

How Tor's privacy was (momentarily) broken, and the questions it raises

December 10 2015, by Steven J. Murdoch



Will Tor's chopped onions lead to tears?

Just how secure is [Tor](#), one of the most widely used internet privacy tools? Court documents [released](#) from the Silk Road 2.0 trial suggest that a "[university-based research institute](#)" provided information that broke

Tor's privacy protections, helping identify the operator of the illicit online marketplace.

Silk Road and its successor Silk Road 2.0 were run as a Tor hidden service, an anonymised website accessible only over the Tor network which protects the identity of those running the site and those using it. The same technology is used to protect the privacy of visitors to other websites including [journalists reporting on mafia activity](#), [search engines](#) and [social networks](#), so the security of Tor is of critical importance to many.

How Tor's privacy shield works

Almost [97% of Tor traffic](#) is from those using Tor to anonymise their use of standard websites outside the network. To do so a path is created through the Tor network via three computers (nodes) selected at random: a first node entering the network, a middle node (or nodes), and a final node from which the communication exits the Tor network and passes to the destination website. The first node knows the user's address, the last node knows the site being accessed, but no node knows both.

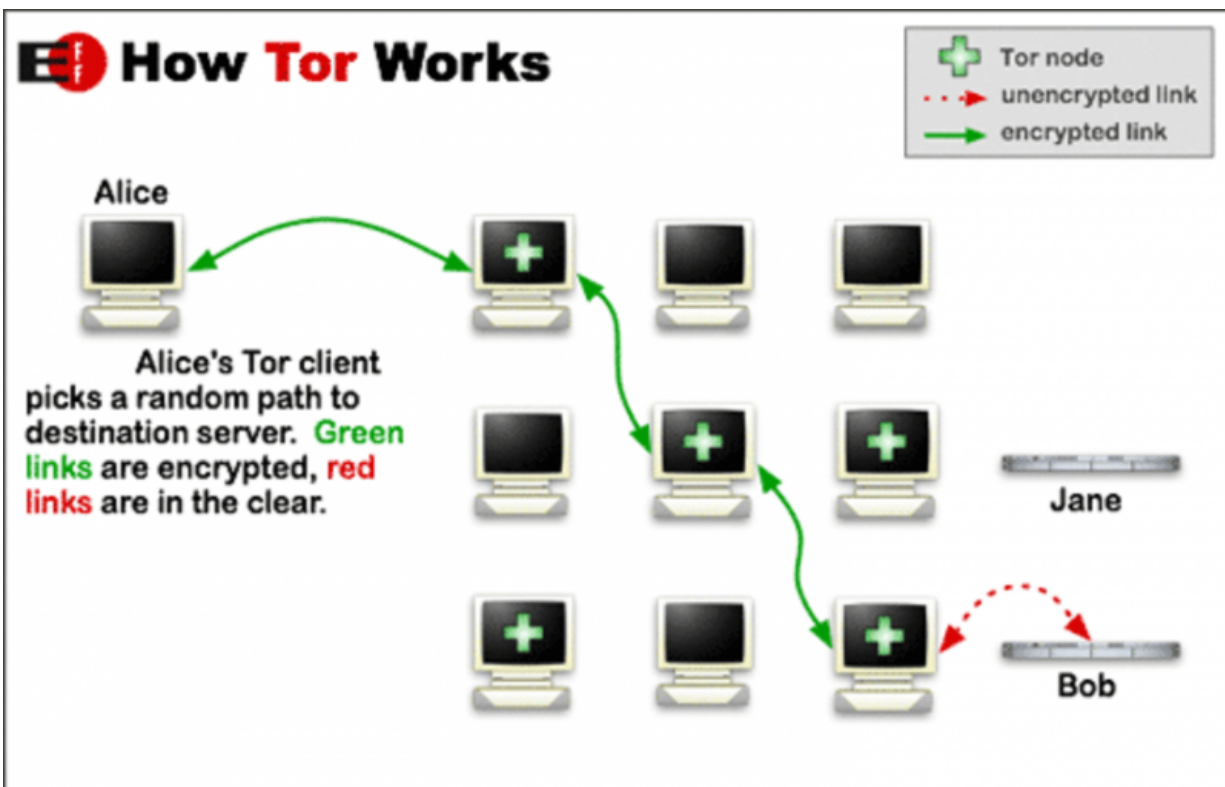
The remaining 3% of Tor traffic is to hidden services. These websites use ".onion" addresses stored in a hidden service directory. The user first requests information on how to contact the hidden service website, then both the user and the website make the three-hop path through the Tor network to a rendezvous point which joins the two connections and allows both parties to communicate.

