

Optimizing Lawful Responses to Cyber Intrusions

Thomas C. Wingfield¹, James B. Michael^{2,*}, Duminda Wijesekera³

¹*Potomac Institute for Policy Studies, Arlington, Va., U.S.A. twingfield@potomacinstitute.org*

²*Naval Postgraduate School, Monterey, Calif., U.S.A. bmichael@nps.edu*

³*George Mason University, Fairfax, Va., U.S.A. dwijesek@gmu.edu*

*** Corresponding author**

Paper no. 290

Optimizing Lawful Responses to Cyber Intrusions

Thomas C. Wingfield¹, James B. Michael², Duminda Wijesekera³

¹*Potomac Institute for Policy Studies, Arlington, Va., U.S.A. twingfield@potomacinstitute.org*

²*Naval Postgraduate School, Monterey, Calif., U.S.A. bmichael@nps.edu*

³*George Mason University, Fairfax, Va., U.S.A. dwijesek@gmu.edu*

Abstract: Cyber intrusions are rarely met with the most effective possible response, less for technical than legal reasons. Different rogue actors (terrorists, criminals, spies, etc.) are governed by overlapping but separate domestic and international legal regimes. Each of these regimes has unique limitations, but also offers unique opportunities for evidence collection, intelligence gathering, and use of force. We propose a framework which automates the mechanistic aspects of the decision-making process, with human intervention for only those legal judgments that necessitate human judgment and official responsibility. The basis of our framework is a pair of decision trees, one executable solely by the threatened system, the other by the attorneys responsible for the lawful pursuit of the intruders. These parallel decision trees are interconnected, and contain pre-distilled legal resources for making an objective, principled determination at each decision point. We offer an open-source development strategy for realizing and maintaining the framework.

Key words: Law, information warfare, intrusion response, decision support, open source, Schmitt analysis

1. Dangers of Oversimplified Responses

When a person of ill intent, which we shall refer to as a rogue actor, intrudes into a computer system, misuses a computer system, or attacks a computer system, the owner of that system or the owner's agent needs to know something about the rogue actor in order to develop a tailored response to the rogue actor's behavior. Applying a "one-size-fits-all" response, such as always terminate all interaction with the rogue agent or always respond in kind, can be an ineffective or worse, illegal, response in some cases. For instance, terminating interaction with a rogue actor may prevent the collection of evidence for criminal prosecution, counter-targeting for military response, or collection for a counterintelligence operation. By responding in kind, or conducting some form of cyber vigilantism as described in [Jayaswal 2002], the owner or the owner's agent may violate domestic laws, or if the attack is deemed to be a "use of force," may contravene the customary rules of war (accepted as authoritative law by the United States and punishable under 18 U.S.C. §1097).

We have approached this problem in earlier work, examining the need for a legal framework in dealing with computer attacks on high-profile systems [Michael 2002b, 2003a] and for cyber and kinetic attacks on the Washington, D.C. metro system [Michael 2003b]. In those case studies, the lack of adequate legal and operational preparation made it difficult if not impossible to formulate a timely, lawful, and effective response. Furthermore, the unique legal aspects of cyber attacks require both a return to first principles and a mechanism for developing new analyses. We extend this work by addressing the fundamental legal concern in this entire area: providing owners and agents with sufficient information in order to make informed decisions when formulating responses to rogue actors. The specific problem we address is how to address the central question: "What do attorneys need to know about a rogue actor in order to apply the correct legal regime within which to advise their clients about alternative responses to the rogue actor?" We assume that owners and their agents want to defend their computer systems without violating domestic and international law.

2. The Need for Legal Preparation

Both the rate and intensity of attack in cyberspace can be high, affording little time to respond before the cyber battle is over. Similarly, what may initially appear to be a minor intrusion or misuse of a computer system can ultimately result in damage to or destruction of property, or even human injury or loss of life. In either case, the owner and the owner's agents must be prepared to respond to such attacks with plans and mechanisms in place to gather and process information to answer the aforementioned question. In other words, the owner and agent need to tighten their Observation-Oriented-Decision-Action (OODA) loop [Boyd 1986] in order to gain a competitive advantage over the rogue agent. In order to achieve this, the owner and agent need to be operationally prepared.

