# Tracking and Tracing Cyber-Attacks: Technical Challenges and Global Policy Issues

Howard F. Lipson, Ph.D.
CERT® Coordination Center

*November 2002*

SPECIAL REPORT
CMU/SEI-2002-SR-009

Carnegie Mellon
Software Engineering Institute

Pittsburgh, PA 15213-3890

# Tracking and Tracing Cyber-Attacks: Technical Challenges and Global Policy Issues

CMU/SEI-2002-SR-009

Howard F. Lipson, Ph.D.
CERT® Coordination Center

*November 2002*

**Networked Systems Survivability Program**

Unlimited distribution subject to the copyright.

This work is sponsored by the U.S. Department of State. The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2002 by Carnegie Mellon University.

Requests for permission to reproduce this document or to prepare derivative works of this document should be addressed to the SEI Licensing Agent.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. Requests for permission to reproduce this document or prepare derivative works of this document for external and commercial use should be addressed to the SEI Licensing Agent.

This work was created in the performance of Federal Government Contract Number F19628-00-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

For information about purchasing paper copies of SEI reports, please visit the publications portion of our Web site (http://www.sei.cmu.edu/publications/pubweb.html).

® CERT and CERT Coordination Center are registered in the U.S. Patent and Trademark Office.

# Table of Contents

# List of Figures

# Acknowledgements

I'd like to express my gratitude to Dr. Sven Dietrich, my colleague at the CERT/CC, for many in-depth discussions on track and trace technology and policy, numerous pointers to the literature, and his valuable comments on earlier drafts of this report.

I'd like to thank Dr. Tom Longstaff, manager of research at the CERT/CC, for his technical insights and for continuing to provide a stimulating and productive research environment. I'd also like to thank Dr. Albrecht Funk, at the Heinz School of Public Policy and Management, Carnegie Mellon University, who reviewed this work and provided valuable insights on policy issues. I'm grateful to my CERT/CC colleagues Martin Lindner, Dr. John McHugh, and Stephanie Rogers for insightful technical and policy discussions, and to Thomas Grasso, Jr. of the FBI and Cornelius Tate of the U.S. Secret Service for providing a most helpful law enforcement perspective. I appreciate the very helpful comments provided by William Fithen of the CERT/CC after his review of an earlier version of this report. Bob Rosenstein, the SEI's Account Executive for the State Department, was invaluable in keeping the project on course. Finally, I'm grateful to Pamela Curtis for editing the final report under a tight deadline and to Mindi McDowell for her editing help with earlier drafts.

# Abstract

In the cyber world, the current state of the practice regarding the technical ability to track and trace Internet-based attacks is primitive at best. Sophisticated attacks can be almost impossible to trace to their true source using current practices. The anonymity enjoyed by today's cyber-attackers poses a grave threat to the global information society, the progress of an information-based international economy, and the advancement of global collaboration and cooperation in all areas of human endeavor.

Part I of this report examines the current state of the Internet environment and the reasons why tracking and tracing cyber-attackers is so difficult. Part II examines some promising research on technical approaches that may greatly improve the ability to track and trace cyber-attacks to their source. Also discussed are some policy considerations with regard to privacy, information sharing, liability, and other policy issues that would be faced by the U. S. State Department in negotiating international agreements for cooperation and collaboration in the tracking and tracing of cyber-attacks. The report concludes with a closer look at technical and policy considerations for next-generation Internet protocols to enhance track and trace capabilities.

# Part I: Technical Challenges in Tracking and Tracing Cyber-Attacks

# 1 Introduction

The creation and phenomenal growth of the Internet has spawned the emergence of a global information society. Businesses possessing highly distributed information assets can function internationally with great efficiency, exchanging information quickly and seamlessly among their divisions, partners, suppliers, and customers. Governments use the Internet to provide information to their citizens and to the world at large, and they will increasingly use the Internet to replace manual methods of collecting information and providing government services. Governmental use of the Internet will increasingly extend to international information sharing and collaboration. The scientific, engineering, and educational communities are all using the Internet as an indispensable tool for collaboration and rapid dissemination of information on advances in research and practice at all levels of scientific and engineering endeavor.

Critical national infrastructures supporting such vital areas as power, transportation, communications, banking and finance, and defense are growing progressively more dependent upon Internet-based applications. The older, often manual, closed, and proprietary methods of providing the essential services that societies depend on are gradually disappearing as they are replaced by cheaper, open, more efficient, and highly distributed Internet applications.

All of the benefits that the Internet and the global information society can provide, including support for the most basic and essential services that nations depend on, are subject to disruption by Internet-based cyber-attacks that use the global computer network to cross international boundaries with ease. Historically, attacks on a nation's essential services typically required a physical attack that crossed the nation's borders slowly enough that it was subject to recognition and interception by that nation's military. At the very least, some physical evidence would likely be left that would allow for the tracking, tracing, and identification of the perpetrators and the tools or weapons used in the attack. Today, cyber-attackers use the speed and global connectivity of the Internet to make national boundaries irrelevant, and sophisticated attackers leave little in the way of electronic evidence that can be used to track or trace them.

The domestic and international implications of an increasingly critical societal dependence on the Internet makes necessary the ability to deter, or otherwise minimize, the effects of cyber-attacks. The capability of a nation (or a cooperating group of nations) to track and trace the source of any attacks on its infrastructures or its citizens is central to the deterrence of such attacks and hence to a nation's long-term survival and prosperity. An acknowledged ability to track and trace both domestic and international attackers can preempt future attacks through fear of reprisals such as criminal prosecution, military action, economic sanctions, and civil lawsuits. The process of trac-

ing an attack back to its source may also help to uncover sufficient details of the attack techniques to allow for the development of defensive measures that could prevent similar attacks in the future. Moreover, tracking and tracing an attack quickly could enable the interruption of an attack in progress. As the critical nature of Internet-based applications and services continues to increase, the ability to deter, prevent, or interrupt attacks in progress will be of greater value to society than assigning blame and collecting damages after a disaster has occurred.[1]

In the cyber world, the current state of the practice regarding the technical ability to track and trace Internet-based attacks is primitive at best. Sophisticated attacks can be almost impossible to trace to their true source using current practices. The anonymity enjoyed by today's cyber-attackers poses a grave threat to the global information society, the progress of an information-based international economy, and the advancement of global collaboration and cooperation in all areas of human endeavor.

[1] Of course, the perpetrator of a particular incident might never be identified.

# 2 A Brief History of the Internet

The Internet (originally known as ARPANET) began its life in 1969 as a research network sponsored by the Advanced Research Projects Agency (ARPA) of the Department of Defense (DoD). By the end of that year, the ARPANET consisted of four nodes, connecting four U.S. universities,[2] and by 1973 the first international connections to the ARPANET were made.[3] In 1982, the earliest definitions of an *internet* as an interconnected network of networks began to appear, along with the establishment of a standard *network protocol*[4] suite to support inter-networking communications. This protocol suite, comprised of an *Internet Protocol* (IP) and a *Transmission Control Protocol* (TCP), became widely known as "TCP/IP," which still forms the foundation of network communications on today's Internet. In December 1988, as a direct result of the first major computer security incident on the Internet (the Morris Worm in November 1988), DARPA[5] founded the CERT® Coordination Center (then known as the "Computer Emergency Response Team") to provide a central place for coordinated responses to Internet cyber-attacks.

Today, the Internet is an interconnected network of networks comprised of approximately 150 million hosts worldwide.[6] The number of computer security incidents handled by the CERT Coordination Center (CERT/CC) has grown from 6 in 1988, to 52,658 in 2001. By the end of September 2002, the CERT/CC had already seen over 73,000 incidents.[7] And yet, despite serious security shortcomings, TCP/IP is still the standard protocol suite for network communications on the Internet, greatly limiting our ability to track and trace Internet cyber-attacks to their source.

---

2 The first four nodes of the ARPANET connected the University of Southern California (UCLA), Stanford Research Institute (SRI), the University of California, Santa Barbara (UCSB), and the University of Utah.

3 University College of London (England) via NORSAR (Norway)

4 A *network protocol* is a common language for communicating across a network. The protocol specifies the rules for data format and transmission.

5 By then ARPA had changed its name to DARPA (Defense Advanced Research Projects Agency).

6 As of January 2002, the number of hosts advertised in the Domain Name Service (DNS) was 147,344,723. Source: Internet Software Consortium (http://www.isc.org/).

7 For additional CERT/CC statistics, see http://www.cert.org/stats/.

# 3 A Brief Tutorial on Internet Technology

At the highest level of abstraction, the Internet is a network of interconnected networks comprised of a myriad of host computer systems joined together by communications links (wired and wireless). Host computer systems communicate by sending messages to each other over the communication links, where a standard network protocol suite called TCP/IP specifies the data formats and transmission rules.

