Searching Securely: Technical Issues with Warrants for Remote Search

[Extended Abstract]

Steven M. Bellovin
Columbia University
smb@cs.columbia.edu

Matt Blaze
University of Pennsylvania
mab@crypto.com

Susan Landau
Worcester Polytechnic Institute
susan.landau@privacyink.org

1. INTRODUCTION

In computer science we first develop theory (e.g., algorithms) which we aim to put into practice (e.g., through protocols). In law, there is an analogous process: first there is legislation, which is followed by rules that govern the implementation of the legislation. Of course, neither technology nor law are static. As we improve algorithms, or discover problems in theory or implementation, we adjust. The change from SSL to TLS is one example, and there are myriad others. The law also has processes for such accommodations. As technology and the world changes, so, too, do the rules for implementing legislation.

In the U.S., the Committee on Rules of Practice and Procedure of the Judicial Conference of the United States meets and examines how various legal procedural rules should be changed. Proposals are made, and after much discussion, amendments are proposed to current rules. There is then a period of public comment and testimony, after which the judicial advisory committee decides whether to submit proposed changes to the Standing Committee on Rules of Practice and Procedure, then to the Supreme Court. Finally, Congress may disapprove changes. All of this can be a several year process.

This paper analyzes, from a technical perspective, a set of proposed changes to *Rule 41* of the Federal Rules of Criminal Procedure, which governs the processes for authorizing searches and seizures. The proposed new rules were moved forward in 2014 by the Committee, and are currently moving forward through the process. We are concerned here with rules relating to remote computer searches under two conditions: when "anonymizing software" hides the location of a computer has been used and when the investigation involves botnets. Under today's rules, a magistrate judge can issue a warrant to search a computer located within his or her district. The proposed changes would grant judges the authority to issue a single warrant to cover remote searches of computers in other districts if the location of the computer has been concealed or if the computers to be searched are located in five or more jurisdictions [9, pp. 326-327].

Specifically, the proposal states:

> [A] magistrate judge with authority in any district where activities related to a crime may have occurred has authority to issue a warrant to use remote access to search electronic storage media and to seize or copy electronically stored information located within or outside that district if: (A) the district where the media or information is located has been concealed through technological means; or of (B) in an investigation of a violation of 18 U.S.C.§1030(a)(5),[1] the media are protected computers that have been damaged without authorization and are located in five or more districts.[9, pp. 338-339]

Part (A) is intended to apply to situations where a criminal or spy seeks to hide their activities through disguising the location of their device. Part (B) says when protected computers in at least five districts have been damaged, a magistrate judge can authorize remote searches via a single warrant. According to the committee, this aspect of the proposed changes, which is intended to apply to botnet investigations, is meant to be used in a "limited class of investigations" [9, p. 338]—but there is no explicit limitation in the proposal that would make it so.

"Protected computer" is a legal term of art, defined in 18 U.S.C. §1030(e)(2): a computer intrusion is a Federal offense if and only if the computer is "protected". The definition given is quite broad, though; it includes government computers, financial institution computers, and any other computer "which is used in or affecting interstate or foreign commerce or communication, including a computer located

[1] That notation means "Section 1030(a)(5) of Title 18 of the U.S. Code". The U.S. Code is the orderly compilation of Federal statutes. Bills passed by Congress are, in CS terminology, "diff" files to it. 18 U.S.C. §1030 is more commonly known as the *Computer Fraud and Abuse Act*, the Federal anti-hacking statute; the original version was passed in 1984. You can find it at `https://www.law.cornell.edu/uscode/text/18/1030`.

outside the United States that is used in a manner that affects interstate or foreign commerce or communication of the United States". In other words, if a machine is connected to the Internet, it is "protected". Thus, while "protected computer" might appear to refer only to a limited, special set of machines, in fact, it is quite a broad category. This means that the proposed changes to Rule 41 would have very wide implications, potentially applying to a large number of cases in which computer evidence is at play.

It is worth stepping back briefly to put the discussion of search in context. In the US, this begins, of course, with the Fourth Amendment to the Constitution:

> The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no warrants shall issue, but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

Applying the Fourth Amendment in specific situations can be complex, and there is a large body of often counterintuitive and perhaps occasionally contradictory law here. For example, a search warrant is generally required for searching a private home. But searching a car (say, incident to a traffic stop) often does *not* require a warrant. Searching an arrestee for knives and such does not require a warrant, but searching their seized phone does [33]. The rules for when warrants are required, what they cover, etc., are embodied in the *Federal Rules of Appellate, Bankruptcy, Civil, and Criminal Procedure,* with specific circumstances often requiring interpretation by the Supreme Court.

We are concerned about the interactions of the proposed changes regarding remote computer search with the realities of the technology itself. We are specifically concerned about issues of jurisdiction, chain of custody and authenticity of evidence, specificity of search, and notice. Note that we do not take a position on the legality or wisdom of law enforcement remote computer searches here. Rather, we focus on how their proposed implementation will have serious consequences if not resolved.

In this paper, we discuss those issues. We aim not to break new technical ground here, but instead examine how technical details, *many of which are well known to the security community,* play out in a legal framework.

2. HOW ELECTRONIC SEARCHES ARE CONDUCTED

Although the FBI has been comducting remote computer searches for over a decade, the bureau has said very little about how remote searches are performed today. The stated reason for their reticence is fear of alerting suspects to their techniques.

The first public mention such tools goes back to 2001 [37]. No one seemed to pay much attention until 2007, when there was a court filing on the FBI's use of the "Computer IP Address Verifier" (CIPAV) package, software that collects IP and MAC addresses, open ports, running programs, default browser and version, default OS and version, current logged-in user name [26, 31, 1]. CIPAV was used to track down a student threatening to bomb a high school in Lacey Washington.

CIPAV has been used in multiple cases across the country. The tool is complicated; first it must be downloaded onto a target machine, search the machine for certain information (such as that described above), report it, and then download spyware onto the machine to capture particular data. An FBI memo notes "a good deal of uncertainty under what authority is required to deploy an IPAV (sic)." [26] After consulting with their Office of General Counsel and the National Security Law Branch, the FBI opted for a two-step legal process: a search warrant for the computer intrusion, and a so-called "Pen Register/Trap-and-Trace" order for the subsequent monitoring.

