Conducting Cybersecurity Research Legally and Ethically

Aaron J. Burstein
University of California, Berkeley (School of Law)
aburstein@law.berkeley.edu

Abstract

The primary legal obstacles to conducting cybersecurity are not outright prohibitions but rather the difficulty of determining which of a large set of complex statutes might regulate a given research project. Privacy, computer abuse, tort, and contract law are all potentially applicable. Moreover, even when the law permits a research activity, researchers may wonder whether it is ethically permissible. This paper seeks to clarify these issues by explaining the areas of law that are most generally applicable to cybersecurity researchers and offering guidelines for evaluating ethical issues that arise in this area of research.

1 Introduction

Research occupies a central role in cybersecurity policy in the United States. It may provide ways to reduce and mitigate the increasingly serious threats to the computers and networks that the United States (and other highly developed countries) have come to rely upon so heavily. Funding this research has been a priority for Congress as well as the National Science Foundation, DARPA, the Department of Homeland Security, and other agencies [11]. As networked information systems become pervasive, this commitment to research is essential.

But a fog of legal and ethical uncertainty hangs over cybersecurity research. A variety of federal and state statutes either prohibit activities that would provide cybersecurity researchers with data about real systems and real attackers, or cast such doubt on research activities that researchers modify their programs or conduct them with a sense of uncertainty as to their legality. Cybersecurity researchers (and officials within the organizations that employ them) may also suspect that certain things are illegal when, in fact, they are not; but researchers nonetheless avoid certain paths. Conversely, researchers may view the legality of a certain course of research as license to pursue it without regard to ethical considerations.

Ethical questions lurk beyond these legal issues and also deserve researchers' attention. Though the statutes discussed here contain expansive prohibitions on certain kinds of conduct, they do not address all instances in which researchers may find themselves wondering, "Is this the *right* thing to do?" In addition, many cybersecurity researchers present their data collection and analysis plans to institutional review boards (IRBs) and information officers (e.g., CISOs) for approval. These individuals and bodies often are unfamiliar with cybersecurity research in general and the problems that research face collecting data in particular. They will often wonder about how proposed research affects individual privacy and the security of the organization's information systems. The better researchers can explain how their activities will affect these interests, the easier they may find it easier to obtain approval and cooperation.

The overall argument in this paper is twofold. First, though U.S. law does not permit everything that cybersecurity researchers would like to do, relatively few research activities are flatly prohibited.[1] Nonetheless, uncertainty among researchers about what the law actually says, as well as doubt about the ethics of some activities, may hold back certain research efforts. Though privacy is an important part of this picture, computer abuse, copyright, tort, and contract law pose issues as well. Second, this paper emphasizes that cybersecurity researchers work within organizations whose interests typically include far more than improving cybersecurity. Thus, this paper strives to provide ways to allow cybersecurity re-

[1] **Disclaimers:** First, this paper considers U.S. law only. Other nations' laws are part of a more complete picture of cybersecurity research legal issues, but, given the limited space available and the complexities of U.S. law, it is impossible to address international law in a helpful manner here. Second, though the author of this paper is an attorney, nothing in this paper constitutes legal advice. Researchers who believe they are encountering issues similar to those discussed here should discuss their individual circumstances with an attorney.

searchers to think through the legal and ethical dimensions of their research, so that they may better explain it to non-experts and discuss how it is consistent with an organization's overall interests. The discussions in this paper revolve around general problems that cybersecurity researchers face, rather than particular research efforts. The hope is that whatever is lost by avoiding discussion of specific research will be recovered by preventing embarrassment to researchers and encouraging a frank discussion within the cybersecurity research community.

Section 2 reviews previous work examining legal issues in cybersecurity research. Section 3 explains the legal and ethical issues surrounding collecting and sharing network datasets, ending with a proposal to create a cybersecurity research exception to federal communications privacy laws. Section 4 discusses issues associated with running malicious code on research machines. Section 5 analyzes the law and ethics of mitigating attacks, while Section 6 does the same for publishing results. Finally, Section 7 concludes with a few suggestions for action by cybersecurity researchers with respect to their own research, within their organizations, and within the political arena.