Looking at a greater level of detail, one sees that the Internet is a *packet-switched* network. Messages to be transmitted across the Internet are broken into manageable chunks called *IP packets*, which contain the data to be sent (such as part of an email message or web page content); the *destination* address (i.e., the "TO" field); the *source* address (i.e., the "FROM" field); a *port* number (which represents a specific type of service offered by a host, such as mail transfer, file transfer, or web browsing); and other miscellaneous header information that supports the reliable transmission of data. A vast array of network devices called *routers* forward each packet, so that it moves from router to router until it arrives at its desired destination, or until the number of routers touched by the packet exceeds a maximum allowed value.[8] The forwarding of a packet from one router to a second router is called a *hop*. Each router has a table of routing information (containing a snapshot of the network topology) that it builds and updates by communicating with other routers. The router uses this routing information in an attempt to choose the best path for sending a packet towards its desired destination. In simplest terms, a router looks at the destination address of a packet and sends the packet to another router that moves the packet closer to its final destination.

Typically, there are many alternative paths to a destination, which enables packets to be routed around communications links or routers that are out of service due to attack, accident, or even maintenance. Thus, this packet-switched design allows the Internet to be robust in the face of accidents or external physical attacks on the routing infrastructure. This, in fact, was one of the original design goals of the Internet.

TCP attempts to ensure that the packets are successfully delivered, in the proper order, and it will retransmit packets in the event of their loss. TCP allows two hosts to establish a connection or a

[8] One of the header fields in the packet contains a numerical counter called Time To Live (TTL) which is decremented by one for each hop (i.e., each time a router processes the packet) to prevent infinite looping. When the counter reaches zero, the packet is no longer forwarded, and a control packet (i.e., a notification) is sent back to the source address.

session for a period of time to exchange streams of packets to support a service such as email, FTP, or web browsing.

Each host on the Internet has an *IP address*, which allows packets to be delivered to and received from a specific host. An IP address consists of four decimal numbers in the range 0–255, separated by periods.[9] For example, the IP address of the CERT Coordination Center's web server is 192.88.209.14. However, it is clearly more mnemonic and convenient for a human to enter a name like "www.cert.org" in a web browser window than to enter a purely numeric address. An Internet-wide Domain Name System (DNS)[10] supports the translations of IP addresses into hostnames, using a huge distributed database and a vast number of cooperating systems (*DNS name servers*) to provide the DNS name translation service.

[9] Each decimal number represents 8 bits of a 32-bit binary representation of the IP address.

[10] For further information on DNS, see http://www.isc.org/products/BIND/.

# 4 Problems with Internet Security

Perhaps the greatest threat to the Internet today is the abysmal state of security of so many of the systems connected to it. There are many contributing factors, including commercial off-the-shelf (COTS) software, in which the number of features and rapid time to market outweigh a thoughtful security design. New vulnerabilities are continually being discovered in such software. The widespread use of many COTS products means that once a vulnerability is discovered, it can be exploited by attackers who target many of the thousands or even millions of systems that have the vulnerable product installed. A lack of security expertise by most Internet users means that vendor security patches to remove the vulnerabilities will not be applied promptly, if at all. As a result, systems with unpatched vulnerabilities can be easily compromised, in large numbers, by motivated attackers, who will then use these systems as launching points to concentrate an attack against better-protected systems and to hide the tracks of the attacker.

Figure 1 shows that although the sophistication of Internet attacks has increased over time, the technical knowledge of the average attacker is declining, in the same manner that the technical knowledge of the average user has declined. What the graph represents can best be explained by the fact that sophisticated attackers routinely build attack scripts and toolkits that the novice attacker can use with the click of mouse, with devastating effects. Hiding the tracks of the attacker and expunging or concealing any related evidence is an integral part of many attacker toolkits today.

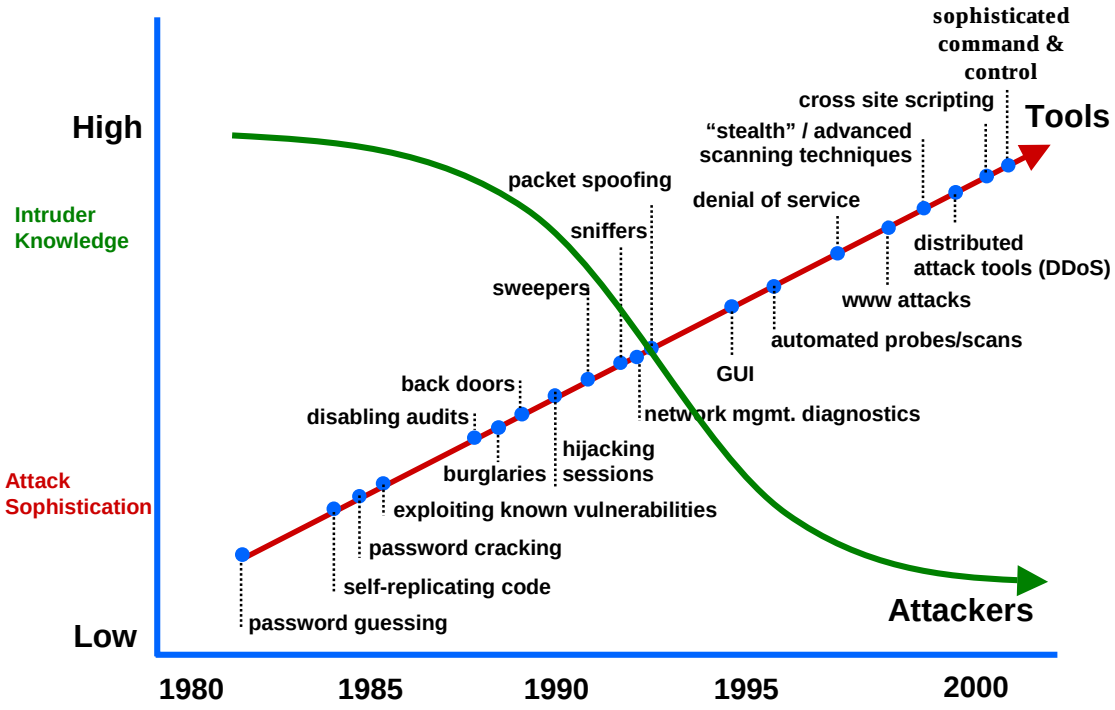


*Figure 1: Attack Sophistication vs. Intruder Technical Knowledge*[11]

Figure 2 illustrates the vulnerability exploit cycle, in which the height of the graph represents the number of security incidents related to a given vulnerability. In the beginning, a sophisticated attacker discovers a new vulnerability, and a few systems are compromised in an early attempt to exploit the vulnerability. Eventually, the widespread availability of exploit tools leads to a peak in the number of incidents. The release of security advisories and the availability of vendor patches help to finally reduce the number of incidents related to the vulnerability.

[11] Source: CERT Coordination Center, © 2002 by Carnegie Mellon University.

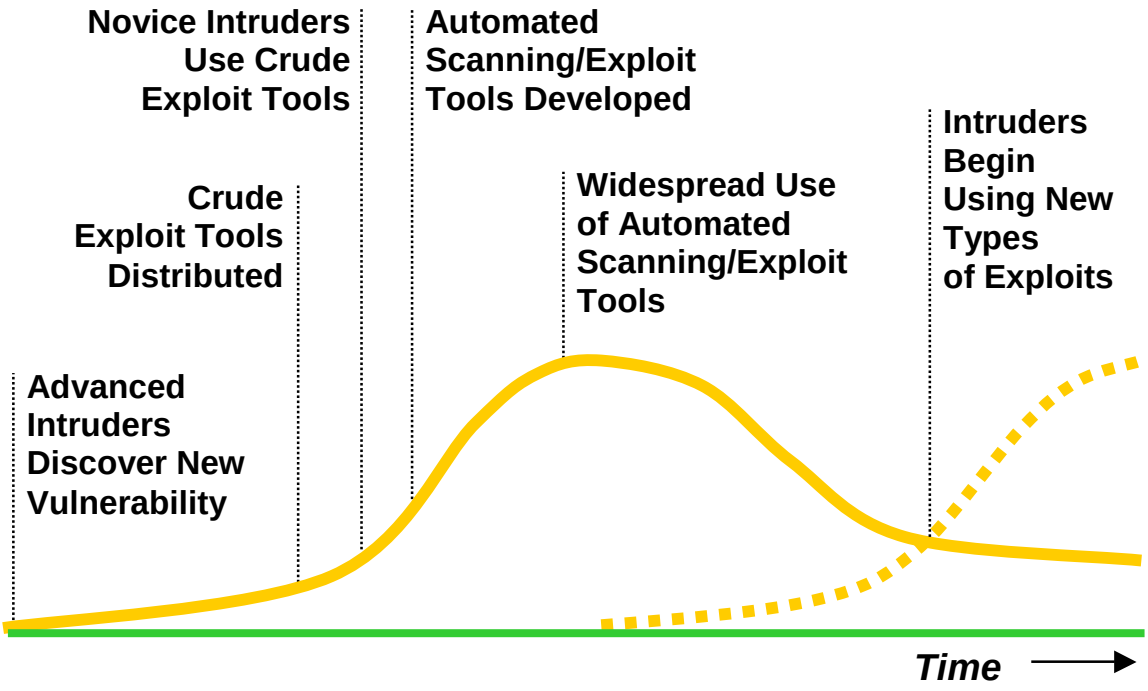


*Figure 2: Vulnerability Exploit Cycle*[12]

The vulnerability exploit cycle provides convincing evidence that a purely defensive "patch and pray" approach to computer security will never be sufficient to protect the extremely critical applications that are increasingly becoming a part of the Internet landscape. For such critical applications, a cyber-attack can mean significant economic loss and even the loss of human life. This means that the deterrent value of successfully tracking and tracing attackers will become increasingly vital to the survival of the Internet and the nations that depend on it.

