Is Tor still secure after Silk Road?

6 February 2015, by Steven J. Murdoch

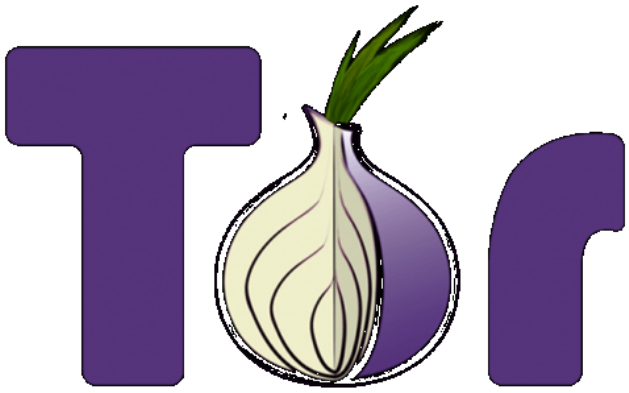

Tor, your online an-onionising software. Credit: Tor Project, CC BY-SA

The Silk Road trial has concluded, with Ross Ulbricht found guilty of running the anonymous online marketplace for illegal goods. But questions remain over how the FBI found its way through Tor, the software that allows anonymous, untraceable use of the web, to gather the evidence against him.

The development of anonymising software such as Tor and Bitcoin has forced law enforcement to develop the expertise needed to identify those using them. But if anything, what we know about the FBI's case suggests it was tip-offs, inside men, confessions, and Ulbricht's own errors that were responsible for his conviction.

This is the main problem with these systems: breaking or circumventing anonymity software is hard, but it's easy to build up evidence against an individual once you can target surveillance, and wait for them to slip up.

The problem

A design decision in the early days of the internet led to a problem: every message sent is tagged with the numerical Internet Protocol (IP) addresses that identify the source and destination computers. The network address indicates how and where to route the message, but there is no equivalent indicating the identity of the sender or intended recipient.

This conflation of addressing and identity is bad for privacy. Any internet traffic you send or receive will have your IP address attached to it. Typically a computer will only have one public IP address at a time, which means your online activity can be linked together using that address. Whether you like it or not, marketers, criminals or investigators use this sort of profiling without consent all the time. The way IP addresses are allocated is geographically and on a per-organisation basis, so it's even possible to pinpoint a surprisingly accurate location.

This conflation of addressing and identity is also bad for security. The routing protocols which establish the best route between two points on the internet are not secure, and have been exploited by attackers to take control of (hijack) IP addresses they don't legitimately own. Such attackers then have access to network traffic destined for the hijacked IP addresses, and also to anything the legitimate owner of the IP addresses should have access to.

This is why those looking for cat videos on YouTube on February 24, 2008 found themselves at Pakistan Telecom instead, why hackers made off with US$83,000 worth of bitcoin between February and May 2014 by impersonating the legitimate owners, and why hundreds of organisations found their communications mysteriously routed via computers in Belarus and Iceland in 2013.

Redesigning with security in mind

Onion routing was developed to correct these mistakes, separating identity and address so that it's possible to communicate through the internet without revealing the IP address used. Originally a

US Navy Research Laboratory project, the latest implementation of onion routing is known as Tor and is independently developed by the non-profit Tor Project.

Tor routes internet traffic through three or more intermediate computers called nodes, which prevents anyone listening in – and any website the traffic connects to – from knowing the source of the traffic or working out who is communicating with whom. Even Tor nodes aren't individually aware of the details of which user, where, is connecting to what. The first node sees the user's IP address, the last knows which site is being accessed, but unless both the first and last nodes are controlled maliciously these two facts won't be linked.

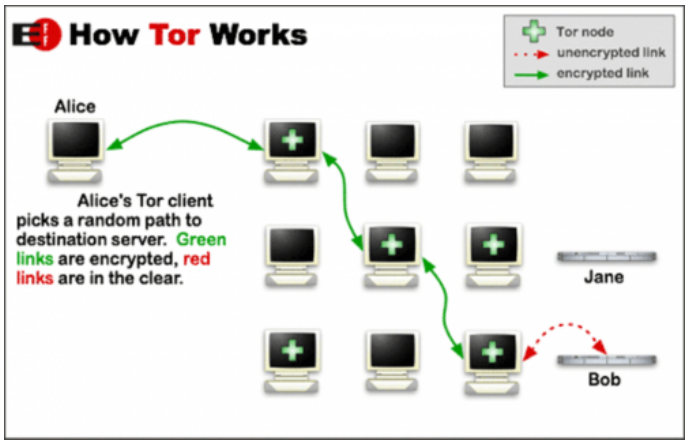

How Tor works. Credit: Tor Project/EFF, CC BY

Who uses Tor?

There are all sorts of reasons to use Tor to protect privacy: law enforcement monitoring criminals, firms studying potential takeover targets, those who don't like advertisers profiling them, or political activists in authoritarian states. An increasing number use Tor to access websites that are blocked in their country, as Tor's anonymisation prevents the censor's software from detecting the traffic is destined for a banned website.

As well as websites on the everyday internet, Tor allows the creation of hidden services: websites accessed only through the Tor network, of which the Silk Road is an example. This ensures privacy and security by identifying sites not with an IP address and domain name but with a cryptographic key. Without this key, there's no way for a would-be eavesdropper to impersonate the real website and intercept traffic directed to it. These are represented by a URL ending in .onion – accessed with a Tor-enabled browser, Facebook is at facebookcorewwwi.onion.

An evolving architecture

Tor isn't perfect. It can't protect traffic that has left the Tor network, for example, where traffic becomes vulnerable to all the usual attacks. The solution to this is end-to-end encryption – preventing monitoring or tampering not only for Tor users but for everyone else on the internet too.

Another potential weakness is flaws in other software used with Tor. The FBI distributed malware to every visitor to a group of hidden services, some of which claimed to distribute child abuse images. The malware exploited a vulnerability in the Firefox web browser in order to send the real IP address of the user and other identifying information back to the FBI.

It's been suggested that a flaw in the software behind the Silk Road gave the FBI the breakthrough that let them discover the IP address and so the real location of the Silk Road's servers. But the lack of detail on this from the FBI, compared to the other evidence gleaned from Ulbricht's server and laptop computer, has led some to ask whether the FBI used techniques it doesn't want openly discussed – such as the involvement of the National Security Agency and its vast surveillance infrastructure.

It could also have been what's called a traffic confirmation attack, where the entry and exit Tor nodes are compromised or monitored. Our own research has shown that this can allow

communications to be de-anonymised, and researchers including myself are working on how to address this – by making it difficult or unlikely for any one person to control the entry and exit nodes, and by reducing the potential damage if this occurs.

The internet has come a long way since its beginnings, and simplifications and rationalisations made for good reasons at the time need to be re-visited in light of what's been learned in the 40 years since.

Conflating addressing and identity is one of these decisions. Tor – useful in its own right – also indicates how the internet's architecture should provide strong assurance of identity when needed, and strong privacy when not. Given enough resources, attackers will be able to de-anonymise at least some of Tor's users, some of the time, but it's still the best web privacy solution we have today. The next generation of systems will be better still.

Source: The Conversation