There are other techniques in use as well. In one well-publicized (but never officially acknowledged) case, the FBI apparently hacked into an Irish child pornography server in Ireland and patched it to serve malware to visitors running a particular version of Firefox over Tor [18]. The malware did nothing except to send an alert with the real IP address of the machine to a server located in Virginia; this, of course, is an important step in finding the users of this site.

A search warrant almost certainly legally suffices for the penetration, and is quite likely necessary as well. (In one recent case, the Supreme Court held that a warrant, supported by probable cause, is needed even to affix an external GPS tracker to a car [39]. The case was decided partially on trespass and partially on violation of reasonable expectation of privacy. Both rationales would apply to the case of a remote computer search.) The second order, the pen register/trap-and-trace order, authorizes the the FBI to collect ongoing information on the endpoints of new communications during this period. Unlike search warrants, these orders do not require probable cause; they merely require "a certification by the applicant that the information likely to be obtained is relevant to an ongoing criminal investigation." Of course, if private information or communication content is to be obtained during a remote search, then a warrant or Title III "super warrant" would be needed.

There is one important exception to the warrant requirement: the protections of the Fourth Amendment stop at the border. If the FBI wishes to plant a CIPAV on a foreign computer, there is no obstacle in U.S. law to them simply doing so, as long as no "U.S. persons" are involved. In one case a U.S. magistrate noted that under certain circumstances he was not even authorized to issue a search warrant to hack into a computer outside the U.S. [35]; the relevant rules simply did not grant him that authority. Indeed, these changes to Rule 41 were proposed partially in response to his opinion. Of course, hacking into a foreign computer may violate the law of that country.

Little is known about just how CIPAV or other tools are deployed. Most likely, they use ordinary hacking techniques. Poulsen notes that in the 2007 bomb threat case, the FBI used MySpace's internal messaging system to deliver the

malware [31, 1]. It isn't clear just how matters proceeded from there, but there were certainly numerous browser bugs that could have been exploited.

There is one further legal step that may be necessary. As we noted in an earlier work [4], it is often necessary to do some preliminary reconnaissance of a target machine in order to prepare an exploit. Internet Explorer has different vulnerabilities than does Firefox; Firefox on Windows demands different techniques than does Firefox on MacOS or Linux. Depending on the circumstances, there are many ways to do this reconnaissance; some of these may demand their own search or wiretap warrants.

3. SPECIFICS OF THE PROPOSED RULE 41 CHANGES

In this section we examine the various technical problems posed by the proposed Rule 41 changes. Some are particular problems in the case of botnets, others are problems for remote search generally.

3.1 Searches of Victim Computers

Botnets pose a complex challenge to law enforcement. Botnet size is one problem; and the fact that the machines that have been taken over are *victims'* devices is another. Because of the multiplicity of victims, law enforcement sought to simplify the situation by using a single warrant approach to search multiple machines participating in a botnet. There are serious problems with this approach, which we discuss below.

The proposal suggests using a "common scheme to infect the victim computers with [law enforcement] malware" [9, p. 325]. This is a dangerous approach from both a technical and legal standpoint.

From a technical standpoint, the danger is that such a common scheme may easily go out of control. Current malicious botnet technology is generally relatively simple: the malware is typically virtually the same on all victims' machines, and thus it is easy to know where to find it and how to disable it. *There is no technical reason why, in the future, botnet malware may not be far more sophisticated, but the proposed rule is focused on the current state of criminal practice, and not on how technology might change practice.* In particular, botnet malware could be configured in a multiple of different ways that would not necessarily be easily predictable. What this means is that the "common scheme to infect the victim computers with malware" may fail, and not simply fail by not working. Such a scheme could easily fail by damaging the victims' computers in unpredictable and unexpected ways. As we know from such examples as Stuxnet, malware downloaded on victims' machines must be carefully tailored to the device [14, 44]. This is both to prevent the malware from damaging other parts of the victims computer (important for the uses being prescribed in the change to Rule 41) and also to prevent the law enforcement malware from causing damage should it escape the victim's computer.

From a legal standpoint, the lack of specificity is highly problematic. A technically sophisticated criminal could hide data in victims' machines in different places on different machines. (Specificity of search can be very important. In one case [40], a judge suppressed evidence of child pornography because the investigating agent specifically looked for child pornography on the suspect's computer, even though the warrant did not authorize him to do so.) If furthermore, the botnet information were to be encrypted—*and thus not visible in plain sight*—the resulting search would be essentially indistinguishable from a general warrant, since it would require searching the entire computer for a very few files.

In combination, these two sets of reasons make the multiple-victims-one-search-warrant approach exceptionally dangerous.

3.2 Location and Jurisdiction

Remote search creates complexities and potentially serious problems for location and jurisdiction. The Fourth Amendment specifies that warrants "particularly describ[e] the place to be searched." Apart from the legal issue of determining from which judicial district a valid warrant may be issued, finding the location of an arbitrary computer is not an easy task. This is true even if its IP address is known.

"IP geolocation" attempts to map an IP address to a location. This can be done using the *whois* database and DNS records (limited in value since many sites use a third party as host), using Internet topology, or even manual information, such as examining the language used in a webpage [17]. But while such techniques are often "good enough" for some purposes, IP geolocation is often legally problematic. In some cases, IP geolocation can be incorrect relatively frequently, and sometimes undetectably so. For example, because cellular carriers use carrier-grade NAT, IP geolocation information is least accurate with smartphones, and IP location—not cell tower—is picked up only coarsely. Indeed, one of us has seen a situation where a phone located in Singapore was identified as being in Kuwait; apparently, the geolocation mechanism being used relied on the registration address of the cellular company.

The practical inability of IP geolocation to be reliably accurate can create serious legal ambiguities. In a 2013 District Court case in the Southern District, Texas, *In re Warrant* [20], Judge Stephen Smith ruled that because "the current location of the Target Computer is unknown, it necessarily follows that the current location of the information on the Target Computer is also unknown. This means that the Government's application cannot satisfy the territorial limits of Rule 41(b)(1)."[20, p. 2].