2 Background

A few legal scholars have examined some of the legal issues facing cybersecurity research. Liu, for example, has examined the effects of the Digital Millennium Copyright Act (DMCA) on cryptography research [13]. He concluded that the DMCA's prohibitions on circumventing "technical protection measures" on copyrighted works are so broad, and the encryption research exception is so narrow, that researchers are justified in fearing liability for researching and publishing about vulnerabilities in certain encryption schemes.

Research using honeypots and honeynets raises significant questions about liability under the federal Computer Fraud and Abuse Act (CFAA) and communications privacy statutes (including the Wiretap Act and Pen Register/Trap and Trace Devices Act). Salgado analyzed a range of honeynet set-ups and found that the risk of liability under the communications privacy statutes can best be reduced by incorporating honeynets into production systems and networks[20]. He did not, however, give much attention to researcher liability under the CFAA, the possibility of which must be taken into account given that more recent honeynet designs involve more interaction with attackers.

Finally, Ohm et al. examined statutory communications privacy (including the Stored Communications Act in addition to the statutes named above) issues arising in conjunction with collecting, publishing, and using network traces [17]. They argued that these statutes are sufficiently vague to make it unclear whether a given trace collection will violate one or more of them. Nonetheless, they argued, legislative reform of these laws is probably unnecessary and, in any event, would be unlikely to add much clarity for cybersecurity researchers.

3 Obtaining Data from Networks

Data from real networks is critical to several areas of cybersecurity research. Intrusion detection research, for example, depends on access to large volumes of network traffic in order to generate signatures of attacks while minimizing false positives and false negatives. The stresses of real systems may also be necessary to test the performance of real-time collection and analysis technologies. In addition to their importance to individual research efforts, datasets can contribute to a broad picture of the Internet when shared among researchers [6].

3.1 Collecting Network Traces

As many cybersecurity researchers are aware, however, federal communications privacy laws limit access to the traffic on computer networks.[2] In particular, federal law provides the following:

- Wiretap Act [1]: Prohibits real-time interception of the *contents* of electronic communications. A "provider exception," however, permits the employees of a network operator to intercept and record communications to the extent necessary to protect the "rights and property" of the operator.

  Unfortunately, the distinction between "content" and "non-content" information is not always clear. In particular, the distinction is not as simple as the separation between packet header and payload. The contents of a communication are defined to mean the "substance, purport, or meaning" of the communication, while non-content information refers to both addressing information as well as records pertaining to a network user, e.g., billing information. Under these definitions, courts have held IP addresses (both sender and receiver) and the To: and From: fields in e-mail messages to be non-content information [25], while the Subject field is commonly regarded as contents [15]. The same definitions of "contents" and "non-content information" apply to the two statutes discussed below.

- Pen Register/Trap and Trace statute [5] (commonly referred to as the "Pen/Trap statute"): Prohibits

[2] Many states have their own versions of these laws. In particular, many have their own version of the Wiretap Act, and in some states, the law is more strict with respect to consent. In California, for example, *both* parties to a communication must consent to its interception.

real-time interception of the non-content portions of electronic communications. The Pen/Trap statute contains a provider exception that is similar to the one provided under the Wiretap Act. Once non-content data are stored, analysis and disclosure of the data are subject to the Stored Communications Act.

- Stored Communications Act (SCA)[4]: Prohibits providers of "electronic communications service to the public" from knowingly disclosing the contents of customers' communications, as well as non-content records relating to customers' communications. The SCA imposes little, if any, restrictions on uses of data within the organization that collects them. Publishing or sharing the same data with employees of other organizations, however, implicates the more restrictive disclosure rules discussed in Section 3.2.

Taken as a whole, there are two salient features of this complex set of laws. First, they contain no research exceptions. This is in contrast to other privacy statutes, such as the Health Insurance Portability and Accountability Act (HIPAA), which restricts disclosures of personal health information but provides means for researchers to obtain such information both with and without individual consent. The provider exceptions to the Wiretap Act and Pen/Trap statute are the closest that these laws come to a research exception. Making use of this exception requires close cooperation between researchers and officials from their institutions.

