

A REMARKABLE DATE FOR THE VIRGIN BIRTH OF THE SILK ROAD INVESTIGATION

As Wired first reported, there's been an interesting exchange in the Silk Road prosecution. In September, the former FBI Agent who helped to bust accused Silk Road operator Ross Ulbricht, Christopher Tarbell, submitted a declaration explaining the genesis of the investigation by claiming the FBI got access to the Silk Road server because it became accessible via a non-Tor browser. In response, Ulbricht lawyer Joshua Horowitz submitted a declaration claiming Tarbell's claims were implausible because the FBI wouldn't have been able to get into Silk Road's back end. The government responded by claiming that even if it did hack the website, it would not have been illegal.

Given that the SR Server was hosting a blatantly criminal website, it would have been reasonable for the FBI to "hack" into it in order to search it, as any such "hack" would simply have constituted a search of foreign property known to contain criminal evidence, for which a warrant was not necessary .

On Friday, Judge Katherine Forrest rejected Ulbricht's efforts to throw out the evidence from the alleged hack, accepting the government's argument that Ulbricht had no expectation of privacy on that server regardless of when and how the government accessed it.

The temporal problems with the government's story

Most of the coverage on this exchange has focused on the technical claims. But just as interesting are the temporal claims. Horowitz

summarizes that problem this way:

[S]everal critical files provided in discovery contain modification dates predating the first date Agent Tarbell claims Icelandic authorities imaged the Silk Road Server, thereby casting serious doubt on the chronology and methodology of his account;

The government claims that server was first imaged on July 23, 2013.

As I'll lay out below, Horowitz and Tarbell provide a lot of details suggesting something – perhaps the imaging of the server, perhaps something more – happened six weeks earlier.

But before we get there, consider the date: June 6, 2013.

June 6, 2013 was the day after the afternoon publication of the first Snowden leak, and the day before the Guardian made it clear their leak included cyberwar materials.

That is, the FBI claims to have officially “found” the Silk Road server at the same time the Snowden leaks started, even while they date their investigation to 6 weeks later.

The June 6 materials

FBI's Tarbell is much vaguer about this timing than Ulbricht's team is. As Tarbell tells it, on some unknown date in early June 2013, he and a colleague were sniffing Silk Road data when they discovered an IP not known to be tied to Tor.

In or about early June 2013, another member of CY-2 and I closely examined the traffic data being sent from the Silk Road website when we entered responses to the prompts contained in the Silk Road login interface.

That led them to look further, according to Tarbell. When he typed the IP into a non-Tor

browser, he discovered it was leaking.

When I typed the Subject IP Address into an ordinary (non-Tor) web browser, a part of the Silk Road login screen (the CAPTCHA prompt) appeared. Based on my training and experience, this indicated that the Subject IP Address was the IP address of the SR Server, and that it was “leaking” from the SR Server because the computer code underlying the login interface was not properly configured at the time to work on Tor.

That led the government to ask Iceland, on June 12, to image the server. Iceland didn’t do so, according to the official narrative, until the next month.

The defense doesn’t buy this – in part, because Tarbell claims he didn’t adhere to forensics standard procedure by keeping copies of his packet sniffing.

Failure to preserve packet logs recorded while investigating the Silk Road servers would defy the most basic principles of forensic investigative techniques.

[snip]

[T]he government’s position is that former SA Tarbell conducted his investigation of Silk Road, and penetrated the Silk Road Server, without documenting his work in any way.

According to the government, the only record of Tarbell’s access to the server from this period is from access logs dated June 11.

[A]n excerpt of 19 lines from Nginx access logs, attached hereto as Exhibit 5, supposedly showing law enforcement access to the .49 server from a non-Tor IP address June 11, 2013, between 16:58:36 and 17:00:40. According to the

Government, this is the only contemporaneous record of the actions described by the Tarbell Declaration at ¶¶ 7-8.9

Given that this bears a particular date, I find it all the more curious that Tarbell doesn't date when he was doing the packet sniffing.