There are other reasons besides the difficulty of IP geolocation that make reliably ascertaining a target's location extremely difficult or impossible under some circumstances. Virtual Private Networks (VPNs) create one difficulty, while Tor ("The Onion Router"), which is specifically *designed* to obscure location, creates another. Needless to say, geolocating Tor endpoints is at best extremely difficult [13, 38].

Open standards and procedures for making location determination are essential. The proposed rule is problematic, though. Section (b)(6)(A) provides that any magistrate in a district affected may issue a warrant if "the district

where the media or information is located has been concealed through technological means." Would this mean that a warrant can be issued against a VPN user simply on the basis of their address-hiding efforts? Surely that was not the intent of the rule makers, yet this is how the proposed rule has been written.

What should happen to the fruits of a search in event of erroneous location determination is a purely legal issue that we are not qualified to opine on; we nevertheless note that such outcomes are not at all improbable, even when no concealment has been attempted. We also note the "forum-shopping" issues raised by Orin Kerr regarding the transformation of physical searches into remote ones [21].

In a minor vein, we note that the current text of Rule 41 requires that warrants generally be executed during "daytime" in the subject's local timezone [15, Rule 41(e)(2)(A)(ii)]. This is not meaningful in a situation where location is difficult to determine; obviously, if a location is incorrect, the timezone may be incorrect as well. Presumably, this would be dealt with by an explicit exemption in the warrant itself, as is permitted by the current rules.

The fact that a target machine may be abroad makes this even more critical. While U.S. law may permit such searches, the law of the host country almost certainly does not. Thus what is needed is coordination with other signatories to a mutual legal assistance treaty (MLAT).

Given the current post-Snowden environment, it is unlikely that the searches being proposed under the changes in Rule 41 would be easily accepted. Before such changes are promlugated, law enforcement must be sure that American criteria for remote access are valid abroad. Some countries, in fact, prohibit such activity. Russia has charged an FBI agent with hacking for a remote search; the German courts have held that their constitution prohibits remote search entirely [5]. It is not clear that these issues have been properly considered in putting forth the proposed rule.

3.3 Danger and Intrusiveness

A remote search carries many risks, including the ones stemming from software errors. One fact that every working computer programmer or system administrator learns early on is that software often fails. This is especially true of patches or modifications to existing code. To give just one example, a recent release of iOS broke the ability of some iPhones to make calls [11]. The key word is "some": Apple presumably tested the iOS 8.0.1 update before shipping it, but on *some* machines it had serious side-effects.

Testing *cannot* be comprehensive; there will *always* be some situation that will occur on deployed code that was never tried in the test lab. Therein lies danger: all too often, an unsuspected failure can occur.

Remote search software is not immune. In fact, given some of its characteristics—it must run as a privileged ("root" or "administrator") program, in order to hide and to override file protections and examine hidden parts of the machine—it is more likely to cause unanticipated problems. Furthermore, errors in privileged programs can cause more damage; the same privileges that let them read protected files will also let them overwrite or delete files.

Two incidents widely attributed to intelligence agencies illustrate this point. In the "Athens Affair", someone subverted the lawful intercept mechanism on a mobile phone switch operated by Vodafone Greece [32]. Over a period of ten months, about a hundred phones were tapped, including the Prime Minister's. The penetration was detected because a programming error by the intruder caused a switch malfunction: text messages weren't being delivered properly. It is quite striking (and not at all surprising to the technical community) that the flaw affected a part of the switch not directly involved in the tap.

A second case is the Stuxnet attack on the Iranian nuclear centrifuge plant in Natanz [14, 44]. The direct impact on the centrifuges was not noticed; however, some of the PCs behaved so suspiciously that one was sent to a security firm in Belarus for examination. This company found the attack software.

We are certainly not asserting that remote search software will always fail, nor even that it will do so most of the time. However, if it is used on enough machines, e.g., when doing a large-scale search of bots, there almost certainly will be problems on some of them. This creates two serious problems. The first is the issue of the government causing further damage to victims' computers, a situation that is all too likely to occur on occasion. The second is that too much interference with their targeted computers' operation might render the search invalid. In one case [10], the 9th Circuit held that turning a car's telecommunications system into a remote bug violated the requirement in 18 U.S.C. §2518(4) for a "minimum of interference with the services." While this holding, pertaining to wiretap law, was based on statutory language, and was highly fact-specific, it does suggest that there is a threshold of interference beyond which law enforcement should not normally go. The rules for executing search warrants are also intended to minimize excess interference with the subject's normal life; consider the the normal restriction to daytime execution [15, Rule 41(e)(2)(A)(ii)]. Searches that have a significant chance of causing damage to victims' computers are an even larger problem.

3.4 Discussion of Techniques

With the exception of national-security investigations that do not result in evidence used in court, under U.S. law wiretap investigations must be disclosed to the target. For example, if a wiretap is conducted under Title III (federal law governing wiretaps for criminal investigations), the target must be informed of the search within thirty days after the conclusion of the wiretap. Yet the surreptitious searches being proposed create certain serious problems for the openness that lies at the heart of U.S. jurisprudence.

Surreptitious collection of evidence by compromising computers (and computerized devices such as mobile telephones) is an inherently technical endeavor, involving the use of methods that will vary widely depending on the particular hardware and software used by the target. Over time, these techniques will change to adapt to new target devices

and to circumvent new countermeasures. In practice, we would expect these tools to be constantly evolving, often quite rapidly.

It is natural to expect law enforcement and prosecutors to resist disclosing the specific tools and techniques they use to obtain access to their targets, citing the desirability of preserving sensitive "sources and methods" that might be used against other targets in the future. However, this goal must be balanced against a number of other risks, whose significance may not be immediately apparent to a non-technically trained judge.