The second point to note about the electronic communications privacy statutes is that they create a patchwork of prohibitions and exceptions that are difficult for researchers and research organizations to navigate. As the summaries above indicate, the rules for accessing communications contents are different from those governing access to addressing information; and access to data in real-time versus in storage introduces still more variations in the law.

Thus, the Wiretap Act and Pen/Trap statute pose obvious hurdles to cybersecurity researchers. Consider the issue of consent under the Wiretap Act. Given that testing, say, intrusion detection algorithms may require access to traffic at a university's gateway, obtaining individual consent is probably unworkable. Universities typically inform their network users, through banner notices or terms of use, that the network is monitored. It is unclear, however, whether these notices cover all situations of interest to researchers (e.g., large-scale packet trace collection). Even if a university obtains broad consent to monitors its network users, administrators are likely to give considerable weight to other institutional interests (e.g., student or faculty backlash) that may cut against increasing researchers' access to network data. An empirical study of institutions' policies and practices could shed light on this area.

Making use of the provider exception to the Wiretap Act or the Pen/Trap statute obviates the need for consent, but it requires coordination with the appropriate officials within the institution that operates the network. For large organizations, the key official is likely to be a chief information security officer (CISO) and his or her staff. Convincing a CISO that research that involves tapping into the contents of communications on the institution's network is likely to involve more than an assertion that an appropriately structured research project is *legal*. The CISO will also want to ensure that the fits the institution's mission and policies. It is here that attention to ethical considerations may be valuable.

The question that researchers and institutional officials must confront is: Even if it is legal to allow research that involves real-time monitoring and analysis of communications, why *should* the institution allow it? The broader background of communications privacy law and policy provides a few answers.

First, research that fits within the provider exception is, by definition, potentially applicable to protecting the institution's network. A close relationship between researchers and staff with responsibility for keeping a network operational may bring immediate benefits—improved security—to the network and its users.

A second answer is based on a more basic look at the interests that the Wiretap Act was intended to protect. Giving cybersecurity researchers access to real-time communications streams would do little to undermine these interests. When the Wiretap Act was first enacted in 1968, and even when it was expanded in 1986 to cover electronic communications, intercepting communications in real time was by far the easiest—and perhaps the only—way of obtaining their contents. The advent of essentially unlimited storage of email and other forms of electronic communications, however, has made it possible for law enforcement officials and private parties to obtain contents from *stored* communications. The individual informational privacy interest is in the contents of a communication, rather than the mode in which it was obtained.

In addition, the Wiretap Act was framed against the assumption that a person might have one of a few reasons for intercepting a communication without authorization, all of which merit some control under the law: gathering evidence for a criminal investigation, gathering material to embarrass another person, or simply satisfying a curiosity in the affairs of other people. Cybersecurity researchers do not (or should not) pursue these ends when they make use of real-time communications streams. Instead, for the most part, they subject the communications to automated analysis. To be sure, it may sometimes be necessary for researchers themselves to examine the contents of communications to debug software, improve experimental set-ups, or to explain anomalous or unexpected results. Researchers should be frank about this possibility when discussing proposed projects with institutional officials, and they specify which investigators would have access to individual communications and how they would keep the communications confidential.

3.2 Sharing and Publishing Network Traces

A second general problem that cybersecurity researchers face in the realm of communications privacy is that of sharing publishing network traces. The scientific bases for sharing these data are compelling: common datasets can provide meaningful comparisons between competing research approaches; simulated data are inadequate for some uses; and existing datasets may not reflect present-day threats or traffic characteristics [18].

The Stored Communications Act (SCA), introduced above, poses a significant barrier to sharing these data. Some additional detail about this law is warranted at this point.

