The disruptive effects of user privacy

Chris Mitchell

www.chrismitchell.net

Agenda

- Where are we now?
- The privacy goal
- Disruption?
- What will really happen?

A lack of privacy

- We are all accustomed to the idea that what we do online is not very private.
- We may not know exactly who knows what, but we do know from personal experience:
  - service providers monitor our *activity* and use it to target advertising;
  - *activity* includes where we browse on the web, our past purchases, the contents of our emails, and other factors we may not be aware of …

Some things aren't obvious

- The means by which we are tracked is not so clear, at least to most internet users.
  - A minority of us understand how cookies can be used to track repeated visits to the same website, and also, through the referrer field and links embedded in web pages, how advertisers can track us.
  - A smaller minority understand that, even if cookies are disabled, fingerprinting techniques enable web servers to uniquely identify platforms (see next slide).
  - Of course, IP addresses help with fingerprinting, but the use of anonymising routers doesn't stop fingerprinting.

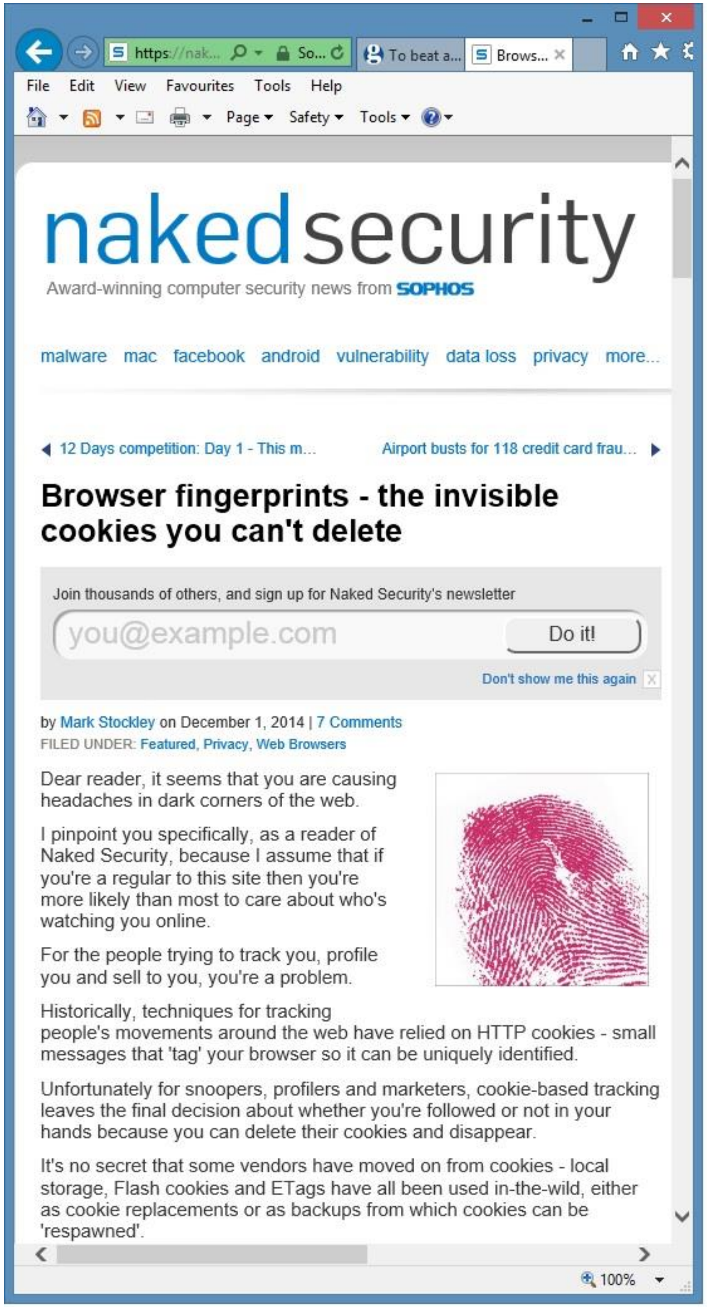

nakedsecurity

Award-winning computer security news from SOPHOS

malware
mac
facebook
android
vulnerability
data loss
privacy
more...
12 Days competition: Day 1 - This m...
Airport busts for 118 credit card frau...