[12] Source: CERT Coordination Center, © 2002 by Carnegie Mellon University.

# 5 Shortfalls in the Current Internet Environment

This chapter highlights many of the problems with the current Internet environment that make tracking and tracing attackers an extremely difficult task.

## 5.1 The Internet was never designed for tracking and tracing user behavior

Unlike the telephone system, which has an effective tracking and billing capability based on the need to charge users of its services on a per-call basis, the Internet has no standard provisions for tracking or tracing the behavior of its users. The original Internet model never envisioned billing on the basis of single host-to-host connection or some fine-grained unit of delivery of service (such as a small charge for each email message sent or received). Originally, access to the Internet was sponsored by ARPA and was free to its end users. Current Internet charges are based on large-grained service provisions, such as monthly connectivity, connection speed, and storage capacity, all of which do not require Internet service providers (ISPs) to track the fine-grained behavior of their customers in order to bill them.

Moreover, the Internet was originally designed to link together a cooperative and collaborative community of researchers. It was designed to explore the many possible uses of networking, and in particular to facilitate research collaboration. Facilities for tracking and tracing the behavior of such benign users were never a consideration, and a tracking infrastructure was never designed or implemented. ARPA funds were instead focused on improving the functionality and performance of the Internet.

## 5.2 The Internet was not designed to resist highly untrustworthy users

One of the original design goals of the Internet was that the network be survivable in the face of external physical attack or accident that damaged some significant portion of the routing infrastructure. With its vast array of communications links and routers, the current Internet readily meets that design goal. The routing infrastructure has the inherent ability to route traffic along alternative paths that bypass damaged portions of the network.

Although the original design goals of the Internet recognized the need to be robust in the presence of external attack, there was no equivalent concern with regard to the possibility of internal[13] cyber-attacks by the Internet's own users. After all, the original users of the Internet were university researchers, seeking to collaborate with their colleagues and openly share and exchange information for the benefit of all. Thus, the Internet designers' view was that the user community was essentially benign and trustworthy, so simple security measures (such as passwords) would be sufficient to ensure the protection of the Internet community. Under this worldview, provisions to track and trace malicious user behavior were never designed or implemented.

## 5.3 A packet's source address is untrustworthy, which severely hinders tracking

One of the consequences of Internet protocols being designed under the assumption of a benign and trustworthy user community was the lack of any provision for cryptographic authentication of the information contained in IP packets. Therefore, an advanced user can readily modify any information in an IP packet and, in particular, can forge the source address of a packet, effectively hiding its true origin. For a one-way communication, an attacker needs only to insert a false address in the source field, a relatively straightforward task. For a two-way communication, forging the source address is more complex, but the techniques are well known to the attacker community.

Another way to obscure the true origin of an attack on a desired target host is to first compromise a number of intermediate hosts and to then use them as stepping stones on the way to the final target. It's not uncommon for a large number of stepping stones to be used, and from the target's perspective, the attack will appear to have originated from the last stepping stone (i.e., the compromised host that transmitted the attack packets directly to the target).[14] This "packet laundering" technique is quite effective at thwarting traceback attempts, particularly if there is a significant time delay between any attacker activities[15] involving the stepping stones and the attack on the final target. Moreover, an attacker may arrange things[16] so that the packets transmitted from stepping stone *A* to stepping stone *B* may be quite different in nature from the packets transmitted from stepping stone *B* to stepping stone *C* on the way to the final target *T*, making traceback even more difficult because attempts to trace the attack by correlating similar packets will fail.

---

[13] From the perspective of the Internet as a whole, any user (including a malicious user) is an insider once he or she is on the network.

[14] Yin Zhang and Vern Paxson have developed a traffic analysis algorithm that uses packet size and timing to help detecting stepping stones [Zhang 2000]. Zhang and Paxson note that there are legitimate reasons for using stepping stones, so the algorithm must be supplemented by a policy component to help distinguish legitimate uses from malicious ones.

[15] These activities may include the act of compromising a stepping stone or signaling a stepping stone to attack another host.

[16] This could by done by varying attack techniques in moving from stepping stone to stepping stone during the compromise phase, or by varying (or encrypting) the signals/commands sent to the stepping stones during the attack on the final target.

A specific type of stepping stone, often called a *zombie*, provides a "remote control" capability for an attacker and can severely obscure tracking. One common type of remote control attack is known as a distributed denial-of-service (DDoS) attack. During the preparation phase of a DDoS attack, the attacker compromises a large number of insecure machines (preferably having high bandwidth connections) and installs software on each of them that will later be used to target the ultimate victim during the attack phase. During that attack phase, the attacker sends a command to each of the zombies, commanding them to flood the victim with packets, overwhelming the victim's system with a distributed coordinated attack and putting it out of service. Alternatively, the attack software on the compromised machines (i.e., the zombies) may have a timer that triggers a coordinated attack on the victim. The small number of command packets used to trigger the attack are difficult to trace (and there are none to trace if the remote timer option is used), so if the attacker takes care not to arouse suspicion (or leave evidence, such as in log files) during the preparation phase, then determining the true source of the attack will be close to impossible.

Forged source addresses are also the key to a new and very potent type of DDoS attack, called a *reflector attack.* In the most potent version of this attack, packets containing the victim's IP address in the source address field are transmitted by the attacker to a huge number of servers and/or routers[17] situated throughout the Internet. The servers or routers see the packets as a request for service, and send (or "reflect") replies directly to the source address of the packet (i.e., directly to the victim). The overall effect is to overwhelm the victim's system with an enormous flood of packets, which are difficult to trace because the reflected packets are coming from so many different machines, each sending only a relatively small number of packets diffused throughout the Internet. An attacker that rotates the attack among some subset of a vast "inventory" of reflectors makes traceback even more difficult.

## 5.4 The current threat environment far exceeds the Internet's design parameters

As previously stated, the Internet's fundamental technology was developed to serve the needs of a benign community of researchers and educators. In the context of an environment that posed little or no threat, there was scant motivation for the development of significant security capabilities. Password protection was seen as sufficient to deter inappropriate behavior in a perceived threat environment consisting solely of occasional pranksters. The Internet's original design goals never included the support of today's mission-critical and high-stakes applications in business, energy, transportation, communications, banking and finance, and national defense. These high-stakes Internet applications pose enormously tempting targets of opportunity for criminals, terrorists, and hostile nations, and they overly stress the technology designed for a more innocent time. The migration of critical applications from expensive, closed, arcane, and proprietary technology to less costly, open, highly-distributed Internet applications, based on widely available COTS or

[17] An attacker can obtain a list of IP addresses for these machines by the use of scanning tools.

public-domain software, makes these applications much more accessible and more easily attacked by malicious adversaries.

In this high-threat, target-rich environment, the technical ability to reliably track and trace intruders (supported by international agreements on monitoring and information sharing, and by international law proscribing attacks and specifying sanctions and punishment) is an indispensable element for enabling the continued use of the Internet to provide many of the essential services that societies depend on.

## 5.5 The expertise of the average system administrator continues to decline

In the early days of the Internet, a relatively small number of systems were attached to the network, and these machines were typically administered by individuals possessing a reasonable amount of skill in configuring and maintaining at least a basic level of system security. A huge and growing percentage of the computer systems attached to the Internet today are operated by users with little or no security or system administration expertise (e.g., the vast majority of ordinary consumers who own a desktop PC or laptop). Such machines are easy prey for attackers, who can, in turn, use these machines as stepping stones to attack more lucrative and well-defended targets.[18] By moving from stepping stone to stepping stone on their way to the final target, attackers can obscure the true origin of the attack, making tracking and tracing the attacker an extremely difficult task. Just as a fire-safe building situated in a neighborhood of poorly maintained firetraps can be overwhelmed by a conflagration caused by its neighbors, the generally poor security of the majority of Internet systems makes even well-protected targets more vulnerable and severely hinders any attempt to track and trace an attack to its true origin.