- Entities Subject to the SCA. The relevant sections of the SCA do not cover all network providers, but rather providers of electronic communications services "to the public." Commercial e-mail providers and ISPs generally are thought to be covered by the SCA, while private businesses that provide Internet access to their employees for work purposes likely are not covered by the SCA. Universities may fall somewhere in the middle, or even have some networks governed by the SCA and some that are not. For example, if a university operates an open wireless network, records pertaining to that network might well be covered by the SCA. A research network that is available only to students, staff, and faculty, however, might not be a service "to the public"; and hence the SCA might not apply to content and records pertaining to that network. To reiterate, the question of whether an entity provides service to the public is critical; if it does not, the disclosure provisions of the SCA do not apply.

- Disclosures regulated by the SCA. A service provider subject to the SCA may not disclose content records to another person or entity without consent (or the appropriate court order).

  Moreover, a covered service provider may not disclose *non-content* records to any "governmental entity" without consent or the appropriate order. The meaning of "governmental entity" is quite broad; it might refer to *any* government agency and its employees [27], including public universities. The term is not limited to law enforcement or intelligence agencies and officials.

For those entities covered by the SCA, the prohibition against divulging non-content records to governmental entities makes an unrestricted public release of data a risky proposition. Putting a dataset on a public website, for example, would make it possible for anyone to obtain the data. Though a case could be made that this mode of disclosure does not meet the statutory standard of *knowingly* divulging non-content records to a governmental entity, researchers (and their institutions) are probably will not want to rely on this argument.

As discussed above, the SCA only applies to providers of communications services "to the public." Others may disclose non-content records. For these entities, the question becomes an ethical one that researchers and institutions must confront: *should* they publish network traces?[3]

The SCA's history and structure points toward some answers. The baseline of statutory protection for non-content records is quite low. The SCA primarily protects against government intrusions into the privacy of non-content records, as is evident from the prohibition on disclosure to governmental entities, which includes (among many other things) law enforcement agencies that have the power to use such information to surveille or prosecute individuals. Though the threat of government surveillance has not abated, private firms now rival, if not surpass, the government's power to analyze network data at the individual level; and the SCA leaves monitoring and analysis by the private sector essentially unregulated. This legal structure allows commercial datamining, behavioral targeting and other practices that are particularly offensive to some conceptions of individual informational privacy to go forward. It is against this background that sharing non-content network traces should be evaluated in privacy terms; carefully anonymized datasets reveal far less about individuals than organizations learn from the data that they control and use for commercial purposes. (Compare Allman and Paxson's description of anonymized packet traces and NetFlow records in [6] with Solove and Hoofnagle's description of commercial datamining in [22] and Solove's description of government datamining in [21]. Yet public and private investment are heavily tilted toward supporting these invasive forms of analysis.

A more general solution to the barriers to research posed by electronic communications privacy laws would

[3] For the purposes of this discussion, it is assumed that only non-content (i.e., packet header) traces are in question, and that releasing the contents of communications raises insurmountable privacy issues.

be to create a cybersecurity research exception to them. A full proposal for such an exception is discussed in [8].

4 Running Infected Hosts

This section discusses legal and ethical issues that arise in two situations that involve running hosts that are infected with malicious software. First, it may be necessary to allow attackers to remotely exploit hosts in order to collect malware and observe the behavior of both the attackers and the software [19]. Second, researchers may run malware in testbeds in order to observe the software's behavior in a controlled environment.

4.1 Testbeds

The primary legal concern with running malware in testbeds is liability from accidental exfiltration of malicious traffic beyond the testbed. The exfiltration pathway might be a link from the testbed to the Internet that is provided to allow users to run experiments remotely. The Computer Fraud and Abuse Act (CFAA) would be the most likely legal theory for holding researchers liable [2].

The CFAA prohibits a wide variety of conduct directed against essentially any computer connected to the Internet. It prohibits not only targeted break-ins of specific computers, but also knowingly transmitting a program—such as a worm or virus—that damages another computer connected to the Internet.[4] Though this provision would appear to cover code that escapes from a testbed, it is important to note that the CFAA also requires intentional harm to another computer in order to find an offense. A researcher who accidentally allows malicious traffic to escape containment is highly unlikely to possess this intent.