First, it is imperative that any judge or magistrate authorizing a technical computer intrusion understand certain aspects of the specific technology that will be used to conduct the intrusion. This is necessary in order to meaningfully analyze the scope of the intrusion (what other information besides the evidence being sought will be exposed) and the risks that the technique to be employed might exceed the scope of the authorization. (We note that in [40], the judge wrote "the Court *now* understands that it is simple to make selections that allow the user to take advantage of the utility of the FTK program to categorize and sort out common known files such as program files, etc., without being required to flag the KFF alerts for child pornography files as part of the process" (emphasis added). In other words, the judge originally did not understand how FTK could be configured to include or exclude certain files.) This is particularly important when, as is often the case, the target's device is used for real-time communication (with content covered by the wiretap statutes) as well as for processing and storing information.

A defendant, similarly, will often require detailed technical information about how an intrusion was conducted in order to raise challenges as to whether a search improperly exceeded its authorization. Forensic examination of a possibly-hostile computer is difficult [23], and software bugs in the examination process can affect the results. We note that the Federal Rules of Evidence state that "But the expert may be required to disclose those facts or data on cross-examination" [16, §705]. Similarly, expert testimony must be "the product of reliable principles and methods" [16, §702(c)]. It is impossible to meet these conditions without disclosing the tools that extracted that data and making them available to the defense for examination.

The techniques used to obtain access to a computer can also have bearing on the authenticity, provenance, and context of the evidence collected. For example, it is possible that, depending on the technical details, a law enforcement intrusion could expose the target's computer (and any evidence collected from it) to tampering by others. Such claims can only be raised by the defense (or refuted) through analysis, possibly involving expert testimony, of the specific tools and techniques used. Other fields of forensic examination have been plagued by bad science [19, 28]; the best assurance of quality in the U.S. court system is the adversarial process.

The courts have not always agreed on the importance of the defendants' being able to view source code [36, 29]. We believe the courts were wrong, and because of the number of users involved in the remote search rule being proposed, it is imperative that as much information as possible about the technology used to conduct a remote search be disclosed to the judge authorizing the search as well as to the defense in any case in which such evidence is used. Declaring someone guilty "beyond a reasonable doubt", without examining the software that provided crucial evidence, is just wrong.

3.5 Chain of Custody and Authenticity of Evidence

Just as U.S. jurisprudence requires open processes, it requires that evidence be uncorrupted. It is much harder to maintain the integrity of evidence during a remote search than in a normal search done on a physically seized computer. Normal forensic procedures require that all analysis be done on a copy of a seized disk. Kerr describes the process well [22, pp. 540–541]:

> To ensure the evidentiary integrity of the original evidence, the computer forensics process always begins with the creation of a perfect "bitstream" copy or "image" of the original storage device saved as a "read only" file. All analysis is performed on the bitstream copy instead of the original. The actual search occurs on the government's computer, not the defendant's.
>
> A bitstream copy is different from the kind of copy users normally make when copying individual files from one computer to another. A normal copy duplicates only the identified file, but the bitstream copy duplicates every bit and byte on the target drive including all files, the slack space, Master File Table, and metadata in exactly the order they appear on the original. Whereas casual users make copies of files when their machines are running, analysts generally create bitstream copies using special software after the computer has been powered down. The bitstream copy can then be saved as a "read only" file so that analysis of the copy will not alter it.
>
> The accuracy of the bitstream copy often is confirmed using something called a "one way hash function," or, more simply, a "hash." A hash is a complicated mathematical operation, performed by a computer on a string of data, that can be used to determine whether two files are identical. If two nonidentical files are inputted into the hash program, the computer will output different results. If the two identical files are inputted, however, the hash function will generate identical output. Forensic analysts can use these principles to confirm that the original hard drive and the bitstream copies are identical.

There are a number of very important points in this excerpt. First, proper handling procedure for evidence requires that an "image copy" be made of the target disk. One reason for doing an analysis on a read-only image copy is that normal mechanisms for examining files change some of the metadata. Simply displaying a file will change that value. Second, and perhaps more important, working on a read-only

copy avoids any chance of making other inadvertent changes to the original.

Kerr notes that image copies also include the "slack space"—the free space—on the disk. This is very important for forensic analysis: when a file is deleted, its data is generally *not* overwritten; rather, the disk blocks are simply returned to the list of free storage. Indeed, as a Department of Justice manual notes, information can be concealed there deliberately: "Even if the agents know specific information about the files they seek, the data may be mislabeled, encrypted, stored in hidden directories, or embedded in 'slack space' that a simple file listing will ignore" [30, p. 76].

Finally, Kerr notes that the image file and the original device should be "hashed" to ensure that the two are identical. But hashes are not likely to help in practice here. A difference of a single bit will change the hash output. It is not possible to calculate a useful hash of a disk drive that is booted, even when the computer is idle; in most file systems, there are continual (and hard-to-notice) changes occurring in the the file system image through normal operating system activities.

All of this is important for evidentiary reasons. If a defendant challenges the authenticity of prosecution evidence, the case is much stronger if these procedures are followed. In a hearing in the "Silk Road" case, precisely such challenges were made.[2]

Current technology simply does not match our needs here, and this is not likely to change in the foreseeable future. Simply making an image copy from a machine can take hours under ideal conditions and with the cooperation of the machine's owner. Creating such an image copy of a non-trivial size disk is generally infeasible for surreptitious remote search; disks are too big and communications lines are too slow. Consider a two terabyte disk (a typical size on current-generation desktop computers) and a 25 Mbps Internet link (a typical fiber home or small business broadband connection). Even consuming the full bandwidth of the connection and assuming zero disk access or software bottlenecks, the theoretical minimum time to copy the entire drive is 640,000 seconds — more than a full week. And that is the absolute *best* case. Real throughput rarely exceeds half the link speed; furthermore, latency—the round trip time between the source and the destination, which is limited by the speed of light—is inversely proportional to the effective bandwidth [27, p. 68]. Copying a disk from San Francisco to Washington is inherently slower than a similar copy from New York, simply because of the distance. The issue of the difficulty of creating an image copy has been ignored in the discussion of the proposed rule changes, yet it is extremely important.

3.6 Specificity

Sometimes a difference in scale can be a difference in kind, and we believe it is in the case of searches of the victims of botnets. The proposed rule change is not about a single victim, or even a handful of victims, but potentially millions of such targets. Allowing broader seizures of information from millions of machines simply because they were the victims of computer crime seems wrong. Per our comments in Section 3.1, we suggest an explicit requirement that all remote search software be configured extremely narrowly when used on victim computers.

