# THE YAHOO SCAN: ON FACILITIES AND FISA

There are now two competing explanations for what Yahoo was asked by the government to do last year.

## Individual FISA order or 702 directive?

NYT (including Charlie Savage, who FOIAed all the FISC opinions and then wrote a book about them) explains Yahoo got an individual FISA order to search for a "signature" that the FBI had convinced the FISA Court was associated with a state-sponsored terrorist group.

> A system intended to scan emails for child pornography and spam helped Yahoo satisfy a secret court order requiring it to search for messages containing a computer "signature" tied to the communications of a state-sponsored terrorist organization, several people familiar with the matter said on Wednesday.
>
> Two government officials who spoke on the condition of anonymity said the Justice Department obtained an individualized order from a judge of the Foreign Intelligence Surveillance Court last year. Yahoo was barred from disclosing the matter.
>
> To comply, Yahoo customized an existing scanning system for all incoming email traffic, which also looks for malware, according to one of the officials and to a third person familiar with Yahoo's response, who also spoke on the condition of anonymity.
>
> With some modifications, the system stored and made available to the Federal Bureau of Investigation a copy of any

> messages it found that contained the digital signature.

Reuters — in a story emphasizing the upcoming debate about reauthorization — says that the order was a Section 702 order.

> The collection in question was specifically authorized by a warrant issued by the secret Foreign Intelligence Surveillance Court, said the two government sources, who requested anonymity to speak freely.
>
> Yahoo's request came under the Foreign Intelligence Surveillance Act, the sources said. The two sources said the request was issued under a provision of the law known as Section 702, which will expire on Dec. 31, 2017, unless lawmakers act to renew it.
>
> The FISA Court warrant related specifically to Yahoo, but it is possible similar such orders have been issued to other telecom and internet companies, the sources said.

Yet it also reports that both Intelligence Committees are investigating more about this request (which tells you something about Reuters' potential sources and how much the spooks' overseers actually know about this).

> The intelligence committees of both houses of Congress, which are given oversight of U.S. spy agencies, are now investigating the exact nature of the Yahoo order, sources said.

For what it's worth, at least until 2012, I think NSA and FBI might have been able to request this scan under 702; there are a bunch of court decisions, including one associated with what got reported as an upstream violation in 2012, that we haven't seen on this point

though. But particularly given Reuters' discussion of a "warrant" — which is more often used with traditional FISA — I suspect NYT is correct on this.

## "Hard" and "soft," and "upstream," "about," and "PRISM" are confusing the debate

The source of the confusion seems to stem from two separate sets of vocabulary that are unhelpful in understanding how FISA works.

The first set has to do with "hard" and "soft" selectors, language used in XKeyscore, which basically conducts boolean searches of buffered Internet traffic. Hard selectors are name, email, or phone identifiers associated with a specific person. Soft selectors are characteristics that can range from geographic location to specific code — so a search might ask for users of the encryption tool Mujahadeen Secrets in Syria, for example, which will return a bunch of people whose identities may not be known but whose activities warrant interest. Soft selectors can include searches on what counts as "content," but they also search on what counts as metadata.

I think the hard/soft distinction is misleading because — as far as I know — FISA has always operated on single selectors, not boolean searches. NSA isn't asking providers — whether they're phone companies or Internet providers — to go find people who are in interesting places and use interesting crypto (though AT&T may be an exception to this rule). Rather, they're asking for communications obtained by searching on specific selectors.

To be sure, for each target, there will be a range of selectors, often a huge number of them. Even for one person, as I have noted, NSA and FBI probably know of at least a hundred selectors. One Google subpoena response I

examined, for examined, included 15 “hard” identifiers for just one person (and multiply that by any major Internet service a person used). For a targeted organization like “Russian GRU hackers,” the NSA will probably have still more. But — again, as far as we know — FISA providers are asked to return data based off known selectors. But as I’ll show below, they’ve been asked to return data off selectors that would count as both hard and soft under XKeyscore.