In both cases, if a malicious operator simultaneously controls both the first and last nodes to the Tor network then it is possible to link the incoming and outgoing traffic and potentially identify the user. To prevent this, the Tor network is designed from the outset to have sufficient diversity in terms of who runs nodes and where they are

located – and the way that nodes are selected will avoid choosing closely related nodes, so as to reduce the likelihood of a user's privacy being compromised.

This type of design is known as [distributed trust](#): compromising any single computer should not be enough to break the security the system offers (although compromising a large proportion of the network is still a problem). Distributed trust systems protect not only the users, but also the operators; because the operators cannot break the users' anonymity – they do not have the "keys" themselves – they are less likely to be targeted by attackers.

Unpeeling the onion skin



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

With about [2m daily users](#) Tor is by far the most widely used privacy system and is considered one of the most secure, so research that demonstrates the existence of a vulnerability is important. Most research examines how to increase the likelihood of an attacker [controlling both the first and last node](#) in a connection, or how to [link incoming traffic to outgoing](#).

When the 2014 programme for the annual [BlackHat conference](#) was announced, it included [a talk by a team of researchers from CERT](#), a Carnegie Mellon University research institute, claiming to have found a means to compromise Tor. But the talk [was cancelled](#) and, unusually, the researchers did not give advance notice of the vulnerability to the Tor Project in order for them to examine and fix it where necessary.

This decision was particularly strange given that CERT is worldwide coordinator for [ensuring software vendors are notified](#) of vulnerabilities in their products so they can fix them before criminals can exploit them. However, the CERT researchers gave enough hints that Tor developers were able to investigate what had happened. When they examined the network they found someone was [indeed attacking Tor users](#) using a technique that matched CERT's description.

The multiple node attack

The attack turned on a means to tamper with a user's traffic as they looked up the .onion address in the hidden service directory, or in the hidden service's traffic as it uploaded the information to the directory.

When traffic bound for a hidden service directory enters the Tor network, the .onion address of the hidden service is visible. This traffic was tampered with in a way that wouldn't disrupt the request, but would

leave a trace, different for each .onion address, which could be detected when the traffic left the network. If both the first and last node selected by the Tor software for communicating with the hidden service directory were run by the attacker, the .onion address could be linked via the pattern to the identity of the user's computer, or to the computer where the hidden service is hosted.

While this technique is unreliable because it requires the attacker to control both the first and last hop, given enough time it will eventually succeed – and part of the attack was to register many new nodes to the Tor network to make it more likely they'd control both first and last node. Because hidden services are always available it's a case of repeatedly connecting to the target until the attack succeeds. This brute force attack only works with hidden services and is why they're less secure than using Tor to anonymise access to standard websites.

A lesson for the future

Carnegie Mellon has [refused to answer questions](#) over whether its researchers were involved in any attack or had any contact with the FBI. No evidence has been revealed of such contact, but the timing and technique of the attack has prompted some to ask questions about their involvement in the FBI's pursuit of Silk Road 2.0. It might be that the researchers were legally compelled to assist the FBI under some kind of warrant, although the [FBI has denied](#) it paid the researchers US\$1m for the attack. The university has stated only that it "[abides by the rule of law](#)".

The vulnerability was [fixed in July 2014](#), but protecting Tor hidden services remains an inherently difficult problem.

The affair also raises [questions about research ethics](#) and the control of surveillance by government agencies. CERT, an autonomous, federally

funded research institute, may not be [subject to the ethics review requirements](#) in the same way that university researchers are, for example. And the attack went further than just the normal practice of proof of concept: rather than taking steps to protect innocent users, the attack on Tor potentially exposed every user and hidden service operator at the time.

Research, when carried out ethically is key to improving internet privacy, and the Tor Project has always assisted researchers and [given them the benefit of the doubt](#) when experiments show up as unusual network activity, but given these events, suspicious behaviour may now be blocked when detected.

To help avoid situations that may put people at risk we need to be able to [validate experimental results without involving real people](#), and where that's impossible to have better procedures for [protecting network users](#).

This story is published courtesy of [The Conversation](#) (under Creative Commons-Attribution/No derivatives).

Source: The Conversation

Citation: How Tor's privacy was (momentarily) broken, and the questions it raises (2015, December 10) retrieved 10 April 2024 from <https://phys.org/news/2015-12-tor-privacy-momentarily-broken.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--