Browser fingerprints - the invisible cookies you can't delete

Join thousands of others, and sign up for Naked Security's newsletter

you@example.com

Do it!

Don't show me this again

by Mark Stockley on December 1, 2014 | 7 Comments

FILED UNDER: Featured, Privacy, Web Browsers

Dear reader, it seems that you are causing headaches in dark corners of the web.

I pinpoint you specifically, as a reader of Naked Security, because I assume that if you're a regular to this site then you're more likely than most to care about who's watching you online.

For the people trying to track you, profile you and sell to you, you're a problem.

Historically, techniques for tracking people's movements around the web have relied on HTTP cookies - small messages that 'tag' your browser so it can be uniquely identified.

Unfortunately for snoopers, profilers and marketers, cookie-based tracking leaves the final decision about whether you're followed or not in your hands because you can delete their cookies and disappear.

It's no secret that some vendors have moved on from cookies - local storage, Flash cookies and ETags have all been used in-the-wild, either as cookie replacements or as backups from which cookies can be 'respawned'.

Demands for greater privacy

- Whilst our activities can be readily tracked using a variety of means, there is also great pressure to change this, e.g. from
  - **legislators**, e.g. the European Union, who wish to protect citizen’s privacy;
  - **pressure groups** of many types, arguing in favour of greater end user privacy;
  - **standards** and other **guidelines**, which set down codes of behaviour and best practices for websites.

Privacy technologies

- Supporting these demands for greater privacy are a range of technologies that help support privacy, e.g. including:
  - encryption;
  - good practice schemes such as the ‘do not track’ HTTP header field;
  - anonymising routers;
  - anonymous credential systems and other *special* cryptographic schemes;
  - homomorphic encryption.

Technology versus regulation

- In practice we tend to largely rely on regulatory/legal compliance solutions.
- This assumes those with access to our personal data will behave in accordance with law/regulation.
- This may be a questionable assumption.
- Some in the academic community advocate a purely technological solution, arguing that technology could prevent any misuse of personal data, for whatever reason.
- However, the consequences of such an approach, if it could ever be realised (which is a big if), are profound.

Agenda

- Where are we now?
- The privacy goal
- Disruption?
- What will really happen?

No traces

- Perhaps the ultimate goal of privacy advocates is to enable us all to leave no identifiable trace of our activities, if that is what we want.
- Some would suggest that such an arrangement should even be the default, given that many users have limited technical expertise.

Difficulties in definition

- However, defining what *no trace* means is problematic.
- To some extent almost everything we do partly identifies us, e.g. we indicate our language, interests, …
- Some activities automatically reveal our unique identity, e.g. when we use a credit card for payment.
- Perhaps the key property is linkability of activities, or rather unlinkability.

Anonymisation & pseudonymisation

- These difficulties highlight the difficulties in effectively anonymising personal data.
- Such anonymisation has clear benefits, allowing large data sets to be analysed, e.g. to identify new treatments for illness, new solutions to complex problems, etc.
- However, the risk of de-anonymisation is always present, so anonymisation needs to done with great care.
- Anyway, this is a bit of a side track …

Agenda

- Where are we now?
- The privacy goal
- Disruption?
- What will really happen?

Let’s imagine …

- Suppose the privacy advocates are completely successful, and by default all our activities are unlinkable (except where necessary).
- That is, suppose we can all use the Internet knowing that, unless we choose to reveal who we are, it is technologically impossible to link our various interactions with third parties.

Who would be impacted?

- The service providers would lose their ability to link one user interaction with another, severely limiting their ability to target advertising.
- It would also have an impact on security (of both users and service providers) in a variety of ways.
- We next look at these impacts in a little more detail …

No more free stuff

- Many of the free web services we use on a daily basis are funded through advertising, e.g. search, cloud storage, social networks, messaging (email and instant), voice over IP, …
- Loss of targeted advertising could severely impact revenues for these service providers.
- Perhaps we will have to start paying for all these services?
- Maybe service providers will simply vanish?

Degraded intrusion detection

- Network intrusion detection systems (NIDSs) typically examine DNS messages.
- If DNSsec encryption is deployed, enhancing privacy, then such messages become opaque to the NIDS.
- That is, by concealing traffic, detecting intrusions becomes more difficult.