There are a number of other details that point back to that June 6 date. Perhaps most significant is that Iceland imaged a server Silk Road had earlier been using on June 6.

There are a total of 4 tarballs in the first item of discovery: home, var, all, and orange21 – all contained in .tar.gz files. The mtime for orange21.tar.gz is consistent with the July 23, 2013 image date. However, the other 3 tarballs have an mtime of June 6, 2013, as shown below²²:

- *root 30720 Jun 6 2013 home.tar.gz*
- *root 737095680 Jun 6 2013 var.tar.gz*
- *root 1728276480 Jun 6 2013 all.tar.gz*
- *root 22360048285 Jul 23 2013 orange21.tar.gz*

The modification date of the tarballs is consistent with an imaging date of June 6, 2013, a full six weeks before the July 23, 2013, imaging of the .49 Server, a fact never mentioned in the Tarbell Declaration.

Though – as the defense points out – Tarbell didn't mention that earlier imaging. He notes an earlier "lead" on the Silk Road server that resolved by May, and he notes that after Ulbricht's arrest they obtained record of him noting leaks in the server.

5 After Ulbricht's arrest, evidence was discovered on his computer reflecting that IP address leaks were a recurring problem for him. In a file containing a log Ulbricht kept of his actions in administering the Silk Road website, there are multiple entries discussing various leaks of IP addresses of servers involved in running the Silk Road website and the steps he took to remedy them. For example, a March 25, 2013 entry states that the server had been "ddosd" – i.e., subjected to a distributed denial of service attack, involving flooding the server with traffic – which, Ulbricht concluded, meant "someone knew the real IP." The entry further notes that it appeared someone had "discovered the IP via a leak" and that Ulbricht "migrated to a new server" as a result. A May 3, 2013 entry similarly states: "Leaked IP of webserver to public and had to redeploy/shred [the server]." Another entry, from May 26, 2013, states that, as a result of changes he made to the Silk Road discussion forum, he "leaked [the] ip [address of the forum server] twice" and had to change servers.

[snip]

7 Several months earlier, the FBI had developed a lead on a different server at the same Data Center in Iceland ("Server-1"), which resulted in an official request for similar assistance with respect to that server on February 28, 2013. See Ex. B. Due to delays in processing the request, Icelandic authorities did not produce traffic data for Server-1 to the FBI until May 2013. See Ex. A. By the time the FBI received the Server-1 traffic data, there was little activity on Server-1, indicating that it was no longer hosting a website. (As a result, the FBI did not request

that Icelandic authorities proceed with imaging Server-1.) There was still some outbound Tor traffic flowing from Server-1, though, consistent with it being used as a Tor node; yet Server-1 was not included in the public list of Tor nodes, see supra n.4. Based on this fact, I believed, by the time of the June 12 Request, that the administrator of Silk Road was using Server-1 as a Tor “bridge” when connecting to the SR Server, as indicated in the June 12 Request. See Ex. A, at 1. (A Tor “bridge” is a private Tor node that can be used to access the Tor network, as opposed to using a public Tor node that could be detected on one’s Internet traffic. See Tor: Bridges, available at <http://torproject.org/docs/bridges>.) To be clear, however, the traffic data obtained for Server-1 did not reflect any connection to, or otherwise lead to the identification of, the Subject IP Address. The Subject IP Address was independently identified solely by the means described above – i.e., by examining the traffic data sent back from the Silk Road website when we interacted with its user login interface.

The two other details that point to June 6 may not actually exonerate Ulbricht. Silk Road’s live-ssl config file was altered on June 7, which is the earliest date for the site configuration provided in discovery (though page 23 has some additional dates).

The mtime for the live-ssl configuration file provided in Item 1 of discovery is June 7, 2013, and the phpmyadmin configuration is July 6, 2013.⁸

⁸ Since Item 1 is the oldest image provided in discovery the defense does not have site configuration data prior

to June 7, 2013.