However, operational preparedness is only part of the equation; one also needs to be legally prepared. One cannot, without undue risk, respond without first considering the legality of the response. Against opponents who disregard any laws which are not immediately and effectively punitive, the default response of inadequately counseled operators is to forego otherwise lawful and effective defensive strategies. In other words, the vast legal gray area which exists today operates in favor of the attacker. A clearer and timelier picture of the operational legalities of the situation would provide the defender with more, rather than fewer, options.

3. Complexity and Scalability

The scale of the problem increases as the cardinality of interaction between parties changes from that of one-to-one to one-to-many or many-to-many. For example, multiple rogue agents could attack a single system or network of systems that have a single owner or defending agent, or multiple rogue agents could misuse, such as in a distributed denial-of-service attack, a network of computers that are owned or defended by different parties. For instance, suppose, in the latter case, that there are three rogue agents who launch a coordinated attack against a U.S. Government computer network: a U.S. military officer who has legitimate access to the computer network but misuses the computer with the intent to allow foreign nations to attack the network, a foreign information warrior who is given the assignment by his government to attack the network, and a U.S. citizen who is funded by a foreign government to launch covert attacks on the network. In this case, the owner of the computer network is the U.S. Government, and its agents for responding to the attack include actors from the military, law enforcement, and intelligence communities.

The law enforcement personnel, in this case, are the "first responders," so the observed rogue-like behavior is treated as a law enforcement situation—absent otherwise lawful presumptions, one must use the most restrictive legal rule set at the outset of a response. After additional information has been gathered, the U.S. Government may be able to transition to a more appropriate rule set to deal with spies, terrorists, soldiers, and other specific types of rogue actors. As law enforcement agents learn more about the rogue actors, they may discover the source of the attacks or even something about the attackers. This information can then be used to determine, based on domestic and international law, what role the other responders can play in responding to the rogue agents: the intelligence community to address the role of the foreign national (but not on U.S. persons), and the military to assist in all aspects of the response except for law-enforcement duties such as apprehending the U.S. noncombatant (*i.e.*, private citizen). Note that the responder must know what laws apply to each party involved in the interaction.

4. Legal-Technical Interaction

Progress has been made in devising technical mechanisms for sensing, processing, and reporting information in real- or near-real time, in addition to offline (for forensics purposes), about computer intrusions, computer misuse, and computer attacks. For instance, [Michael 2002a] describes a general class of mechanisms, known as intelligent software decoys, for deceiving rogue actors into revealing information about themselves. These active defense mechanisms are programmed with rules for engaging a rogue actor for the purpose of automatic data collection and active response—either dissuading further interaction or prosecuting an armed response (in the legal sense). The intelligent software decoys report their progress to human owners and agents so that the human can make decisions manually, where appropriate, on how to respond to the behavior of a rogue actor [Michael 2002b]. However, to support legal and operational preparedness goals, the strategy and tactics employed by intelligent software decoys need to be driven by the requirements for answering the central question. In the remainder of this paper, we describe a computational framework couched in terms of the legal- and operational-preparedness goals, from which any class of automated or manual response mechanisms can be fashioned. Our overriding goal is to provide for computers and humans to respond in tandem once the defender “knows” enough about the identity and intent of the rogue agent.

5. Analytical Framework

Today attorneys answer the central question using manual means. There are two components to determining the categorical legal identity of the rogue agent: (i) presumptions (“any one who comes into the system is assumed to be trespassing,” etc.) and (ii) specific actions of individual rogue agents. What we propose is to build a model of domestic and international law as it applies to cyber intrusions, consisting of two interconnected decision trees, one for computers to execute autonomously and at high speed, and a second requiring human decision making at considerably lower speed. While the computer tree will be “hardwired” for independent execution of clearly discernable, objectively verifiable criteria, the human tree will have pre-selected sources available to assist the attorney in deciding each of the “gray area” judgment calls requiring human reflection and creativity.

Table 1. Comparison of computer and human decision trees

Attribute	Computer tree	Human tree
Speed of decision making	High	Low
Need for human reflection and creativity	Low	High
Reliance on clearly discernable, objectively verifiable criteria	High	Low

5.1 Sources of information

It will be necessary to assemble a comprehensive selection of sources to append to each decision point, but it will be vital, for speed and clarity, to include no more than is required to answer the question at hand. These sources may be grouped as constitutional, legislative (statutes), executive (regulations), judiciary (cases), and international. These five categories must be further subdivided into primary (e.g., the case or statute itself), and secondary (analytic and synthetic commentary, such as law review articles). These ten categories could contain any legal source needed to address any given question.