As noted, the meaning of "specificity" for electronic searches remains the subject of continuing constitutional debate [9, p. 341]. This issue becomes particularly serious when victim computers are the targets of remote search warrants. As the Preliminary Draft observed, botnets "may range in size from hundreds to millions of compromised computers" [9, p. 325]. While no one seriously calls into question whether or not a police officer, taking a crime report from a victim, should act if contraband is in plain sight, the meaning of "plain sight" in a computer search is by no means clear.

Because searching a victim's computer for botnet malware exposes the victim, a non-suspect, to an unwitting search, it is particularly crucial to limit the reasons under which such a search might be conducted. There would seem to be only three legitimate objectives for doing so: to demonstrate that a crime has indeed taken place (and even that is debatable, since arguably probable cause would be sufficient), to find pointers to the individual responsible for the botnet, and to ascertain the extent of the damage We can separate this into two cases: when the behavior of the botnet is understood, and when it is not.

When dealing with known botnets, law enforcement should be able to develop a clear understanding of exactly how the malware in question works. In particular, the computer security community has had great success studying botnets and locating their "command and control" nodes without hacking into other victim computers. The computer security community uses so-called "honeypot" systems—machines intended to be infected, and that engage in the same sort of risky behavior as unwitting machines do—that can be instrumented and monitored [25]. While law enforcement needs evidence to prove guilt beyond a reasonable doubt, the use of honeypots provides a less intrusive method of investigation, and law enforcement should use this type of approach first. Even if this does not suffice, the evidence will be in a very few, easy-to-locate places. It is thus feasible to construct search software that looks precisely and solely for the necessary indicia, rather than rummaging more broadly through the computer.

The alternative situation involves a more sophisticated sort of attack, where the necessary evidence may not be in a single, easy-to-examine place. A sophisticated attacker may, for example, split a contraband file into several pieces and stash them in different places, using, for example, Shamir secret-sharing [34]. While we haven't heard of criminals actually using such sophisticated techniques, it is certainly possible. That sort of scenario will likely require an examination that is less easily automated. But the complexity of the search *involving many locations on a victim's machine* would indicate that the victim should be necessarily be informed prior to downloading malware to track the attack. Given the sophistication of the attack, and the problems

[2]The case is 1:14-cr-00068-KBF, U.S. District Court for the Southern District of New York. The judge did not rule on the merits of the argument. *See* [24] for a description of the technical dispute.

that could conceivably ensue on the victim's machine, we suspect that most victims would be quite willing to cooperate at ridding their own systems of the infection—once law enforcement properly authenticated itself of course.

There is an alternative to searching the victims' machines for evidence: one could instead find such evidence at the ISP used by the victims. ISPs have been experimenting with sending notices to owners whose machines appear to be infected by a botnet; the ISP uses their knowledge of the machine's IP address to associate this with a billing address and thus can send an out-of-band mailing. An approach using Internet Service Providers (ISPs), discussed briefly in a paper by one of us [6], has the advantage that it also provides law enforcement with a better way to inform the victim of the problem. ISPs might also be used to detect infection, though this also raises privacy issues that deserve a thorough policy vetting.

We thus suggest that language mandating narrow searches, especially of victim machines, be added to the rule:

> An application for a warrant issued pursuant to (b)(6)(B) must include a statement specifying precisely which data is to be be seized. The warrant itself must limit the investigation to those specific facts.

To do otherwise would be to turn a phishing attack into a fishing expedition.

3.7 Notice

Search warrants generally require notice to the target, including a receipt for items seized [15, Rule 41(f)(1)(C)]. As noted in the proposal, this is problematic for remote search [9, p. 327]. We feel that the problem is even more difficult than indicated.

We can think of only four feasible mechanisms for notifying the target of a search: a file left on the computer; a pop-up window; an email message; or a physical letter. All are problematic, especially for mass searches.

A file left on a computer probably won't be noticed, but the most serious concern is that the user has no way to determine the authenticity or provenance of such a note. If such files were actually to become a legitimate form of communication, hackers would immediately start emailing files that looked just like the real ones, except with a URL to click on "to acknowledge the message". Naturally, these URLs would not be benign.

Email, of course, would have similar problems. The FBI itself has warned of malicious spam email purporting to be from them.[3] There are, at least in theory, technical solutions involving digitally signed messages and a Public Key Infrastructure. Experience with both Web browsers and phishing emails suggest that these do not work in the absence of careful training of users.

Hackers will abuse law enforcement-generated pop-up messages in similar ways. Indeed, they already have abused similar mechanisms, to serve ads [43]. Furthermore, there is little evidence that people would pay attention to such boxes; indeed, one online source jokingly defines a "dialog box" as "A window in which resides a button labeled 'OK' and a variety of text and other content that users ignore."[4]

Physical mail might suffice, but it will often be too time-consuming and expensive. While we do not have precise cost figures for criminal investigations, reports indicate that ISPs find such requests burdensome and charge accordingly.[5] Physical mail is also very difficult when dealing with unknown search targets. While a more extensive search of the target computer might yield a physical address, per the discussion in the prior section such a search would be extremely intrusive.

The language in the proposed rule—"reasonable efforts"—is probably correct; given these difficulties, we do not know how it can be done. We thus suggest that the Department of Justice develop and (after suitable public comment) promulgate binding regulations for how this should be accomplished.

The fact that all possible forms of notice are problematic is exactly the point. Consider the following "meatspace"[6] example.

> A radio announcement is made to all residents of San Jose, California (population approximately one million) saying that because of high amounts of ether, ammonia, and acetone—by-products of methamphetamine production—emanating from some apartments in the city, all homes the next day will be subject to a police inspection. Some of the emanations are, of course, the result of high methamphetamine production elsewhere (including next-door buildings or apartments); that is, some of the high sources are from victims, not

[3] *See* `http://www.fbi.gov/scams-safety/e-scams`:

> Ransomware Purporting to be from the FBI is Targeting OS X Mac Users
>
> 07/18/13—In May 2012, the Internet Crime Complaint Center posted an alert about the Citadel malware platform used to deliver ransomware known as Reveton. The ransomware directs victims to a drive-by download website, at which time it is installed on their computers. Ransomware is used to intimidate victims into paying a fine to "unlock" their computers. Paying the fine does nothing to solve the problem with the computer; do not follow the ransomware instructions. The ransomware has been called "FBI Ransomware" because it uses the FBI's name...