An alternative theory of liability for exfiltrated code is based on tort law, an area of common law, i.e., based on court-created doctrines rather than statutes. One potential tort-based theory is negligence, which is the doctrine that courts apply to compensate injured parties after accidents.[5] Another theory is nuisance, which would involve proving that the leak of malicious code caused "an unreasonable interference with a right common to the general public" [12]. A third possibility is tort liability for ultrahazardous activities, which is governed by a standard of strict liability. In contrast to negligence, which requires proof that a defendant failed to take precautions appropriate to prevent harm (discounted by the probability of harm), strict liability does not involve any notion of fault: if strict liability applies to an activity (a big if) and an accident occurs, the person conducting the activity is liable for injuries to others.

These theories remain hypothetical; no cases have been brought against testbed operators or users, perhaps because of a lack of accidents involving testbeds. Still, should this situation change, each theory discussed above would face significant hurdles. The negligence theory, for instance, would require proof that the testbed did not have adequate measures in place to prevent exfiltration. Since testbed designers take pains to keep open a minimum number of channels of communication between the testbed and the Internet, the chances of finding such a breach of duty seem slim [10]. A second weakness, which also applies to the nuisance theory, is that it is an open question whether testbed operators or users owe a duty of care to other Internet users in the first place. It is worth noting that none of these theories have been successfully used to sue software vendors for harm arising from security vulnerabilities in their software [7]. Finally, strict liability applies to activities that are, among other things, uncommon and pose a risk of accidents that due care cannot prevent, such as blasting with dynamite in urban areas [23]. Though running malicious code on a testbed may not be within the experience of most Internet users, one could argue that that is the wrong frame within which to judge commonality: Internet users are constantly exposed to malicious traffic. Thus, releasing malicious traffic might not be considered uncommon. Strict liability for accidental exfiltration of malicious code from a testbed thus seems unlikely.

4.2 Non-Isolated Hosts

Research that makes use of hosts that are allowed to interact with attackers present a few additional legal considerations. One concern that researchers might have is that allowing a computer to become infected with malware that causes the host to join a botnet violates the CFAA or other laws. Allowing the infection (or collecting malware) itself probably is not illegal under the CFAA, as the researcher does not obtain unauthorized access to another computer. Allowing the infected host to communicate with an attacker via IRC or other means is

[4] Specifically, 18 U.S.C. § 1030(a)(5)(A)(i) prohibits:

> [K]nowingly caus[ing] the transmission of a program, information, code, or command, and as a result of such conduct, intentionally caus[ing] damage without authorization, to a protected computer.

A "protected computer," in turn, means any computer owned by a financial institution or the U.S. government, or any computer used in interstate commerce. 18 U.S.C. § 1030(e). The interstate commerce portion of this definition is sufficiently broad to bring any computer connected to the Internet within the definition of "protected computer."

[5] A successful negligence suit requires proving that (1) the defendant owed the plaintiff and duty of care; (2) the defendant breached the duty; (3) the breach caused harm; and (4) the harm is a legally recognized form of injury.

more subtle. The contents of the commands, such as instructions to request data from a third-party victim, may not be illegal. But responding to these commands—by sending a flood of traffic to an innocent third party as part of a distributed denial of service attack, for example—would raise the concern that the research system is participating in an attack. Deciding on the appropriate balance between collecting information and potential liability under the CFAA thus deserves careful, case-by-case analysis.

A second question is whether researchers could be liable for data, such as copyrighted works or child pornography, that attackers place on their hosts. Attackers might even deliberately target researchers with such materials, if they discover the identity of a research host and wish to cause trouble for the researcher.

Consider the copyright question first. The concern for researchers is that merely possessing an unauthorized copy of a work (music, a movie, a book, etc.) could expose them to liability for infringement. This situation could arise for researchers investigating peer-to-peer systems. Under the Copyright Act (Title 17 of the U.S. Code), if a person takes no action to infringe one of the exclusive rights of a copyright holder, then there is no infringement. In this case, if an attacker downloads infringing copies of copyrighted works to a researcher's computer without the researcher's knowledge, then the researcher is probably not liable for copyright infringement. This situation could change, however, if the researcher analyzes the contents of materials that attackers send. In that case, the researcher may become aware that he or she is in possession of infringing copies; and analysis of the copies could constitute infringement of one or more exclusive rights (e.g., the right of reproduction[6]). Researchers would have a strong argument that such reproduction is a fair use (17 U.S.C. § 107) of the work; but a full analysis of that argument is beyond the scope of this paper. Unless analyzing these materials is important for the underlying research, researchers would be better off deleting such materials or preventing attackers from downloading data in the first place.