The other set of confusing vocabulary comes from public debates about FISA (including PCLOB’s report on Section 702). Some debates have made a distinction between “upstream” and “PRISM.” Upstream is when NSA gives the telecoms a selector to collect information from scans conducted at switches, but it fundamentally refers to *how* something is collected, not who does it (and it’s possible there are backbone providers we haven’t thought of who also participate). PRISM is when NSA/FBI give Internet providers selectors to return activity on; it’s a description of *from whom* the information is collected. But even there, a PRISM provider will provide far more than just the email associated with a given selector.

Sometimes “upstream” collection is referred to as “about” collection. That’s misleading. “About” collection — that is, communications that contain a selector in what counts as content areas of the communication — is a subset of upstream collection. But what is really happening is that when the telecoms sniff packets to find a given selector, they need to sniff both the header and content to get all the communications they’re after, which is what PCLOB is saying here.

> With regard to the NSA’s acquisition of “about” communications, the Board concludes that the practice is largely an inevitable byproduct of the government’s efforts to comprehensively acquire communications that are sent to

> or from its targets. Because of the manner in which the NSA conducts upstream collection, and the limits of its current technology, the NSA cannot completely eliminate "about" communications from its collection without also eliminating a significant portion of the "to/from" communications that it seeks. The Board includes a recommendation to better assess "about" collection and a recommendation to ensure that upstream collection as a whole does not unnecessarily collect domestic communications.

One hazard of using "about" to refer to "upstream" collection is it leads people to forget that the NSA needs to use upstream collection to comprehensively collect non-PRISM Internet traffic, even when working just from "hard" selectors like email addresses. Some of this collection (as the PCLOB passage above makes clear) is just looking for any emails involving a target, not emails talking "about" that target. But at least according to PCLOB, because of the way this collection is done, even if NSA is only searching for a hard selector email, it will get "about" traffic.

As you can see, however, this language is already going to be insufficient to discuss the Yahoo request, which is effectively an "upstream" search on a PRISM providers' content (though I'm not clear whether it happens at the packet level or not). We also don't yet know whether the signature involved counts as content, but the filters Yahoo adapted for the process clearly scan the content.

## Public discussions have hidden how 702 includes non-email selectors

But the bigger problem with this discussion is that people are confused about what FISA permits the government to search on.

One huge shortcoming of the PCLOB report — one I pointed out at the time — is that it pretended that Section 702 was not used for cybersecurity. That's unfortunate because cybersecurity is the area where Section 702 most obviously includes non-email selectors, what would be called "soft" selectors in XKeyscore. When I first confirmed that NSA *was* using 702 for cybersecurity back when I briefly worked at the Intercept, it was based off the search on a cyber "signature," not an email. The target was a (state-sanctioned) hacker, but the search was not for the hacker's email, but for his tools.

Here's how PCLOB briefly alluded to this activity.

> Although we cannot discuss the details in an unclassified public report, the moniker "about" collection describes a number of distinct scenarios, which the government has in the past characterized as different "categories" of "about" collection. These categories are not predetermined limits that confine what the government acquires; rather, they are merely ways of describing the different forms of communications that are neither to nor from a tasked selector but nevertheless are collected because they contain the selector somewhere within them.

The Semiannual reports are one place where the government has officially admitted that it searches on more than just email addresses.

> Section 702 authorizes the targeting of non-United States persons reasonably believed to be located outside the United States. This targeting is effectuated by tasking communication facilities (also referred to as "selectors"), *including but not limited to telephone numbers and electronic communications accounts*, to Section 702 electronic communication service

> providers. [my emphasis]

As I said, the Snowden documents confirm that NSA has searched on malware signatures. Given the obvious application and the non-denials I have gotten from various quarters, I would bet a great deal of money that NSA has also searched on some signature associated with AQAP's Inspire magazine, effectively allowing it to track anyone who downloads (or decrypts) the magazine.