Several of us have received other spam messages purporting to be from the FBI.

[4] `http://www.w3.org/2006/WSC/wiki/Glossary`.

[5] See [2] for a news story about a civil case where plaintiffs were offered a limited number of subpoenas per month at the discounted price of $95 apiece. For a discussion of the technical difficulties ISPs face when fielding such requests see [7].

[6] "Meatspace" refers to the world outside the network.

producers.

Notice of such a broad search could only be given in a broadcast manner. This would be an inappropriate, extreme, and unworkable overreaction to the city's meth production problem. The fact that the inspection we are describing is an online one does not change this point.

This analogy, of course, is not especially precise. For many legal reasons, not the least of which is the lack of particularity, no such search warrant would be valid. That said, the notification issue is a close analog to what is being proposed here.

3.8 Remote Access and Security Mechanisms

While not directly addressed in the proposed rules, the proposal anticipates, at least implicitly, that surreptitious remote computer searches will become an increasingly prevalent law enforcement technique in the future. We agree that this is likely, and it is important that rules of evidence and criminal procedure address them. However, these methods also raise a number of policy issues that will need to be addressed by the courts and by lawmakers. We have previously raised some of these in our recent papers on the subject [4, 3].

Law enforcement reliance on remote computer intrusions exposes a conflict between solving some crimes by collecting evidence and preventing other crimes by better securing computers. Whether due to a software flaw or an explicit "backdoor," virtually any vulnerability that can be exploited by law enforcement for investigative purposes has the potential for illicit exploitation by criminals and foreign intelligence services. And the computer software, hardware, and devices used by criminals (and from which evidence is collected) are also used by thousands—or millions—of innocent citizens to store, process, and communicate the most important and sensitive details of their lives and businesses.

This means that that any flaw used by law enforcement for laudable evidence collection purposes also represents a risk to innocent people. As discussed above, it is natural to expect law enforcement to hold information about exploitable flaws closely, to maximize their useful lifetime for investigative use. But other public policy goals must be weighed against this. In addition to the rights of defendants to use information about these techniques to challenge evidence (discussed above), there is the broader question of reporting the vulnerabilities that law enforcement exploits to vendors so they can be fixed [4, 12]. That is, the use of vulnerabilities for law enforcement must be balanced against the need to protect citizens from criminals who might exploit them themselves.

4. WHAT HAS TRANSPIRED

In early November 2014, the Judicial Conference's Advisory Committee on the Federal Rules of Criminal Procedure held hearings on the proposed changes. A number of organizations, including the Electronic Frontier Foundation, the American Civil Liberties Union, the Center for Democracy and Technology, and Google submitted comments. Many were critical, and raised some of the same points we have raised here. The Justice Department even replied specifically to Google's objections [42, 8]. The objections notwithstanding, on March 16, 2015, the Advisory Committee approved the changes.

The full process of amending the rules is complex.[7] The next step is review by the Standing Committee on Rules of Practice and Procedure, probably at its June 2015 meeting; after that, the full Judicial Conference has to approve it, probably in September [41]. From there, it goes to the Supreme Court; Congress then has the right to block the changes. If there are no hold-ups—and some of the groups opposing the change intend to continue contesting it—the rule change would go into effect on December 1, 2016.

5. SECURELY CONDUCTING LAWFUL SURVEILLANCE

Law enforcement's role is largely in the solution of crimes, with prevention a laudable, but usually secondary, goal. This has proved to be particularly problematic in the area of cybersecurity.

When one is tracking down a particular crime or set of crimes, it is difficult to see beyond the immediate short-term goals. Yet whether it is seeking to regulate the braod use of cryptography, or the use of zero-day vulnerabilities in criminal investigations, short-term actions have long-term implications. While we recognize that the policy questions raised by the proposed Rule 41 changes may be beyond the scope of this particular proposal, we believe that it is imperative that they be addressed comprehensively including, and especially, the impact on cybersecurity. A piecemeal solution, such as is proposed here, is likely to leave society more vulnerable rather than less so. Thus any proposal to expand the use of vulnerability exploitation by law enforcement must be accompanied by a broader policy discussions of these inexorably related questions.

5.1 Recommendations

As is undoubtedly clear, we have a number of concerns with the current proposal, which does not appear to have undergone a thorough vetting from the technical side. Because we are not sure of the best way to proceed to satisfy law enforcement's needs, our recommendations are a response to the current proposal rather than a complete set of recommendations. Any proposal to change Rule 41 should satisfy the following recommendations, but there are likely to be other requirements, both technical and legal, that should be met as well.

- We recommend against the use of a single warrant to conduct multiple simultaneous searches on victims' computers. Blanket warrants cover far too many machines, without the necessary specificity; furthermore, they pose a great risk of damage to some of them.

[7] The Administrative Office of the U.S. Courts has a web page detailing the amendment process; see `http://www.uscourts.gov/rules-policies/about-rulemaking-process/how-rulemaking-process-works/overview-bench-bar-and-public`.

- We recommend that when a warrant is issued for searching a victim's computer, the warrant include precise, particularized specifications of the area of the computer that is to be searched.
- Remote search carries significant risk of causing international complications. Guidance to law enforcement, and perhaps the rule itself, should stress this. Except for extremely serious cases, such searches should be done only with the cooperation of the host country.
- As noted in the proposed rules, giving notice of a search is problematic. We suggest a two-pronged approach. First, there needs to be explicit guidance to law enforcement on what mechanisms should be used and under what circumstances; the conditions when notice can be omitted should also be described. Second, the Department of Justice should engage the technical community in an effort to devise better mechanisms.