Unfortunately, the solutions are not as simple in the case of child pornography. Federal law makes it a crime to knowingly possess any image of child pornography [3]. Thus, if a researcher analyzes the contents of materials downloaded by attackers and finds that child pornography is part of the mix, he or she likely meets the definition of this possession crime. The law does provide a defense if a person possesses fewer than three images, reports such possession to a law enforcement agency, and destroys each image. This defense is narrow, and a researcher who stumbles across child pornography planted by an attacker should immediately contact an attorney. As was the case with copyright infringement, the potential for liability should make researchers think seriously about whether projects require allowing attackers to store data on research machines.

5 Mitigating Attacks

Cybersecurity researchers may also find themselves in a position to disrupt or mitigate attacks. After all, their research may yield detailed knowledge of the workings of malware, botnets, etc. This raises the question of what kinds of mitigations are legally permissible, and which steps are ethical. For the most part, mitigation by researchers raises serious legal and ethical questions and should be avoided. To explore these issues, this section makes use of three specific but hypothetical examples.

Example 1. Suppose that a researcher finds that a botnet command and control server is running software that makes it vulnerable to a remote denial of service attack. Taking this server out of commission might seem worthwhile because it would help to disrupt the botnet, if only temporarily. But to the extent that taking down the server would involve sending code or data resulting in unauthorized access to the server, this action could be a violation of the CFAA. (See footnote 4 above for the pertinent text from the CFAA.) The fact that the server is being used for malicious purposes does not matter to an analysis of the proposed mitigation.

Example 2. As a refinement to this example, suppose that messages of a certain format or length cause the command and control program to crash; a researcher (whose computer was infected with malware that the botmaster controls) considers sending crafted messages to effect a crash. In this case, the researcher is communicating via a channel that the botmaster has selected; the botmaster has arguably consented to receive messages from the computers enslaved in the botnet, giving the researcher a stronger argument that the crafted message is "authorized."

Example 3. A final variation to consider on the legal side of mitigation is introducing bogus data (e.g., fabricated usernames and passwords, or fake credit card numbers) into botnets or other networks controlled by malicious actors. In this case, a researcher would simply place the data on hosts that he or she controls and allow attackers to take the data. This research design has the potential to allow researchers to track the flow of data through malicious networks. Still, even bogus data pose legal issues worth considering. The CFAA prohibits trafficking in passwords with intent to defraud and accessing financial records without authorization (18 U.S.C.

[6] Courts have held that copies made in RAM may infringe the exclusive right of reproduction, even if no permanent copy is made. See, for example, MAI Systems Corp. v. Peak Computer, Inc., 991 F.2d 511 (9th Cir. 1993).

§§ 1030(a)(6) and (a)(2), respectively). Even if offering truly fabricated does not meet all elements of these offenses other issues merit consideration. For example, linking the data to an actual brand name, such as a bank or a credit card network, could raise trademark infringement or dilution issues.

There remain ethical considerations for mitigation steps that are legal. Perhaps the most important consideration is whether mitigation fits the role of a cybersecurity researcher. Different researchers will view their roles differently, depending not only on their personal beliefs but also the type of institution for which they work. Whatever these variations may be, a point that seems likely to be constant is that researchers are employed primarily to *study* threats, rather than to take action against them.