Degraded authentication

- Browser fingerprinting has both positive and negative aspects.
- Clearly it negatively impacts user privacy.
- However, it is also widely used as a means of enhancing user authentication, by verifying that a user is working via a known platform.
- That is, if browser fingerprinting was made impossible (actually, very difficult to achieve for anyone other than an expert user) then user authentication would be made less effective.

Degraded forensics/accountability

- As is well known, effective user anonymity makes ensuring that users are held accountable for their actions very difficult, if not impossible.
- That is, efforts to investigate security breaches may be made very much more difficult if all the activity records are anonymised.
- More generally, criminal investigations may be made much more difficult.
- Legal interception may also be made much less valuable to investigators.

Agenda

- Where are we now?
- The privacy goal
- Disruption?
- What will really happen?

A good question …

- Firstly I should say that this talk is not intended as an argument against enhancing user privacy – it is just pointing out some of the implications.
- In fact, implementing complete unlinkability is theoretically possible but very difficult to achieve in the real world. For example:
  - our browsers leak vast quantities of information about us;
  - few of us even know what anonymising routers are or what the threat is that they address, let alone use them;
  - it is not very practical to expect users to start with a clean OS install every time they browse the web.

Likely scenarios I

- Choice of less privacy or payment for service.
- Whilst users say they value their privacy, in practice they appear to be reluctant to spend money to do so.
- AT&T allows gigabit service subscribers to opt out of deep packet inspection – for a $29 fee per month.
- Apparently most users do not pay the extra.
- Not everyone approves (see next slide)!