And, as Horowitz reiterates, the earliest date for which the defense was provided discovery of a server imaging was June 6.

According to the government, the earliest image was captured June 6, 2013, and the latest in November 2013.

From a technical stand point, I'm not sure what to make of this.

A remarkable coincidence

It's clear, however, that FBI was tracking Silk Road well before June, and for some reason decided to make June the official start date (and, perhaps more significantly, official discovery start date; they've refused earlier discovery because it won't be used in trial) of their investigation. At the same time, it seems that Ulbricht's defense seems reluctant to explain why they're asking for earlier discovery; perhaps that's because they'd have to admit Ulbricht was aware of probes of the website before then. Forrest rejected their argument because Ulbricht refused to submit a declaration that this was his server.

But I am rather struck by the timing. As I said, the first Edward Snowden story – the June 5, 2013 Verizon release that could have no tie to the Silk Road investigation and, the next day, the WaPo and Guardian PRISM releases (there were very late Google and Facebook requests that seem like parallel construction, but since Ulbricht is a US citizen, his communications should not have been available via PRISM) – was roughly the day before the day Iceland imaged the other server.

I asked both Glenn Greenwald and Bart Gellman, and it seems the earliest the government could have had official notice of that story may have been late on June 4 though probably June 5 (things get funny with the Guardian, apparently,

because of Greenwich Mean Time). A more relevant leak to the Silk Road investigation was the President's Policy Directive on cyberwar – which Guardian published on June 7 (they may not have warned the government until that morning however).

So it may all be one big coincidence – that the government created a virgin birth for the Silk Road investigation that happened to be the same day that a torrent of leaks on the NSA and GCHQ started, ultimately revealing things like the government's targeting of the Tor network (just days after Ulbricht was arrested on October 2, 2013).

But it certainly seems possible that those investigating Silk Road felt the need to begin to roll up the investigation as that torrent of leaks started, perhaps worrying that the methods they (or GCHQ) were using might be exposed before they had collected the evidence.

Update: A few more points about this. My suspicion is that, if there is a tie between the Snowden leaks and the Silk Road investigation, it stems from the government's recognition that some of the methods it used to find Ulbricht would become known through Snowden's leaks, so it moved to establish an alternate means of discovery before Ulbricht might learn of those actual methods. As one example, recall that subsequent to Snowden's leaks about XKeyscore, Jacob Appelbaum got information showing XKeyscore tracks those who use Tor. While there are a number of things it seems Ulbricht's lawyers believe were parallel constructed (unnamed "law enforcement officers" got warrants for his Gmail and Facebook accounts in September), they most aggressively fought the use of a Title III Pen Register to track IP addresses personally associated with Ulbricht, also in September. It seems that would have been available via other means, especially XKeyscore, especially since by encrypting communication Ulbricht's communications could be retained indefinitely under NSA's minimization

procedures.

Additionally, the language the government used to refuse information on a range of law enforcement and spying agencies sure sounds like they clean teamed this investigation.

The Government also objects to the unbounded definition of the term “government” set forth in the September 17 Requests. Specifically, the requests ask the prosecution to search for information within “not only the United States Attorney’s Office for the Southern District of New York, but also the Offices in all other Districts, any and all government entities and law enforcement agencies, including but not limited to the Federal Bureau of Investigation, Central Intelligence Agency, Drug Enforcement Administration, Immigration and Customs Enforcement Homeland Security Investigations, National Security Agency, and any foreign government and/or intelligence agencies, particularly those with which the U.S. has a cooperative intelligence gathering relationship, i.e., Government Communications Headquarters (“GCHQ”), the British counterpart to the NSA.”

Even in the Brady context, the law is clear that a prosecutor has a duty to learn only of “evidence known to . . . others acting on the government’s behalf in the case.”

The government is not denying they had other means to identify Ulbricht (nor is it denying that it worked with partners like GCHQ on this). Rather, it is just claiming that the FBI officers involved in this prosecution didn’t see those methods.