In a series of tweets yesterday, Snowden confirmed that the scope is even more broad.

> In practical terms, this means anything you can convince FISC to stamp. At NSA, I saw live examples of the following:
>
> The usual suspects (emails, IPs, usernames, etc), but also cryptographic hashes that identify known files (MD5/SHA1), sub-strings from base-64 encoded email attachments (derived from things like embedded corporate logos), and any uncommon artifacts arising from a target's tooling, for example if their app transmits a UUID (like a registration code or serial).
>
> The possibilities here are basically limitless, and we can't infer the specific nature of the string without more info.

The point is, "upstream" collection — whether done at a telecom switch or a tech server — can (and will, so long as FISC will authorize it) search on any string that will return the communications of interest, with "communications" extending to include "cyberattacks conducted by disembodied code."

To understand FISA collection, then, it is best to think in terms of selectors or facilities that will return a desired target. Here's some language from an Semiannual report that explains the distinction between target and facility (and

why the classified numbers in the report are undoubtedly much larger than the unclassified 92,000 "target" number we're given to explain the scope of FISA collection).

> The provided number of facilities on average subject to acquisition during the reporting period remains classified and is different from the unclassified estimated number of targets affected by Section 702 released on June 26, 2014, by ODNI in its 2013 Transparency Report: Statistical Transparency Report Regarding Use of National Security Authorities (hereafter the 2013 Transparency Report). The classified number provided in the table above estimates the number of facilities subject to Section 702 acquisition, whereas the unclassified number provided in the 2013 Transparency Report estimates the number of targets affected by Section 702 (89,138). As noted in the 2013 Transparency Report, the "number of 702 'targets' reflects an estimate of the number of known users of particular facilities (sometimes referred to as selectors) subject to intelligence collection under those Certifications." Furthermore, the classified number of facilities in the table above accounts for the number of facilities subject to Section 702 acquisition during the current six month reporting period (e.g., June 1, 2013 – November 30, 2013), whereas the 2013 Transparency Report estimates the number of targets affected by Section 702 during the calendar year 2013.

As explained above, for any given target, there may be a slew of selectors or facilities that NSA can collect on (though they probably only collect on a limited selection of all the selectors they know; they use the other selectors to make sure they can find all the

online activity of someone). The government tracks this internally by counting how many average selectors or facilities are targeted in a given day. These numbers will get more interesting, by the way, once the numbers incorporate USA Freedom Act compliance, which (in my opinion) significantly serves to require providers to provide all known selectors, that is, to even further expand the universe of known selectors.

## A history of the word "facility"

But to understand the background to the Yahoo thing, it is absolutely necessary to understand how the word "facility" has evolved within FISC (and we only have access to some of this). As far as we know, the meaning of the word started to change in 2004 when Coleen Kollar-Kotelly approved the installation of "Pen Registers" (really, packet sniffers) at switches to accomplish with the Internet dragnet what Stellar Wind had been doing (that is, the collection of Internet metadata in bulk), based on the logic that al Qaeda was using those facilities to communicate. Her ruling changed the definition of facility from meaning an individual user (a phone number or email address) to many users including the target. When Kollar-Kotelly first approved it, she required the government to tell her which specific switches they were going to target — that is, which switches were likely to carry traffic from target countries like Yemen and Afghanistan. But when John Bates reauthorized the Internet dragnet in 2010, he let the government decide on a rolling basis which facilities it would collect metadata from.

Thus, starting in 2004 and expanded in 2010, "facility" — the things targeted under FISA — no longer were required to tie to an individual user or even a location exclusively used by targeted users.