Another ethical consideration is the extent to which mitigation (and other forms of investigation, such as probing networks or running honeynets) might harm the reputation of the researcher's institution. Mitigation may be seen as an action on behalf of the researcher's institution, and the researcher may or may not have this authority. Furthermore, when mitigation would involve action against remote hosts (as was the case with Example 2 above), it raises the possibility of interfering with other efforts to study or disrupt malicious activity, e.g., law enforcement investigations. There may also be a risk of misidentifying the parties responsible for malicious activity; or imperfect or ineffective mitigation might give attackers the opportunity to improve their techniques. For these reasons, researchers should be extremely cautious about taking steps beyond their own networks to mitigate threats. At minimum, they should discuss proposed tactics with IT officers at their institutions and, potentially, with law enforcement officials.

6 Publishing Results

Finally, the topic of publishing results ties together many of the issues discussed so far in this paper. The First Amendment to the U.S. Constitution provides broad protection for publishing cybersecurity-related findings, even potentially damaging disclosures such as zero-day vulnerabilities.[7] Unless a disclosure is part of an agreement with another person to commit some other crime (i.e., it is part of a conspiracy), or is likely to succeed in inciting "imminent lawless action" [26], the First Amendment provides some protection. A publication that merely provide knowledge that might help another person commit a crime is protected speech [28].

The broad protections of the First Amendment, however, are subject to a few qualifications. Perhaps the most important is DMCA's prohibition on trafficking in devices (which includes software), the primary purpose of which is to circumvent a technical protection measure on a copyrighted work. Courts have held that publishing circumvention software, and even linking to a site that offers such software, violates the DMCA [24]. But it is unclear what level of detail triggers the DMCA. For example, after a group of researchers that found vulnerabilities in a digital watermarking scheme was threatened under the DMCA before presenting their work at an academic conference, the U.S. Department of Justice wrote in a court filing that the DMCA did not prohibit publication of the paper or the underlying research [16]. Still, the prospect of liability under the DMCA is sufficiently realistic that researchers who plan to publish about vulnerabilities in software or hardware that protects copyrighted works may wish to consult an attorney before doing so.

Publications also have the potential to harm an institution's reputation by revealing network details that the institution would prefer to keep secret. A strictly legal concern that this raises is a potential breach of contract. Suppose, for example, that an institution holds contracts that specify a network configuration or bandwidth guarantee given to transit or peering partners. Providing details necessary to allow others to understand a data collection set-up or an experiment might reveal that an institution is not living up to its contractual commitments. Again, consultation with information officers in an organization could help allay these concerns. Note that the objective of this coordination is neither to alter the information in a publication nor to force the organization to alter its practices; instead, it is to give an organization an opportunity to identify potential conflicts with contract partners and to plan for remediation.

The possibility that a publication will reveal details about an organization's network also raises issues beyond legal liability. Researchers should also consider whether the papers or datasets that they publish could reveal information that could help adversaries attack the researcher's own network (or other friendly networks). Publishing datasets, as discussed in Section 3.2, is likely to pose a greater risk to an organization's network than a paper; so data releases may deserve a more careful vetting by IT officers than papers do.[8]

The same principles apply to the privacy of users whose network use may be discernible from a dataset. Given recent research demonstrating the difficulty of de-

[7] One exception is for classified systems. Another is for systems examined under a non-disclosure agreement (NDA); a researcher might be liable for damages resulting from a breach of contract if he or she publishes results that violate the NDA.

[8] These officials are usually extremely busy and have limited resources; con vicing them of the benefit of collecting and sharing data that could harm the organization may require considerable relationship-building effort.

vising robust anonymization schemes [9, 14], researchers should be particularly forthcoming about privacy risks before sharing data.

7 Conclusion

The legal environment inhibits cybersecurity research through outright prohibitions and through uncertainties that make some experiments and data collection and sharing efforts too costly to evaluate. Communications privacy laws have also set strong social expectations that network providers will maintain the confidentiality of their data. Though these expectations often do not match reality, they may nevertheless provide a reason that organizations cite to avoid the expense and legal and reputational risk of granting researchers access to network data. Reforming these laws is on the agenda of both privacy advocates and law enforcement agencies. Researchers could participate in reform efforts (e.g., through scholarly meetings and publications, meeting with policymakers, or testifying before them) to make known how the lack of a research exception affects them.