## 5.6 Attacks often cross multiple administrative, jurisdictional, and national boundaries

Before the Internet, computer systems and networks were typically under full central administrative control. System administrators had complete control of all system and network hardware and software in the administrative domain of their organization, including full control of hosts, cabling, hardware add-ons, and the selection of system and application software that was loaded on their systems. Other than dial-in lines that traversed the telephone company's domain, the administrative domain of the organization encompassed the entire computer system or network being managed. Moreover, all of the users of the system were typically members (e.g., employees) of the organization that administered it. Therefore, they could be expected to closely adhere to that

[18] It is an open question as to what extent the owner of a system used as a stepping stone is liable for any downstream damages.

organization’s security policy, and they knew they could quickly be subject to sanctions or dismissal for misuse or abuse of the system.

Although most aspects of such centrally administered systems were not continuously monitored, the system administrator had both the technical ability and the organizational authority to monitor any part of the system at any time, should the need arise. Today, the Internet interconnects an international collection of countless administrative domains, each controlled by an organization or an individual, but no single organization or entity is in charge of the Internet as a whole. There is no central administrative control of the Internet and no global visibility—no single organization or entity has a complete picture of any activity that crosses its administrative boundaries and, in the absence of cooperation, no single organization or entity can monitor, track, or trace packets outside of its administrative domain.

During the execution of Internet applications, IP packets often cross multiple administrative, jurisdictional, and national boundaries. There are no universal technical standards or agreements for performing the monitoring and record keeping necessary to track and trace attacks. Moreover, there are no universal laws or agreements as to what constitutes a cyber-attack, and what punishments, economic sanctions, or liability should ensue. There are no universal international agreements for the monitoring, record keeping, and information sharing necessary to track and trace intruders. No existing privacy laws span the Internet as a whole. Existing international laws and agreements that might touch on these issues were not written for the Internet and need to be tested on cases involving Internet cyber-attacks.

## 5.7 High-speed traffic hinders tracking

Users demand, and are willing to pay for, high speed on the Internet (more precisely, high bandwidth—routing as many packets as possible as quickly as possible). Contemporary Internet applications, such as continuous digital video, require very high bandwidth, so routers are designed to push packets through as quickly as possible. Therefore, it is not considered practical for a router to do much in the way of processing packets other than routing them. The analysis or marking of any significant number of packets to augment our ability to track and trace attackers would slow a router’s throughput to below what is considered to be an acceptable level.

Very high bandwidth also means that packet information that might assist tracking or tracing can’t be stored for very long (e.g., minutes or hours, not days), because the large number of packets passing through the router per unit of time will quickly overwhelm a reasonably sized storage device. As a result, information that might be useful for tracking is very ephemeral—if a traceback is not done quickly, the information needed to carry it out will be gone. This suggests that statistical approaches in which a sampling of packets is saved for a longer period of time might bear fruit. However, a subtle attack (using a small number of packets) would fall below the radar of statistical approaches.

## 5.8 Tunnels impede tracking

IP tunneling (also known as encapsulation) is the wrapping (or hiding) of IP packets within the data portion of new IP packets, often involving the encryption of the original packets. Encrypted IP tunneling is used to create virtual private networks.

## 5.9 Hackers destroy logs and other audit data

Although the details are beyond the scope of this report, hackers routinely attempt to hide their tracks by destroying or modifying system logs or other audit data, and seeding misinformation. Attacker toolkits often help to automate the process.

## 5.10 Anonymizers protect privacy by impeding tracking

A number of commercial and non-profit sites offer "anonymizer" services, functioning as an intermediary between a user and the sites the user wants to connect to. The anonymizer service hides identifying information, in particular the IP address of the user. Such sites typically state in their terms of service that they will reveal user information in response to a request by law enforcement.[19]

Onion routing [Goldschlag 1999, Onion Routing] was a research project that produced a very sophisticated anonymizing technique, in which a large number of intermediate hosts cooperate to protect the anonymity of the user. In onion routing, packets are routed through multiple cooperating hosts, providing multiple layers of encapsulation and encryption (like the layers of an onion) to provide a very high level of protection against tracking and surveillance. The ultimate recipient of a packet sees only the IP address of the last intermediary. Onion routing can also defeat attempts by a small number of intermediate hosts to discover identifying information about the user.

## 5.11 The ability to link specific users to specific IP addresses is being lost

Tracking and tracing attack packets to the machine from which a sophisticated attack originated is indeed a daunting task, and much of Part II of this report will focus on possible solutions to the extremely difficult technical and policy problems. However, as stated earlier in this document, one primary purpose for tracking and tracing attacks is to deter future attacks by punishing or

[19] For example, an anonymizer might embed a token in a packet in order to allow for the later identification of the IP address of the user, if the packet was found to be involved in an attack. The information necessary to link the token to the user's IP address would retained at the anonymizer site for a limited amount of time.

sanctioning the individuals or entities that originated them. To accomplish this, a direct link must be drawn between the IP address of the machine that originated the attack and the individual or entity that set the attack in motion. However, several Internet trends are making it increasingly difficult to link IP addresses of machines to the entities or individuals who use them.

In the early days of the Internet, every connected machine was assigned a *static* (or relatively permanent) IP address. However, the 32-bit address field of the current IP protocol limits the number of possible addresses,[20] and the tremendous growth of the Internet threatens to eventually make IP addresses a scarce resource. Already, organizations consider IP addresses to be a limited resource and are using schemes, such as the *Dynamic Host Configuration Protocol (DHCP)*, to share a pool of IP addresses among their users' machines. For example, an ISP will typically assign a *dynamic* IP address to a dial-in user, from a pool of IP addresses owned by the ISP. The dynamic IP address will remain valid until the user's modem connection to the ISP is terminated. Upon initiating a new dial-in session, the user will most likely be assigned a different dynamic IP address. An attack that is traced to one of an ISP's dynamically assigned IP addresses can only be linked to an individual through the ISP's logs and record keeping and is dependent on the ISP's willingness to collect, preserve, and divulge that information. Many ISPs inform their customers (via "terms of use" agreements) that information about a customer's use of their service may be made available to law enforcement or other governmental authorities at the ISP's sole discretion.

By and large, ISPs have a strong incentive to cooperate, at least domestically, with law enforcement and governmental authority. Far more serious erosions in the ability to link an IP address to a particular user or entity are the appearance of "pay as you go" ISPs and the growth of (and ability to anonymously purchase) mobile wireless computing devices. "Pay as you go" ISPs, and in particular prepaid Internet access cards, allow one to purchase Internet access time without a monthly commitment or long-term contract and, unlike the traditional ISP arrangement, require little or nothing in terms of identification. A potential attacker who pays cash (or uses a stolen credit card) can achieve nearly complete anonymity in accessing the Internet.

Consider a fast-growing mobile digital technology called Global System for Mobile Communications (GSM), which integrates voice, high-speed data, fax, paging, and messaging. GSM mobile devices incorporate a Subscriber Information Module (SIM) "smart" card, which provides the authorization to use the network. An individual can go into a store, put down cash, and anonymously buy a wireless phone/computing device with a prepaid subscriber agreement (Subscriber Information Module). One of the services available under GSM is the General Packet Radio Service (GPRS). Here is a brief description from the GSM Association's web site:[21]

---

[20] Although this is a huge number (4.3 billion), addresses have historically been allocated in large contiguous blocks, so we are already seeing IP addresses being treated as a limited resource. Moreover, a large number of addresses have been reserved for special uses and cannot be assigned as globally routable IP addresses. A proposed new Internet Protocol (IPv6) will greatly expand the number of possible IP addresses.

[21] From <http://www.gsmworld.com/technology/gprs/intro.shtml>.

> *For the first time, GPRS fully enables Mobile Internet functionality by allowing interworking between the existing Internet and the new GPRS network. Any service that is used over the fixed Internet today—File Transfer Protocol (FTP), web browsing, chat, email, telnet—will be as available over the mobile network because of GPRS. In fact, many network operators are considering the opportunity to use GPRS to help become wireless Internet Service Providers in their own right…*
>
> *Because it uses the same protocols, the GPRS network can be viewed as a subnetwork of the Internet with GPRS capable mobile phones being viewed as mobile hosts. This means that each GPRS terminal can potentially have its own IP address and will be addressable as such.*

The tenuous link between a particular IP address and an individual or organization can be almost entirely obscured by using mobile devices and services (such as a Java-enabled PC/phone) that can be purchased anonymously to access the Internet. I use the term *realm switching* to refer to transferring information from one type of communications technology to another, such as moving packets from a mobile realm (e.g., GPRS) to the traditional Internet.[22] The ability to track, trace, and understand (e.g., correlate the different aspects of) an attack that crosses multiple realms is limited in the extreme.[23]

However, the technical ability to track and trace attacks across the Internet to the IP address of origin is still so difficult a task at present that today's attacker has little motivation to take additional steps, such as realm switching or using prepaid Internet access cards, to achieve further anonymity. As the basic technology to trace an attack to its IP address of origin improves, we would expect attackers to take these additional steps to thwart attempts to discover their identity.