Fig. 1 Sources of information

5.2 Levels of abstraction

Multiresolution modeling [Davis 1998] will be needed to support the computer and human decision makers in obtaining the proper balance of speed and depth for specific decision-making tasks, with each source represented at four levels of abstraction: (i) citation (a legal footnote), (ii) précis (a sentence or paragraph paraphrasing what that source has to say about the question at hand), (iii) excerpt (direct quotes from the source which are on point), and (iv) full document (the complete law review article, statute, or case). In other words, the computer or human must be able to adjust the level of fidelity at which it views the data for creating a legal brief and reasoning about the information contained in the brief, in support of making decisions. For instance, in a group decision-making setting, the facilitator must direct the attention of the team of attorneys between detailed and aggregate source material contained in the legal brief, such as when determining which legal regime applies based on the results of conducting a Schmitt Analysis [Schmitt 1998] of the consequences of a cyber attack.

This general information would be distilled into a specific research question in two media: an audit trail, providing a record of each question asked and each answer chosen, and a brief builder, which would augment the audit trail with those portions of the sources selected by the reviewing attorney to support his answer to the question. This would, in effect, be the first draft of a legal brief supporting the selected course of action.

5.3 Open-source approach to developing the framework

Finally, and most crucially, these two interconnected legal trees, and their supporting sources, would be constructed using the open source methodology most famously employed by Linus Torvalds and the Linux operating system. After the process architecture had been established by a core team of attorneys and computer scientists, the trees would be available to legal academia (law students, practitioners, and professors, participating individually and through conferences, courses, *pro bono* projects, and continuing legal education seminars) for part-time analysis and improvement. This approach would provide three strong advantages: First, the best and broadest academic research and analysis could be solicited, providing the most robust possible input; second, the cost of such a daunting project would be drastically reduced by leveraging the efforts of the non-profit-seeking half of the legal profession, harnessing a small portion of the unfocused

academic effort that goes into building and maintaining an intellectually competent bar. A moderately sized management staff could act as the integrators, much as Torvald and his inner circle manage the contributions of thousands. Third, such an approach would be the political antithesis of the ill-fated U.S. Defense Advanced Research Project Agency's Total (or, later, Terrorist) Information Awareness program [Cherry 2003]. The overwhelmingly negative reaction that program received demonstrated the political danger in allowing any such project to be perceived as an extension of "Big Brother" and an unnecessarily closed effort by a national government [NYT 2003].

In contrast to TIA, our approach to developing a framework would allow the greatest possible contribution from informed and capable academics and practitioners in the legal community. Its inherent transparency would define it as the "counter-TIA," and would be much easier to fund, develop, and deploy. It is a fundamental tenet that there is no classified law (as opposed to necessarily classified regulation and operational information), so the legal portion of defense in cyberspace could be accomplished in the open with no decrement to security.

The "white" or unclassified nature of this project would not interfere with its operational usefulness. At regular intervals, a "snapshot" of the two interconnected trees could be taken and downloaded into a "black" or classified computer system, insulated from the white world by an air gap. This tree would then be isolated and usable for operational planning. Doing this regularly would provide constant updates to the unclassified basis for making classified decisions. The legal analysis completed, classified policy options would be clearly open or foreclosed, and the operators, mission planners, intelligence officers, and commanders would have a secure basis for making time-critical decisions while under attack. To complete the cycle, real-world problems could be sanitized and returned to the "white" world for academic analysis, informing the development of the law in such a way as to minimize academic departure from operational reality.

Similar to Torvold's approach in developing and maintaining Linux, we envision that carrying out such a program would require a small core staff of attorneys and computer scientists to design the substantive and procedural architecture of the open source template. However, one might argue that this core group might become a bottleneck, as pointed out by [Lewis 1999]:

There are other labor problems associated with the anarchic open source development model. Simple organizations work best when the product is simple. But when the product becomes complex, an informal organizational structure struggles to keep on top of it. Even Linus Torvold has limits. As Linux grew, Torvold began delegating large components to his trusted lieutenants, who in turn started delegating portions of their area of responsibility to others. The frequency of releases has slowed because the "sheer size of the code base has begun to overrun the resources of Linus...there is a backlog of patches to be merged and often, Linus is becoming the choke point."