Information Security Group

ROYAL HOLLOWAY UNIVERSITY OF LONDON

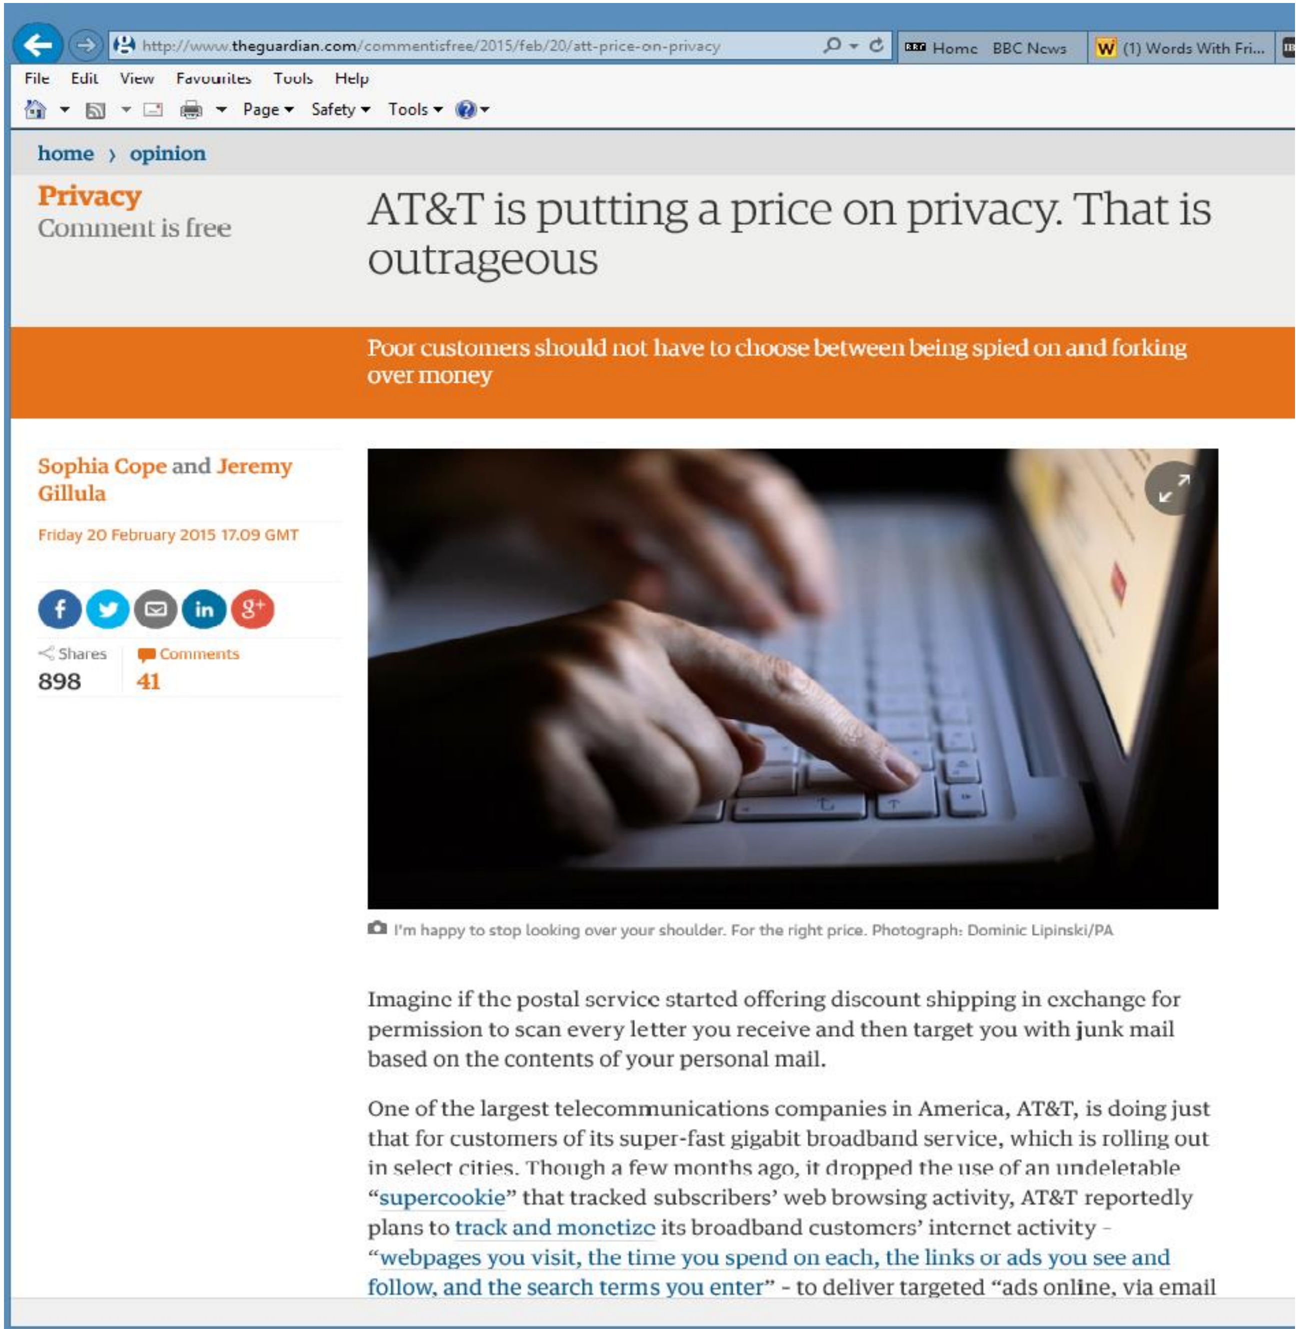

Privacy
Comment is free

AT&T is putting a price on privacy. That is outrageous

Poor customers should not have to choose between being spied on and forking over money

Sophia Cope and Jeremy Gillula

Friday 20 February 2015 17.09 GMT

Shares 898
Comments 41

I'm happy to stop looking over your shoulder. For the right price. Photograph: Dominic Lipinski/PA

Imagine if the postal service started offering discount shipping in exchange for permission to scan every letter you receive and then target you with junk mail based on the contents of your personal mail.