We have noted elsewhere that targeted hacking, with a search warrant and under suitable conditions, is likely to become an increasingly prevalent investigative tool; see [3] and [4]. The former discusses technical aspects; the latter concentrates on the legal and policy issues. However, such searches must be carefully targeted and and their implementations tested, both to comply with legal requirements and mitigate some of the inherent technical risks. For example, despite being narrowly targeted and meticulously crafted, Stuxnet still managed to spread outside its apparent target; fortunately, because it was carefully designed, it does not appear to have actually caused serious damage outside of its target in Natanz.

Depositing law enforcement malware to investigate victims' machines is a very tricky business; it should never be attempted lightly. The proposal, which does not sufficiently attend to complex technical issues, must be substantially reworked to take this concern into account. Otherwise, law enforcement could be creating more damage than that which it is seeking to prevent, an approach that can neither be constitutional nor desired.

We have made recommendations on changes that should be made to the proposal, but we believe that more than simple changes are required. While in this note we have identified a number of specific technical flaws with the proposed changes to Rule 41, there may be others that we have missed. In addition, for the most part, we have not addressed the many legal complexities in this proposal. So we suggest—and we have argued this at greater length earlier [4]—that a legislative fix would be best. There is, to our knowledge, no explicit statutory authority for law enforcement to hack into computers; given the intrusiveness and danger of such activities, there is a need for balance. The legislative process is better suited to address this than the rulemaking process.

We note that while this paper has focused on a specific proposal that applies only to U.S. law, the issues are international. Matters of jurisdiction, proportionality, privacy, intrusiveness, preservation of evidence, and striking the balance between effective law enforcement and risk to the innocent are concerns in all democracies that operate under the rule of law. So while the debate is currently local, the issues, and the stakes, are global.

Let us close with an observation that has been made repeatedly in recent years. Communications technology has grown increasingly complex, and the ability of legislators and judges to understand the technological implications of law and policies has become increasingly difficult as a result. Several of the reports addressing NSA collection recommended that outside technology expertise be provided to the Foreign Intelligence Surveillance Court; such expertise should also be drawn upon as a matter of course in drafting and implementing the rules that govern law enforcement wiretapping as well.

References

[1] *Affidavit for State of Washington, County of King, In the Matter of the Search of any Computer Accessing Electronic Message(s) Directed to the Administrator(s) of MySpace Account "Timberlinebombinfo" and Opening Message(s) Delivered to that Account by the Government*. No. MJ07 - 5114. 2007. URL: http://politechbot.com/docs/fbi.cipav.sanders.affidavit.071607.pdf.

[2] Nate Anderson. "Big Cable fed up with endless P2P porn subpoenas". In: *Ars Technica* (Feb. 4, 2011). URL: http://arstechnica.com/tech-policy/2011/02/big-cable-getting-fed-up-with-endless-p2p-porn-subpoenas/.

[3] Steven M. Bellovin, Matt Blaze, Sandy Clark, and Susan Landau. "Going Bright: Wiretapping without Weakening Communications Infrastructure". In: *IEEE Security & Privacy* 11.1 (Jan.–Feb. 2013), pp. 62–72. ISSN: 1540-7993. DOI: 10.1109/MSP.2012.138. URL: https://www.cs.columbia.edu/~smb/papers/GoingBright.pdf.

[4] Steven M. Bellovin, Matt Blaze, Sandy Clark, and Susan Landau. "Lawful Hacking: Using Existing Vulnerabilities for Wiretapping on the Internet". In: *Northwestern Journal of Technology & Intellectual Property* 12.1 (2014). URL: http://scholarlycommons.law.northwestern.edu/njtip/vol12/iss1/1/.

[5] Susan W. Brenner. "Law, Dissonance, and Remote Computer Searches". In: *North Carolina Journal of Law and Technology* 14 (Fall 2012–2013), pp. 43–92.

[6] D.D. Clark and S. Landau. "The Problem isn't Attribution: It's Multi-Stage Attacks". In: *Third International Workshop on Re-Architecting the Internet*. 2010.

[7] Richard Clayton. "Anonymity and Traceability in Cyberspace". Also published as technical report UCAM-CL-TR-653. PhD thesis. University of Cambridge, Darwin College, 2005. URL: http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-653.html.

[8] *Comment from David Bitkower, U. S. Department of Justice*. Feb. 23, 2015. URL: http://www.regulations.gov/#!documentDetail;D=USC-RULES-CR-2014-0004-0055.

[9] Committee on Rules of Practice and Procedure of the Judicial Conference of the United States. *Preliminary Draft of Proposed Amendments to the Federal Rules of Appellate, Bankruptcy, Civil, and Criminal Procedure*. Aug. 2014. URL: http://www.uscourts.gov/uscourts/rules/preliminary-draft-proposed-amendments.pdf.

[10] *Company v. United States*. 349 F.3d 1132 (9th Cir. 2002).

[11] Andrew Cunningham. "iOS 8.0.1 disabling cellular and TouchID on some phones". In: *Ars Technica* (Sept. 24, 2014). URL: http://arstechnica.com/apple/2014/09/apple-releases-ios-8-0-1-with-healthkit-keyboard-iphone-6-fixes/.

[12] Michael Daniel. "Heartbleed: Understanding When We Disclose Cyber Vulnerabilities". In: *The White House Blog* (Apr. 26, 2014). URL: https://www.whitehouse.gov/blog/2014/04/28/heartbleedunderstanding-when-we-disclose-cyber-vulnerabilities.

[13] Roger Dingledine, Nick Mathewson, and Paul Syverson. "Tor: The Second-Generation Onion Router". In: *Proceedings of the 13th USENIX Security Symposium*. Aug. 2004.

[14] Nicolas Falliere, Liam O Murchu, and Eric Chien. *W32.Stuxnet Dossier*. Symantec Security Response. Version 1.4. Feb. 2011. URL: http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf.

[15] *Federal Rules of Criminal Procedure*. Dec. 1, 2013. URL: http://www.uscourts.gov/uscourts/rules/criminal-procedure.pdf.

[16] *Federal Rules of Evidence*. Dec. 1, 2014. URL: https://www.law.cornell.edu/rules/fre.