When Kollar-Kotelly authorized the Internet

dragnet, she distinguished what she was approving, which did not require probable cause, from content surveillance, where probable cause was required. That is, she tried to imagine that the differing standards of surveillance would prevent her order from being expanded to the collection of content. But in 2007, when FISC was looking for a way to authorize Stellar Wind collection — which was the collection on accounts identified through metadata analysis — Roger Vinson, piggybacking Kollar-Kotelly's decision on top of the Roving Wiretap provision, did just that. That's where "upstream" content collection got approved. From this point forward, the probable cause tied to a wiretap target was freed from a known identity, and instead could be tied to probable cause that the facility itself was used by a target.

There are several steps between how we got from there to the Yahoo order that we don't have full visibility on (which is why PCLOB should have insisted on having that discussion publicly). There's nothing in the public record that shows John Bates knew NSA was searching on non-email or Internet messaging strings by the time he wrote his 2011 opinion deeming any collection of a communication with a given selector in it to be intentional collection. But he — or FISC institutionally — would have learned that fact within the next year, when NSA and FBI tried to obtain a cyber certificate. (That may be what the 2012 upstream violation pertained to; see this post and this post for some of what Congress may have learned in 2012.) Nor is there anything in the 2012 Congressional debate that shows Congress was told about that fact.

One thing is clear from NSA's internal cyber certificate discussions: by 2011, NSA was already relying on this broader sense of "facility" to refer to a signature of any kind that could be associated with a targeted user.

The point, however, is that sometime in the wake of the 2011 John Bates opinion on upstream, FISC must have learned more about how NSA was really

using the term. It's not clear how much of Congress has been told.

The leap from that — scanning on telephone switches for a given target's known "facility" — to the Yahoo scan is not that far. In his 2010 opinion reauthorizing the Internet dragnet, Bates watered down the distinction between content and metadata by stripping protection for content-as-metadata that is also used for routing purposes. There may be some legal language authorizing the progression from packets to actual emails (though there's nothing that is unredacted in any Bates opinion that leads me to believe he fully understood the distinction). In any case, FISCR has already been blowing up the distinction between content and metadata, so it's not clear that the Yahoo request was that far out of the norm for what FISC has approved.

Which is not to say that the Yahoo scan would withstand scrutiny in a real court unaware of the FISC precedents (including the ones we haven't yet seen). It's just to say we started down this path 12 years ago, and the concept of "facilities" has evolved such that a search for a non-email signature counts as acceptable to the FISC.

## If a facility is not a user, then how do you determine foreignness?

*[Update: I realize this discussion is, given the increasing certainty that the Yahoo scan was done under an individual FISA order, irrelevant for the Yahoo case, because FBI has been cleared to collect on signatures in the US. But the issue is still an important one when discussing "facilities" that have been divorced from a geographically located user.]*

There's one final thing we don't have visibility on.

When Kollar-Kotelly started down this path, she

focused on facilities that were foreign-facing. That is, there was a high likelihood messages transiting those switches were one-side foreign, and therefore targetable, certainly for a PRTT. But as I noted, that foreign-facing distinction got badly watered down in 2010. And Yahoo's entire universe of emails would not be particularly foreign focused (though a lot of foreigners use Yahoo).

The question is, if NSA or FBI is targeting a facility that is not tied to a given user, but is instead tied to an organization that is located overseas, how does the government determine foreignness on a signature? NSA's General Counsel would permit analysts to collect on but not target metadata of, say, bots in the US based on the assumption that the ultimate source of the bot was overseas. If the signature that FBI searches on derives from overseas — as in the case where Inspire magazine is produced overseas — does that by itself deem a communication involving that signature to be "located" overseas, and therefore targetable.

I suspect that may be why NYT's sources emphasized that the target of the Yahoo search was a state-sponsored terrorist organization, rather than just a terrorist organization, because *by definition* that state would be overseas. But I also suspect that a lot of the recent troubles at NSA pertaining to "roving" selectors stems from the ambiguity that arises when you start targeting selectors that are not by definition geographically bounded.

The way the government targets facilities is constitutionally problematic in any case. But this question of foreignness seems to present both statutory and constitutional problems.