In contrast to the Linux model of open source modeling, as described earlier we take a view of allowing anarchy to rein on the "white" side, while enforcing discipline on work performed on the "black" side. We believe this separation of the two communities—academia and operations—will help us pass Lewis' "acid test of mainstream viability" of the open source development. The open source model has worked well for quite some time in the legal community: law reviews are but one example.

Once open to academic participation, this cadre would manage inputs and make the final decisions in pruning or grafting new branches onto the trees, and in modifying the choice of sources available at each decision point along the human tree. Properly executed, such a project could reasonably be expected to yield impressive operational, economic, and political results.

6. Conclusion

An academically comprehensive and operationally useful legal framework is needed to address the growing threat of cyber intrusions, particularly against national critical infrastructures and mission-critical systems. The importance of protecting these assets effectively and lawfully is difficult to overstate. We propose a thorough review of the law governing these intrusions, and its distillation into two interconnected decision trees. The first would be executed by the threatened system itself in real time, and would require only the clearest and most objectively verifiable criteria for its decision-making inputs. The second would be for human use, containing at each decision point the legal resources (presented in four levels of abstraction) required to make nuanced, principled decisions in near-real time. This framework would be the basis for the seamless application of the law to criminal, military, and espionage activities in cyberspace. It would be of incredible complexity, but could be built and maintained using an open source architecture. This approach would provide the greatest academic input at the lowest cost, and would provide a methodology clearly distinguishable from politically unpalatable efforts of the past.

Acknowledgements

Conducted under the auspices of the Naval Postgraduate School's Homeland Security Leadership Development Program, this research is supported by a grant from the U.S. Department of Homeland Security. Sponsorship of this research also was provided by the Critical Information Infrastructure Program at George Mason University. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright annotations thereon.

References

- [Boyd 1986] Boyd, J. R. A Discourse on Winning and Losing: Patterns of Conflict. Lecture notes, U.S. Department of Defense, Pentagon, Washington, D.C., Dec. 1986. (Typewritten)
- [Cherry 2003] Cherry, S. M. TIA is dead – Long live TIA. *IEEE Spectrum*, Nov. 2003, p. 22.
- [Davis 1998] Davis, P. K. and Bigelow, J. H. *Experiments in Multiresolution Modeling*. Santa Monica, Calif.: RAND National Defense Research Institute, 1998.
- [Jayaswal and Doss 2002] Jayaswal, V., and Doss, D., Internet Hack Back: Counter Attacks as Self-Defense or Vigilantism? In *Proc. Int. Symposium on Technology and Society*, IEEE (Raleigh, N.C., June 2002), pp. 380-386.
- [Lewis 1999] Lewis, T. The open source acid test. *IEEE Computer*, Feb. 1999, pp. 125-128.

- [Michael 2002a] Michael, J. B., Auguston, M., Rowe, N. C., and Riehle, R. D. Software decoys: Intrusion detection and countermeasures. In *Proc. Workshop on Inf. Assurance*, IEEE (West Point, N.Y., June 2002), pp. 130-138.
- [Michael 2002b] Michael, J. B. On the response policy of software decoys: Conducting software-based deception in the cyber battlespace. In *Proc. Twenty-sixth Annual Computer Software and Applications Conf.*, IEEE (Oxford, Eng., Aug. 2002), pp. 957-962.
- [Michael 2003a] Michael, J. B. and Wingfield, T. C. Lawful cyber decoy policy. In Gritzalis, D., di Vimercati, S. D. C., Samarati, P., and Katsikas, S., eds. *Security and Privacy in the Age of Uncertainty*. Boston, Mass.: Kluwer Academic, 2003, pp. 483-488.
- [Michael 2003b] Michael, J. B., Wingfield, T. C., and Wijesekera, D. Measured responses to cyber attacks using Schmitt Analysis: A case study of attack scenarios for a software-intensive system. In *Proc. Twenty-seventh Annual Int. Computer Software and Applications Conf.*, IEEE (Dallas, Tex., Nov. 2003), pp. 622-627.
- [NYT 2003] The right and wrong stuff of thinking outside a box, *New York Times*, July 31, 2003, p. A17.
- [Schmitt 1998] Schmitt, M. N. *Bellum Americanum*: The US view of Twenty-first Century war and its possible implications for the law of armed conflict. *Mich. J. Int. Law* 19, 4 (1998), pp. 1051-1090.