lead, or wholesale use our root list to secure their HTTPS

implementation. No other trust program is open like ours, and every major CA root administrator follows our mailing lists to understand what our program and the public thinks of policies created and implemented there. Our policies about trust, accountability and transparency underpin the security of HTTPS through this program.

TLS 1.3 and DoH are both network-level protocol changes to the internet that improve security. TLS 1.3 also improved performance significantly. Changes from 1.2 to 1.3 make it harder to intercept and control web traffic without users knowing. This is a protocol that every browser has adopted! DNS over HTTPS or (DoH) is another internet protocol, and it is in the process of being adopted by Chrome and Firefox. DoH is a way of tunneling DNS requests through HTTPS that both encrypts and makes it difficult to eavesdrop or block the requests. Firefox's approach is again to combine a strong privacy policy (which we call our Trusted Recursive Resolver policy) which restricts storage of DNS data by TRRs to 24 hours, and ensures accountability and transparency for any DNS blocking lists that are created, which is very similar to our approach to working with Disconnect on tracking protection. Handling DNS data is very different than blocking trackers, however, and so we're taking our rollout of this feature slow, starting in the US only, while we work to find good TRR partners throughout the rest of the world. Our hope for deployment is that eventually DOH and/or DNS-over-TLS will become the defaults, and encrypted DNS traffic will be available to all people.

We're also working on Firefox Private Network (Currently available to US-users only)— which is currently operating as a free secure proxy into the browser. Our goal here is to provide protection to users by hiding their IP address from trackers. Providing this is part of a layered approach — where we enable Enhanced Tracking Protection by default, continue to implement other privacy enhancing features that reduce surveillance and help people hide their IP addresses while they browse - which are a signficant source of unique identification of users.

So the last bit of anti-surveillance work we do is work to educate people about it - in the product, as well as with educational publications and policy papers. I'm just going to show three things — but I realized when doing this talk that we do quite a bit more and I should create a list somewhere for reference. Which I'll do after the conference.

A new thing that you can see in Nightly and Beta today is our Protection Report. This shows you how many trackers have been blocked by your system, provides links to support.mozilla.org articles and explains more about the kinds of sites we're blocking. We're adding more to this report over time for other services we provide.

Something that the Mozilla Foundation works on is the Internet Health Report. They pick different topics each year, but surveillance was a big part of the 2019 report. These are short pieces with collections of primary sources to give people a fast way to get a summary, and then be able to immediately dive into curated sources of information about an important topic.

I also wanted to share Firefox Monitor. It helps alert you if your email has been compromised in a website breach. This product's purpose is to help people with a common problem that they struggle with — website account compromises — and give them an easy and secure way to keep an eye on it for themselves. We support multiple email account monitoring from a single Firefox Account and we send out alerts when new breaches that affect you occur. We've hit about 4 million subscribers at this point, and it's one of our more popular services.

So we're close to the end here - just a few more slides… Privacy is essential to being human — It's how we grow and change, make our world a better place.

And the web is part of our world — so privacy is an essential part of being human on the web.

But what we know is that the web wasn't designed with privacy in mind. It wasn't designed with surveillance in mind. It was designed without thinking of the threat model, really. The threat models have evolved, and we're still trying to catch up. We now need to disrupt the system that's emerged and we need to protect everyone, including ourselves from buying extra pairs of shoes we don't need, being discriminated against for jobs or housing, and having our essential selves being labeled and nudged in scary or even dangerous ways by feedback loops in advertising systems.

So what can we do to help preserve our privacy?

I believe that browsers can provide understandable privacy properties that fight surveillance and support a sustainable future. Firefox does a lot here on a tech and policy basis, and other browsers do a lot to. I think we could all do more. We can do it through more encryption. We can do this through policies in the browser, standards and by blocking harmful tracking and enabling user choice.

I also believe every person and company can take action to reduce the amount off data we collect. Mozilla has a site that describes our lean data practices, our data stewardship programs and we welcome anyone to reach out to us to learn more about how we implement this at mozilla.

And what about our future — the one that Henri shared yesterday with an ever expanding world that has phones and browsers on phones, but also those phones have apps. What about those?

We know that apps are invading people's privacy basically without restraint at this point. Unlike the browser, there's nothing but monopoly-like app stores controlling policies and reviewing apps for violations. We're really still in the wild-west stage here, and it's likely going to get worse before it gets better. So I want to close with this:

We're at a point where we have an opportunity to make choices, and think deeply about what happens next. We don't have to accept massive surveillance and violation of privacy as

inevitable. What kind of future do you want? We are on a path to a modulated one — modulated by ad tech and profiling. I'm asking you - uninstall apps, and use the browser, especially on the mobile web. Support PWAs and alternatives that provide understandable privacy properties. We can all be part of a better and more private web, so please use the browser. I'd love for you to pick Firefox. But, and I know I'm preaching to the choir here — but choose the web. It's worth your time to think about what happens next in this ecosystem. And we're hiring! You're welcome to join us in this fight.

Thank you to Eric Rescorla, Laura Thomson, Luke Welling, Tanvi Vyas, Steve Englehardt and Mike Taylor for assistance with this talk. And the bird.