## 5.12 Purely defensive approaches will fail, so deterrence through tracking and tracing is crucial

On the Internet, attackers have the same advantages with respect to asymmetrical warfare as terrorists do in the physical world. They can choose the time and place of their attacks, having only to discover and exploit one vulnerability, or at most a limited number of vulnerabilities. Security personnel must cover all the bases (i.e., eliminate or block all the vulnerabilities), leading to a severe inequity in the technical and economic resources required for defense versus those required for offense. Only the threat of retaliation, sanctions, or other punishment for the attacker can even the playing field and provide a deterrent effect. The ability to trace an attack and identify the attacker (or the sponsoring entity) is indispensable if the Internet is to be able to survive

[22] Other examples of realm switching can involve passing through a private branch exchange (PBX) telephone system, or an X.25 packet switched network.

[23] Realm switching can foil attempts at packet correlation temporally, spatially, and logically.

and thrive as the foundation for delivering the essential services needed to support a global information-based society.

# 6 Example: A “Smurf” IP Denial-of-Service Attack

A particularly good description of the technical details of a well-known type of denial-of-service attack (called “smurf”) can be found in CERT Advisory CA-1998-01 [CERT 1998]. Although the described attack method, and this advisory, is over four years old, it is still quite an effective and dangerous attack. It is based on an attacker’s ability to spoof (forge) IP source addresses. In the most abstract sense, the attacker sends a request for return packets to some intermediate network’s “broadcast address,” which in turn automatically relays the request to all the machines on that network. All the machines then reply with a return packet. However, in the original attack packet, the attacker replaces his or her true source address (i.e., the address all the machines will reply to) with the address of the intended victim. The victim is then flooded with replies from all the machines in the intermediate network. Hence, the intermediate network is not the true victim of the attack—it merely serves to amplify the effect of the original attack packets. Clearly, the attacker can send similarly forged packets to other intermediate networks at the same time to further amplify the attack and cause so much network congestion at the victim’s site that it will be impossible for the victim to perform any useful work or provide any useful services.[24]

There is little that victims can do on their own to deal with this attack once it is underway. However, there has been some progress in preventing this type of attack, since it is now considered “best practice” for ISPs (and other network administrators) to filter and discard outgoing packets containing source addresses that are not part of their own network’s allocated and assigned IP addresses and therefore could not be legitimate packets. It is also considered best practice for network administrators to filter and discard packets coming from outside their domain and destined for broadcast addresses inside their network. Logging and following up on the appearance of such packets can help in the process of tracing the source of malicious packets. However, the inability to authenticate the source address using the current Internet standard protocols remains a huge obstacle to tracking and tracing those who craft attack packets, design tools to automate the process, or simply use the automated tools to generate packets with a forged point of origin.

See the “smurf” advisory [CERT 1998] for further details. It is the inability of a recipient to authenticate a packet’s source address that is at the core of this particular attack strategy and of many denial-of-service attacks in general.

[24] A smurf attack is an example of an amplified “reflector” attack. Reflector attacks were described in Section 5.3.

# Part II: Promising Research Approaches: The Search for Near- and Long-Term Solutions

# 7 Overview

In the cyber world, the current state of the practice regarding the technical ability to track and trace Internet-based attacks is primitive at best. Sophisticated attacks can be almost impossible to trace to their true source using current practices. The anonymity enjoyed by today's cyber-attackers poses a grave threat to the global information society, the progress of an information-based international economy, and the advancement of global collaboration and cooperation in all areas of human endeavor.

Part II of this report examines the current best practices and the most promising research approaches in a search for near-term and long-term solutions to the traceback problem. However, it is clear that technical approaches alone can never offer a complete solution to the problem. Along with the proposed technical solutions, the policy implications and issues engendered by the technology are discussed.

First the most basic approaches are described, to serve as a baseline from which more advanced techniques can be judged. Next is an in-depth example of a tracking technique called *backscatter traceback*, arguably the best traceback technique used in practice today. After describing an innovative proposal for a near-term solution, called *CenterTrack*, I highlight some of the most promising long-term research approaches. All but one of these approaches are probabilistic (i.e., statistical) in nature—auditing only one out of every several thousand packets of Internet traffic. Only when the attack traffic is very large can the results be correlated into a successful traceback to the point of origin or to the entry point into an administrative network domain. As discussed in Part I of this report, one of the most devastating classes of attack on the Internet today is the distributed denial-of-service (DDoS) packet flood, for which these probabilistic traceback techniques have a useful role to play. The final research approach attempts to make feasible the traceback of single IP packets, which is one of the ultimate goals of security research. Finally, I discuss some of the global policy issues directly associated with the technical problem of tracking and tracing cyber-attacks, and propose some technical and policy requirements for next-generation Internet protocols to support tracking and tracing.

# 8 Basic Approaches

## 8.1 Hop-by-Hop IP Traceback

The most common and basic method in use today for tracking and tracing attacks is *hop-by-hop traceback.* This method is only suitable for tracing large, continuous packet flows that are currently in progress, such as those generated by ongoing denial-of-service (DoS) packet flood attacks. In a DoS flood attack, the source IP addresses are typically spoofed (i.e., they are forged addresses inserted into the source address field of a packet to disguise the true IP address of the machine that originated the packets), so tracing is required to find the true origin of the attack. For example, assume that the victim of a flood attack has just reported the attack to their ISP. First, an ISP administrator identifies the ISP's router that is closest to the victim's machine. Using the diagnostic, debugging, or logging features available on many routers, the administrator can characterize the nature of the traffic and determine the input (ingress) link on which the attack is arriving.[25] The administrator then moves on to the upstream router (i.e., the router one previous hop away that is carrying attack packets toward the victim).

The administrator repeats the diagnostic procedure on this upstream router, and continues to trace backwards, hop-by-hop, until the source of the attack is found inside the ISP's administrative domain of control (such as the IP address of a customer of the ISP) or, more likely, until the entry point of the attack into the ISP's network is identified. The entry point is typically an input link on a router that borders another provider's network. Once the entry point into the ISP's network is identified, the bordering provider carrying the attack traffic must be notified and asked to continue the hop-by-hop traceback. Often there is little or no economic incentive for such cooperation.

Traceback to the origin of an attack fails if cooperation is not provided at every hop, or if a router along the way lacks sufficient diagnostic capabilities or resources. If the attack stops before the trace is completed, the trace fails. Hop-by-hop traceback is a very labor-intensive, technical process,[26] and since attack packets often cross administrative, jurisdictional, and national boundaries, cooperation can be difficult to obtain. Partial traceback is still of value, since packet filters can be put in place to limit the DoS flood. How anomalous the attack packets are and how well they can

[25] The diagnostic capabilities and the specific procedures vary, depending on the router. For example, input diagnostic functions (such as logging) for Cisco routers are set using access control lists [Cisco 2001].

[26] Tools can be developed to automate the process somewhat, but these can be expected to be specific to a particular brand or type of router, and typically would not be applied across administrative boundaries.

be characterized determines how restrictive the filters have to be. Overly restrictive filters can contribute to the negative effects of a DoS attack.

In summary, hop-by-hop traceback can be considered to be the baseline from which all proposed improvements in tracking and tracing are judged. It is the most basic method for tracing large packet flows with spoofed source addresses, but has many limitations and drawbacks, as described above. For example, DDoS attacks are extremely difficult, if not impossible, to trace via this process, since there are multiple sources of attack packets, multiple paths through the Internet, and potentially a relatively small number of packets coming from each source.

## 8.2 Ingress Filtering

Much of the havoc inflicted on the Internet by attackers is accomplished using attack packets with spoofed source addresses. The occurrence of packets with spoofed source addresses, and their ability to transit the Internet, can be greatly limited through cooperative efforts by ISPs, using a basic packet filtering approach called *network ingress filtering*.

For example, assume that an ISP provides Internet connectivity to a customer network and assigns the customer a fixed set of IP addresses. Assume that the connectivity is provided via the ISP's router *R*. To limit IP source address spoofing, the ISP places an *ingress* (input) filter on the input link of router *R*, which carries packets from the customer network into the ISP's network and onto the Internet. The ingress filter is set to forward along all packets with source addresses that belong to the known set of IP addresses assigned to the customer network by the ISP, but the filter discards (and optionally logs as suspicious) all packets that contain source IP addresses that do not match the valid range of the customer's known IP addresses. Hence, packets with source addresses that could not have legitimately originated from within the customer network will be dropped at the entry point to the ISP's network.