One of the largest telecommunications companies in America, AT&T, is doing just that for customers of its super-fast gigabit broadband service, which is rolling out in select cities. Though a few months ago, it dropped the use of an undeletable "supercookie" that tracked subscribers' web browsing activity, AT&T reportedly plans to track and monetize its broadband customers' internet activity – "webpages you visit, the time you spend on each, the links or ads you see and follow, and the search terms you enter" - to deliver targeted "ads online, via email

Likely scenarios II

- If law enforcement and other government agencies cannot access data via interception, then they are likely to try other methods.
- These other methods may be more intrusive.
- There has been much recent discussion of malware distributed by western governments – see, for example, the next slide …

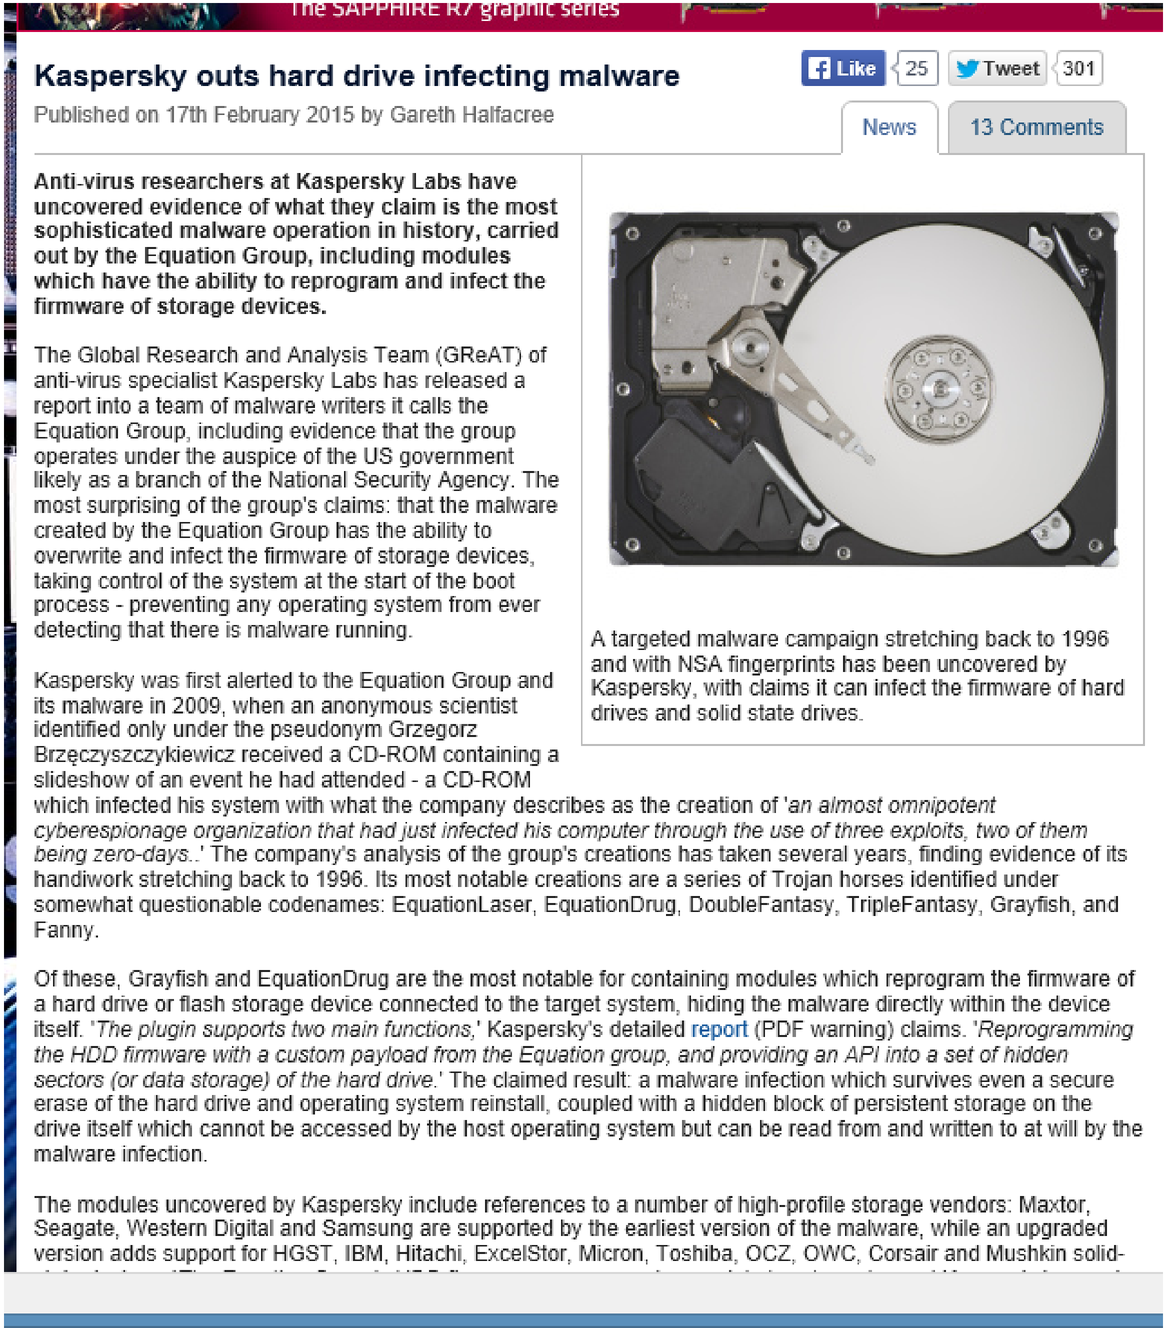

Kaspersky outs hard drive infecting malware

Published on 17th February 2015 by Gareth Halfacree

Anti-virus researchers at Kaspersky Labs have uncovered evidence of what they claim is the most sophisticated malware operation in history, carried out by the Equation Group, including modules which have the ability to reprogram and infect the firmware of storage devices.

A targeted malware campaign stretching back to 1996 and with NSA fingerprints has been uncovered by Kaspersky, with claims it can infect the firmware of hard drives and solid state drives.

The Global Research and Analysis Team (GReAT) of anti-virus specialist Kaspersky Labs has released a report into a team of malware writers it calls the Equation Group, including evidence that the group operates under the auspice of the US government likely as a branch of the National Security Agency. The most surprising of the group's claims: that the malware created by the Equation Group has the ability to overwrite and infect the firmware of storage devices, taking control of the system at the start of the boot process - preventing any operating system from ever detecting that there is malware running.