This paper has also attempted to provide a sense of the interests that the laws relevant to cybersecurity are intended to protect. The hope is that this background will help cybersecurity researchers make decisions about their activities in light of broader ethical considerations. These considerations should include not only the users whose activities may be reflected in network data, but also the reputation of the researcher's own organization and the interests of researchers who have supplied, or would like to supply data. More work is needed to develop the relevant ethical framework.

Acknowledgments

I acknowledge support for this work from TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422). I also thank Deirdre Mulligan and Vern Paxson for many helpful conversations, and Mark Allman, kc claffy, and anonymous referees for helpful comments on this paper.

References

[1] 18 U.S.C. § 2510-2522.

[2] 18 U.S.C. § 1030.

[3] 18 U.S.C. § 2252A.

[4] 18 U.S.C. § 2701-2711.

[5] 18 U.S.C. § 3121-3127.

[6] Mark Allman and Vern bPaxson. Issues and etiquette concerning use of shared measurement data. In *Proceedings of IMC '07*, pages 135–140, October 2007.

[7] Douglas A. Barnes. Deworming the internet. *Texas Law Review*, 83:279–329, November 2004.

[8] Aaron J. Burstein. Toward a culture of cybersecurity research. *Harvard Journal of Law and Technology*, 22, 2008.

[9] S.E. Coull, M.P. Collins, C.V. Wright, F. Monrose, and M.K. Reiter. On web browsing privacy in anonymized netflows. In *Proceedings of the 16th USENIX Security Symposium*, pages 339–352, August 2007.

[10] Emulab. Knowledge base entry: Is emulab firewalled? http://www.emulab.net/kb-show.php3?xref_tag=SI-FW, August 2005.

[11] Seymour E. Goodman and Herbert S. Lin, editors. *Toward a Safer and More Secure Cyberspace*. National Academies Press, 2007.

[12] Practicing Law Institute. *Restatement (Second) of Torts*, page § 821B(1). 1977.

[13] Joseph P. Liu. The DMCA and the Regulation of Scientific Research. *Berkeley Technology Law Journal*, 18:501, 2003.

[14] Arvind Narayanan and Vitaly Shmatikov. How to break anonymity of the netflix prize dataset, 2006.

[15] United States Department of Justice, editor. *Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations*. 2002.

[16] U.S. Department of Justice. Brief of the United States in Support of the Motion Felten v. RIAA (Nov. 8, 2001), CV-01-2669 (GEB) (N.D. Cal.).

[17] Paul Ohm, Douglas Sicker, and Dirk Grunwald. Legal Issues Surrounding Monitoring (Invited Paper). In *Internet Measurement Conference*, October 2007.

[18] Ruoming Pang, Mark Allman, Vern Paxson, and Jason Lee. The devil and packet trace anonymization. *Computer Communication Review*, January 2006.

[19] Moheeb Abu Rajab, Jay Zarfoss, Fabian Monrose, and A multifaceted approach to understanding the botnet In *Proceedings of the IMC*. ACM, October 2006.

[20] Richard Salgado. *Know Your Enemy*, chapter Legal Issues, pages 228–252. Addison-Wesley Professional, 2004.

[21] Daniel J. Solove. Digital dossiers and the dissipation of fourth amendment privacy. *Southern California Law Review*, pages 1083–1167, 2002.

[22] Daniel J. Solove and Chris Jay Hoofnagle. A model regime of privacy protection. *University of Illinois Law Review*, pages 356–403, 2006.

[23] Indiana Harbor Belt Railroad Co. v. American *916 F.2d 1174*. (7th Cir. 1990).

[24] Universal City Studios Inc. v. Corley. *273 F.3d 429*. (2d Cir. 2001).

[25] United States v. Forrester. *495 F.3d 1041*. (9th Cir. 2007).

[26] Brandeburg v. Ohio. *395 U.S. 444*. 1969.

[27] Organizacion JD Ltda. v. United States Dep't of Justice. *18 F.3d 91*. (2d Cir. 1994).

[28] Euguene Volokh. Crime-facilitating speech. *Stanford Law Review*, 57:1095–1222, March 2005.