The widespread use of ingress filtering by all service providers would greatly limit the ability of an attacker to generate attack packets utilizing a broad range of spoofed source addresses, making tracking and tracing the attacker a much easier task. Any attacker located within the customer network, in our example above, would either have to generate packets that carry the attacker's legitimate source address or (at worst) spoof a source address that lies within the set of IP addresses assigned to the customer network. So, even in the worst case, an attack originating within the customer network in our example can be traced to some machine in that customer network, simply by reading the source address on the attack packet. With the help of the administrator of the customer network, the search for the attacker can then proceed in a greatly narrowed search space.

Ingress filtering is described in detail in an Internet Engineering Task Force (IETF) document [Ferguson 2000]. It is one of the IETF documents that specify "Internet Best Current Practices for

the Internet Community." Ingress filtering can be supplemented by *egress filtering* (i.e., exit filtering) on the part of users, such as corporate network administrators, to ensure that packets exiting a network domain contain only valid source addresses that correspond to IP addresses contained within that domain.

## 8.3 Policy Implications

Ingress filtering has yet to be implemented on a very widespread basis. There is a broad element of "citizenship" involved, since many of the attack victims protected by a particular ISP's ingress filters would likely not be the ISP's customers. Also, ingress filtering can break certain network services, and ISPs would have to expend extra effort to get those services to work despite ingress filtering. Whether the liability for an attack extends to the ISP through which the attack originated is an open question, domestically and internationally. The existence of Internet best practices such as ingress filtering may eventually establish a level of "due care" that ISPs and other providers would be expected to fulfill.

# 9 Backscatter Traceback

This chapter gives an in-depth description of arguably the best traceback approach being used in practice.

*Backscatter traceback* [Morrow, Gemberling 2001] is a technique for tracing a flood of packets that are targeting the victim of a DDoS attack. (In this case, the term *backscatter* refers to the slew of "destination unreachable" error packets that are sent back toward the attack packets' source IP addresses as a result of this approach.[27] More details about this later.) This tracking approach was developed by Chris Morrow and Brian Gemberling of UUNET, a major ISP. The UUNET backscatter traceback technique relies entirely on the standard characteristics of existing Internet routing protocols, and although some special router configurations are used, there is no custom modification of protocols or equipment that is outside of current Internet standards.

In a typical DDoS attack, a victim's system is put out of service by a flood of malicious attack packets originating from a large number of zombie machines previously compromised by the attacker. The destination address field of each attack packet contains, of course, the IP address of the victim. The source IP address of each packet is typically spoofed. In contemporary DDoS attacks, the spoofed source address is typically chosen at random (by an attacker support tool) from the universe of all possible IP addresses. However, large blocks of the Internet address space have not yet been allocated for global routing,[28] so many of the randomly chosen spoofed addresses are not legitimate source addresses. Hence, all attack packets in a typical DDoS attack contain fake source IP addresses, which give no clue as to the true addresses of the zombie machines that are participating in the DDoS flood. Moreover, a large number of the attack packets in a typical DDoS flood carry invalid source addresses, which could not correspond to any machine on the Internet.

The UUNET backscatter traceback technique makes clever use of the large number of invalid source address that are characteristic of contemporary DDoS attacks. Here is how the technique works:

---

[27] In general, the term backscatter refers to any unsolicited response to the receipt of a packet (such as SYN/ACK packets sent to a spoofed source address). For the Morrow and Gemberling approach, only "destination unreachable" packets are of interest.

[28] The Internet Address Naming Authority (IANA) has reserved large blocks of the Internet address space for special use. These unallocated addresses cannot be used for global routing, but many of these addresses can be used locally (i.e., within an administrative domain).

1. The attack is reported to an ISP.

   The victim of a DDoS attack reports the problem to their ISP. The flood of attack packets has made the victim's Internet connection essentially unusable, putting the victim out of service.

2. The ISP configures all of its routers to reject all packets destined for the victim.

   The ISP uses a standard routing control protocol to quickly configure all of its routers to reject (i.e., filter) packets that are targeted to the victim.[29] In essence, the router configuration is set so that the victim's machine is now unreachable through the routing infrastructure. Note that by rejecting all packets that have the source address of the victim, benign packets carrying legitimate traffic will also be lost; however, by and large, the overwhelming number of packets heading for the victim will be attack packets. If the technique is successful, the total blockade of packets destined for the victim will only be in place for a very short period of time (a matter of minutes).

3. Rejected packets are "returned to sender."

   Well, that's roughly what happens, but here is a more precise explanation: When a router rejects a packet with the destination address of the victim, it sends an Internet Control Message Protocol (ICMP) "destination unreachable" error message packet back to the source IP address contained in the rejected packet.[30] Some of the "return to sender" ICMP error messages will be sent to legitimate users at legitimate addresses whose benign packets have been rejected along with the malicious ones. However, most of the packets destined for the victim are malicious attack packets.

   As shown in Figure 3, each ICMP "return to sender" error message packet contains, in its **source** IP address field, the address of the router (controlled and configured by the ISP) that blocked and rejected the packet heading for the victim. The router is also the machine that is generating the ICMP message. In its **destination** IP address field, the ICMP "return-to-sender" error message packet contains the source IP address found in the rejected packet that had been heading for the victim. These ICMP error packets are the "backscatter" or "noise" that enable the ISP to trace the attack packets back to their ingress point in the ISP's network.

---

[29] An example of this type of protocol is the Border Gateway Protocol (BGP), which can be used to set up a next hop of "null0" on Cisco routers.

[30] The ICMP error message packet also contains some portion of the original packet for identification purposes. The Internet Control Message Protocol, which is used to signal errors or to send other control information, is defined in an Internet Engineering Task Force (IETF) document: Request for Comments (RFC) 792 [Postel 1981].

| | FROM | TO | |
|---|---|---|---|
| | **Source IP Address:** Address of ISP router that rejected the packet heading for victim | **Destination IP Address:** IP address found in source address ("sender") field of rejected packet | |

*Figure 3: Address Information Contained in ICMP "Return to Sender" Packet*

4. The ISP configures all of its routers[31] to *blackhole* (that is, route for capture) many of the ICMP error packets (i.e., the "backscatter") with illegitimate destination IP addresses.

   As described earlier in this chapter, the Internet Address Naming Authority (IANA) has yet to allocate, for global routing, several very large blocks of IP addresses. One should never see any legitimate packet containing an IP source address from this unallocated address space entering a domain from an external network. However, network administrators can use these addresses for internal routing as long as the packet flow remains entirely within the administrative boundaries of their own network (i.e., fully within the control of their administrative domain). The next step in backscatter traceback is for an ISP to select a large range of IP addresses unallocated by IANA and to configure all of the ISP's routers to send packets destined for these invalid addresses to a specific machine (i.e., a "blackhole" machine) for analysis. Hence, the packets containing these invalid destination addresses have been "blackholed" by the ISP's routing infrastructure. The centermost region in Figure 4 represents the fraction of the overall packets arriving at an ISP's router that are blackholed for analysis. Since packets with these invalid destination addresses cannot have been routed into the ISP's network from an external source, these packets can only be some of the ICMP "destination unreachable" error message packets generated internally by the ISP's routers, which have been configured to reject all packets destined for the victim.

5. Analysis by the blackhole machine quickly traces the attack to one or more routers at the outermost boundary of the ISP's network.

   A human or program at the blackhole machine need only look at the source address of each ICMP error packet to determine the address of the router that sent it. Typically only a single router, or a small number of routers, will be identified as the ingress (i.e., entry) point of the attack into the ISP's network.

[31] It advertises a route for say 96.0.0.0/3 to a central blackhole system to collect the information. It restricts the route to only those routers in its administrative domain or serious problems will result.