Kaspersky was first alerted to the Equation Group and its malware in 2009, when an anonymous scientist identified only under the pseudonym Grzegorz Brzęczyszczykiewicz received a CD-ROM containing a slideshow of an event he had attended - a CD-ROM which infected his system with what the company describes as the creation of '*an almost omnipotent cyberespionage organization that had just infected his computer through the use of three exploits, two of them being zero-days.*.' The company's analysis of the group's creations has taken several years, finding evidence of its handiwork stretching back to 1996. Its most notable creations are a series of Trojan horses identified under somewhat questionable codenames: EquationLaser, EquationDrug, DoubleFantasy, TripleFantasy, Grayfish, and Fanny.

Of these, Grayfish and EquationDrug are the most notable for containing modules which reprogram the firmware of a hard drive or flash storage device connected to the target system, hiding the malware directly within the device itself. '*The plugin supports two main functions,*' Kaspersky's detailed report (PDF warning) claims. '*Reprogramming the HDD firmware with a custom payload from the Equation group, and providing an API into a set of hidden sectors (or data storage) of the hard drive.*' The claimed result: a malware infection which survives even a secure erase of the hard drive and operating system reinstall, coupled with a hidden block of persistent storage on the drive itself which cannot be accessed by the host operating system but can be read from and written to at will by the malware infection.

The modules uncovered by Kaspersky include references to a number of high-profile storage vendors: Maxtor, Seagate, Western Digital and Samsung are supported by the earliest version of the malware, while an upgraded version adds support for HGST, IBM, Hitachi, ExcelStor, Micron, Toshiba, OCZ, OWC, Corsair and Mushkin solid-

What should we be aiming for?

- Getting the balance right is very difficult!
- Even the most strident advocates of technological privacy solutions do not suggest the legal/regulatory/compliance approach should be abandoned, and this surely will continue …
- … as will development of best practice guidelines/standards.

Impact of privacy technology

- Privacy technology will continue to advance, and some will no doubt be deployed.
- However, I fully expect security agencies and others to continue to develop ways round these deployed technologies.
- Of course, highly skilled and highly determined individuals can, as now, make their activities pretty private, but they are essentially irrelevant to the argument.
- So probably not very much will change (and the promised disruption won’t happen), unless legislators demand it.
- But the potential for huge disruption remains …