[17] Mark Gondree and Zachary N.J. Peterson. "Geolocation of Data in the Cloud". In: *CODASPY '13*. Feb. 2013. URL: http://znjp.com/papers/gondree-codaspy13.pdf.

[18] Dan Goodin. "Attackers wield Firefox exploit to uncloak anonymous Tor users". In: *Ars Technica* (Aug. 5, 2013). URL: http://arstechnica.com/security/2013/08/attackers-wield-firefox-exploit-to-uncloak-anonymous-tor-users/.

[19] Spencer S. Hsu. "FBI admits flaws in hair analysis over decades". In: *Washington Post* (Apr. 18, 2015). URL: http://www.washingtonpost.com/local/crime/fbi-overstated-forensic-hair-matches-in-nearly-all-criminal-trials-for-decades/2015/04/18/39c8d8c6-e515-11e4-b510-962fcfabc310_story.html.

[20] *In re Warrant to Search a Target Computer at Premises Unknown*. 958 F. Supp. 2d 753 (S.D. Tex. 2013).

[21] Orin Kerr. *Memo to Members of the Rule 41 Committee*. cited in *Advisory Committee on Criminal Rules*, pp 251–252, April 7–8, 2014. Feb. 8, 2014.

[22] Orin S. Kerr. "Searches and Seizures in a Digital World". In: *Harvard Law Review* 119.2 (Dec. 2005), pp. 531–585. URL: http://www.jstor.org/stable/4093493.

[23] Gary C. Kessler. "Anti-Forensics and the Digital Investigator". In: *Australian Digital Forensics Conference*. 2007. URL: http://ro.ecu.edu.au/cgi/viewcontent.cgi?article=1000&context=adf.

[24] Brian Krebs. "Silk Road Lawyers Poke Holes in FBI's Story". In: *Krebs on Security* (Oct. 14, 2014). URL: http://krebsonsecurity.com/2014/10/silk-road-lawyers-poke-holes-in-fbis-story/.

[25] Kirill Levchenko et al. "Click trajectories: End-to-end analysis of the spam value chain". In: *IEEE Symposium on Security and Privacy*. IEEE. 2011, pp. 431–446. URL: http://www.icir.org/christian/publications/2011-oakland-trajectory.pdf.

[26] Jennifer Lynch. "New FBI Documents Provide Details on Government's Surveillance Spyware". In: *EFF Deeplinks Blog* (Apr. 29, 2011). URL: https://www.eff.org/deeplinks/2011/04/new-fbi-documents-show-depth-government#footnote20_ul40kw8.

[27] Matthew Mathis, Jeffrey Semke, Jamshid Mahdavi, and Teunis Ott. "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm". In: *ACM SIGCOMM Computer Communication Review* 27.3 (1997), pp. 67–82. URL: http://dl.acm.org/citation.cfm?id=264023.

[28] Jane Campbell Moriarty and Michael J. Saks. "Forensic Science: Grand Goals, Tragic Flaws, and Judicial Gatekeeping". In: *Judges Journal* 44 (2005), pp. 16–33.

[29] *Nicholas Ray Swendra, petitioner, Appellant, vs. Commissioner of Public Safety, Responder*. Court of Appeals, State of Minnesota, A07-2434. Jan. 13, 2009. URL: http://mn.gov/lawlib/archive/ctapun/0901/opa072434-0113.pdf.

[30] Office of Legal Education. *Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations*. 2009. URL: http://www.justice.gov/criminal/cybercrime/docs/ssmanual2009.pdf.

[31] Kevin Poulsen. "FBI's Secret Spyware Tracks Down Teen Who Made Bomb Threats". In: *Wired* (July 18, 2007). URL: http://archive.wired.com/politics/law/news/2007/07/fbi_spyware.

[32] Vassilis Prevelakis and Diomidis Spinellis. "The Athens Affair". In: *IEEE Spectrum* 44.7 (July 2007), pp. 26–33. URL: http://spectrum.ieee.org/telecom/security/the-athens-affair/0.

[33] *Riley v. California*. 134 S. Ct. 2473. 2014.

[34] Adi Shamir. "How to Share a Secret". In: *Communications of the ACM* 22.11 (1979), pp. 612–613.

[35] Stephen Wm. Smith. *In re WARRANT TO SEARCH A TARGET COMPUTER AT PREMISES UNKNOWN*. 2013 WL 1729765. Apr. 22, 2013. URL: https://s3.amazonaws.com/s3.documentcloud.org/documents/692822/in-re-warrant-to-search-a-target-computer-at.pdf.

[36] *State of Minnesota, Appellant vs. Dale Lee Underdahl Respondent, Timothy Arlen Brunner Respondent*. Court of Appeals, State of Minnesota, A07-2293, A07-2428. May 20, 2008.

[37] Bob Sullivan. "FBI software cracks encryption wall". In: *MSNBC* (Nov. 20, 2001). URL: http://www.nbcnews.com/id/3341694/ns/technology_and_science-security/t/fbi-software-cracks-encryption-wall/.

[38] P F Syverson, D M Goldschlag, and M G Reed. "Anonymous Connections and Onion Routing". In: *IEEE Symposium on Security and Privacy*. Oakland, California, Apr. 1997, pp. 44–54. ISBN: 0-8186-7828-3. URL: citeseer.nj.nec.com/syverson97anonymous.html.

[39] *United States v. Jones*. 132 S. Ct. 945. 2012.

[40] *United States v. Schlingloff*. 901 F. Supp. 2d 1101. 2012.

[41] Dustin Volz. "FBI's Plan to Expand Hacking Power Advances Despite Privacy Fears". In: *Government Executive* (Mar. 17, 2015). URL: http://www.govexec.com/management/2015/03/fbis-plan-expand-hacking-power-advances-despite-privacy-fears/107712/.

[42] Dustin Volz. "Feds Dismiss Google's Fears Over FBI's Hacking Power". In: *National Journal* (Feb. 26, 2015).

[43] Washington State Office of the Attorney General. *Pop-Up Ads*. URL: http://www.atg.wa.gov/InternetSafety/PopUpAds.aspx.

[44] Kim Zetter. *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*. New York: Crown Publishers, 2014.