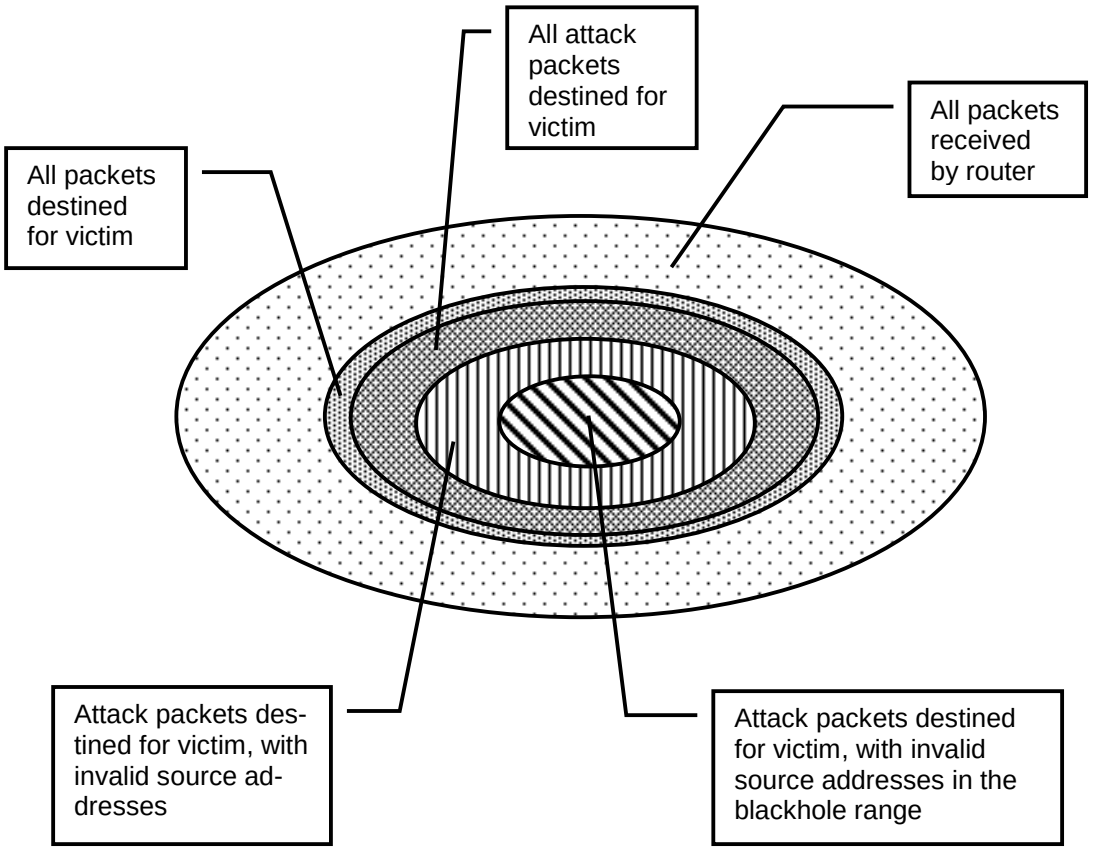


*Figure 4: DDoS Attack Packet Distribution as Seen by Each Router on Attack Path*

6. The ISP removes the filter blocking the victim's IP address from all routers except those serving as the entry points for the DDoS attack.

   The ISP leaves the blocking filter in place at those routers that have been traced as the entry points of the attack into the ISP's network and removes the blocking filter at *all* other routers. The DDoS attack remains blocked, but most of the flow of legitimate traffic to the victim is thereby restored. The entire backscatter traceback process can typically be executed in under a minute.[32]

   Only that portion of the inbound legitimate traffic that passes through the same entry points as the DDoS attack and is intended for the victim's IP address will remain blocked. Further analysis can identify specific characteristics of the attack packets (e.g., similarities in particular fields) that would allow the blocking filter on the attack entry point router(s) to be refined to be more permissive of benign traffic that has followed the same path as the attack packets. This will restore an even higher level of service to the victim.

[32] This assumes that all the preparations required to conduct a trace have been made prior to an attack.

7. The ISP asks neighboring ISPs, upstream of the attack, to continue the trace.

   The ISP further identifies the specific router interface(s) through which the attack is entering the ISP's network[33] and notifies the neighboring ISP(s) directly upstream of the entry points. The neighboring ISP(s) can continue to trace the attack closer to its ultimate source, using the backscatter traceback technique or any alternative tracking method.

## 9.1 Comments on the Backscatter Traceback Technique

The UUNET backscatter traceback technique is a fast and efficient method for countering contemporary DDoS flood attacks and for tracing such attacks to the boundaries of an administrative network domain. Tracking beyond an administrative domain depends on the trustworthiness, cooperation, and skill of upstream ISPs that lie between that domain and the source of the attack (which may merely be a collection of zombie machines, not the ultimate source—the machine of the attacker who is controlling and orchestrating the attack). This tracking technique is not a general-purpose traceback approach but rather depends heavily on the specific characteristics[34] of the DDoS attacks it was designed to defeat. Like many approaches designed to work against DDoS floods, it depends on a large number of attack packets being directed toward the victim. This is not an appropriate traceback mechanism for more subtle attacks. Eventual advances in the attack method can also defeat this approach, such as a DDoS attack tool that ensures that all of the randomly selected spoofed source addresses are valid IP addresses, allocated by IANA. DDoS attacks in which the zombie source addresses are not forged would also defeat this technique.

## 9.2 Policy Implications

The backscatter traceback method is a very good example of a class of traceback and remediation techniques that can be used to track attack packets to their entry point at the boundary of an administrative domain. One may assume that within a typical domain of administrative control an adequate level of trustworthiness among personnel and systems exists. As stated above, tracking beyond an administrative domain of control to the ultimate source of an attack depends on the trustworthiness, cooperation, and skill of upstream ISPs. Clearly, the traceback of a particular attack will likely cross multiple administrative, jurisdictional, and international boundaries. Having international agreements in place to facilitate the required cooperation will be an essential element of a traceback capability. This may include agreements to share traceback technology to raise the overall level of skill needed to complete a trace across multiple ISPs. However, any lack of trust among the entities participating in a given trace can be a major stumbling block to its ef-

[33] That is, the ISP identifies the *peering point*, the connection to a neighboring ISP, through which the attack has entered its network domain.

[34] This includes attack packets with source IP addresses randomly distributed throughout the Internet address space, including IP addresses that are invalid because they have not yet been allocated by IANA.

fective execution. One or more of the entities involved in the traceback may even share responsibility, liability, or merely tacit approval for the attack being traced and may therefore have sufficient incentive to be deceptive. Nothing in the technology of backscatter traceback provides authentication or any other trust mechanism that would improve the trustworthiness of information provided by multiple organizations ostensibly cooperating in a traceback.

# 10 An Overlay Network for IP Traceback

Routers are typically optimized for the high-speed forwarding of packets, not for analyzing the packets or tracing them back to their source. The most basic traceback capability of a router is the ability to report on which of the multiple inputs (i.e., interfaces) to the router a packet arrived, thereby identifying the neighboring router (from the previous hop) that delivered the current packet.[35] Routers on the highest speed links are often particularly lacking in the capabilities or resources to perform even the most basic traceback. This is a marketplace decision based on the desire for the greatest possible speed, and therefore the networks administered by ISPs are typically not well-suited for the effective tracking and tracing of cyber-attacks. Robert Stone has proposed an innovative approach that can improve traceability by adding an overlay network, optimized for hop-by-hop tracing and analysis, on top of an existing ISP network [Stone 2000].[36]

An *overlay network* is a supplemental or auxiliary network that is created when a collection of nodes from an existing network are joined together using new physical or logical connections to form a new physical or logical network on top of the existing one. New physical connections can be quite expensive, so new logical connections (such as IP tunnels) are more common.

The goal of Stone's approach, called *CenterTrack*, is to improve the traceability of the large packet flows associated with DoS flood attacks.[37] The first step in the CenterTrack approach is to create an overlay network, using IP tunnels to connect the edge routers in an ISP's network to special-purpose tracking routers that are optimized for analysis and tracking.[38] The overlay network is also designed to further simplify hop-by-hop tracing by having only a small number of hops between edge routers. In the event of a DoS flood attack, the ISP diverts the flow of attack packets (destined for a victim's machine) from the existing ISP network onto the overlay tracking network containing the special-purpose tracking routers. The attack packets can now be easily traced back, hop-by-hop, through the overlay network, from the edge router closest to the victim, back to the entry point of the packet flood into the ISP's network (i.e., the input connection or interface on the ingress router).

---

35 More precisely, this is the ability of a router to perform input diagnostics or logging to identify which of the router's input interfaces an attack packet has passed through (thereby identifying the previous hop).

36 At the time he wrote the CenterTrack paper, Robert Stone was with UUNET, the major ISP that developed the backscatter traceback technique described in the previous chapter.

37 Although it would be a simple matter to read the contents of the source address field of any of the DDoS flood packets, the source addresses of these attack packets are typically spoofed.

38 Edge routers are routers that have one or more connections to external routers, outside the ISP's administrative domain of control. Hence these routers are at the "edge" of an ISP's network.

## 10.1 Comments on the CenterTrack Method

The CenterTrack method is an interesting short-term research approach for tracing a large packet flow from the entry point in an administrative domain to its exit point. The approach requires no modification of existing Internet protocols. Stone states that probabilistic approaches to packet tracing (i.e., packet marking and ICMP traceback, described in the next chapter) are superior to CenterTrack, but since they require new protocol development, they are longer-term solutions.

Stone outlines some of the advantages and disadvantages of the CenterTrack approach [Stone 2000]. Advantages include the fact that special-purpose tracking and analysis features are only needed on the edge routers and the special-purpose tracking routers. The small number of hops between edge routers on the overlay-tracking network simplifies hop-by-hop tracking. Disadvantages include an increase in overall complexity, which can lead to operational errors (for example, in routing updates). Also, the overhead inherent in creating IP tunnels could amplify a DoS flood's negative effects on the network.[39] The CenterTrack technique does not work with attacks that originate inside an ISP's network,[40] and it is uncertain if this technique is highly scalable for distributed denial-of-service (DDoS) attacks with many entry points into the ISP's network. Stone notes that CenterTrack has been successfully tested in a lab setting, but the feasibility of this approach for real production environments needs further investigation.

The policy implications of this approach are identical to those of plain hop-by-hop traceback. Cooperation is needed at each backward hop for attacks with packet flows that cross administrative, jurisdictional, or national boundaries.

[39] IP tunnels are created by encapsulating packets within other packets, which consumes resources, such as bandwidth (due to the extra space needed for an encapsulating IP header) and additional computation. Stone points out that diverting a DoS flood onto the overlay tracking network, when the attack is targeting a backbone (i.e., very high speed) router, "would cause tunnel collapse or routing loops."

[40] Except that internal attacks that originate only one hop away from an edge router can be handled the same way as external attacks.

# 11 Probabilistic Approaches to Packet Tracing

For very practical reasons (the extremely high overhead), most of the promising research approaches to tracking and tracing attack packets do not attempt to audit all of the traffic passing through a given router. Instead, these research approaches use a probabilistic (i.e., statistical) scheme to sample the packet flow.

## 11.1 Generating Trace Packets (Using Control Messages)

One such approach is called *ICMP Traceback* or simply *iTrace* [Bellovin 2001]. A significant amount of computational resources are required for iTrace, and these resources could best be provided by building iTrace support into the router hardware. The fundamental concept is that about once in every 20,000 packets, a router sends an ICMP traceback message (called an *iTrace packet*) to the same destination address as the sampled packet (or to an outboard monitor). The destination (or monitor) collects and correlates the tracking information, and for large packet flows (as in DDoS attacks) there will be sufficient information to successfully trace the attack. However, an attacker can defeat or disrupt the trace by sending spoofed iTrace packets, so iTrace packets include an authentication field. Which authentication method to use is an open research issue, based on a trade-off between the need for good security (e.g., cryptographic strength) versus the need to minimize the computational resources required. There are also public key infrastructure (PKI) issues associated with who has the right to sign an iTrace packet (an open policy question) and how one validates that signature.

## 11.2 Packet Marking Schemes

One disadvantage of the iTrace scheme is that if one wants to trace a higher percentage of the packet flow, one must increase the probability that an iTrace packet will be emitted, and then the increase in trace packets will unfortunately increase the overall traffic. Other probabilistic tracing approaches, such as those described in “Practical Network Support for IP Traceback” [Savage 2000] and “Advanced and Authenticated Marking Schemes for IP Traceback” [Song 2001] also audit only a statistical sampling of the packet flow. However, instead of generating additional ICMP packets to carry tracking information, these schemes place tracking information into rarely used header fields inside the IP packets themselves. This approach is known as *probabilistic*

*packet marking* (PPM). The tracking information is collected and correlated at the destination of the packets, and, like the iTrace method, for a sufficiently large packet flow there will be enough tracking (path) information embedded in the packets to successfully complete the trace.

As with the iTrace approach, there is the real danger of an attacker tampering with, or spoofing, the tracking information. In response to this threat, Song and Perrig [Song 2001] enhanced the work of Savage [Savage 2000] by adding authentication to the embedded encodings of tracking information. However, it appears that the PKI issues raised by Bellovin for iTrace apply to the packet marking approaches as well.

## 11.3 Comment on Probabilistic Traceback Approaches

All of the probabilistic traceback approaches depend on auditing very sparse samples of large packet flows and thus are well suited for attacks that generate massive packet flows, such as DDoS floods. These approaches are entirely infeasible for tracking attacks that employ only a small number of packets.

In a recent paper, Marcel Waldvogel [Waldvogel 2002] shows how an attacker can effectively generate packets containing false tracking information and thereby defeat the Savage et al. PPM approach. The attack depends on the fact that the PPM tracking information must be split over several packets because of the very limited space within IPv4 headers for encoding the tracking information.

## 11.4 Policy Implications

The most significant policy implication of the probabilistic traceback approaches is the need for international agreements to formalize the necessary cooperation to make the techniques effective. The authentication of tracking information added by the Song and Perrig approach can enhance the trust necessary to make international cooperation effective. In some cases, technical and financial assistance to ISPs to support probabilistic traceback approaches will be required.

# 12 Single-Packet IP Traceback

All of the traceback methods discussed thus far have limited applicability because each of these techniques requires a large amount of attack traffic to support tracking. As devastating as DDoS flood attacks can be, Internet-based attacks that are equally or even more destructive can be far more subtle in terms of traffic flow. In the extreme case, an attack can be carried out with a single packet. Since a single IP packet is the smallest unit of Internet traffic, it has long been the aspiration of security researchers to be able to reliably track a single packet to its source.

In theory, one can address the problem of tracking traffic as small as a single packet by keeping a log at each router in the Internet infrastructure of every packet that passes through it, in a tamper-proof, fully-authenticated manner. When the victim of an Internet-based attack discovers the attack, the victim could then examine their own system logs to determine which packets participated. The victim would then need only to query the routing infrastructure ("Have you seen this packet?") to fully trace the path of the attack.

While appealing in theory, this approach is, of course, impossible in practice. A major technical problem is the need for extremely fast, virtually unlimited storage at each router to record a log of all packets passing through and to maintain that rapidly growing log indefinitely so that attacks not immediately discovered by a victim can still be traced at a later date. High-speed Internet backbone connections operating at OC-192 can transmit on the order of 1.25 gigabytes per second per router interface (i.e., per link).[41] At these speeds, each router can require 75 GB of storage per minute per interface, to record all the tracking information. Hence, a router with 16 links would require up to a massive 1,200 GB (1.2 TB – Terabytes = thousand billion bytes) of storage to keep only a minute's worth of traffic.[42] Even if each of the routers along an attack path had this much storage available, a victim would have to recognize the attack and complete the query in 60 seconds or less, before the relevant logs were overwritten by new traffic.

In addition to the technical infeasibility of this approach, there are very serious policy issues with respect to privacy. Privacy would be nearly impossible to maintain in an environment where complete traffic logs, including all packet contents, were kept at all routers in the Internet infrastructure. These logs would certainly be tempting targets for snooping attackers who could access

[41] OC stands for Optical Carrier, which specifies the speed of fiber optic networks that comply with the SONET (Synchronic Optical Network) standard. Common OC levels range from OC-1 (51.8 Mbps – megabits per second, where mega = million) to OC-192 (10 Gbps – gigabits per second, where giga = billion, which is approximately 1.25 GB – gigabytes per second, since a byte consists of 8 bits).

[42] Snoeren and his colleagues describe this in "Hash-Based IP Traceback" [Snoeren 2001].

massive amounts of stored Internet traffic by compromising a router's security and reading its logs. Trust would also be a primary issue—could you rely on the response to your query from routers outside your administrative domain or national boundaries?

A very promising new research approach called *Hash-Based IP Traceback* [Snoeren 2001], also known as *Single-Packet IP Traceback* [Snoeren 2002], offers the possibility of making the traceback of single IP packets feasible. The fundamental idea is to store highly compact representations of each packet rather than the full packets themselves. These compact representations are called "packet digests" and are created using mathematical functions called *hash functions* [Damgard 1990, Merkle 1990]. Note that the complete original packets cannot be restored from the packet digests, for all practical purposes.

Here's the definition of a hash function from the American National Standards Institute [ANSI 2001]:

> **hash function:** A mathematical function that maps values from a large (or very large) domain into a smaller range, and that reduces a potentially long message into a "message digest" or "hash value" or that is sufficiently compact to be input into a digital signature algorithm. *Note:* A "good" hash function is one that results from applying the function to a (large) set of values that are evenly (and randomly) distributed over the range.

Hash functions play a significant role in cryptography (such as the digital signature algorithm mentioned in the definition above). The only aspect of hash functions of importance for this traceback application, however, is the ability to create highly compact "digests" of long messages (e.g., packets) to greatly reduce the storage requirements at each router. In addition to the use of hash functions, Snoeren and his colleagues make use of a special data structure called a *Bloom filter*, which provides an additional major reduction in the storage requirements needed to uniquely identify each packet. The hash functions and Bloom filter reduce the storage requirement to 0.5% of the link capacity per unit of time, making single IP traceback technically feasible with respect to storage requirements. Moreover, this approach addresses the privacy issues posed by the universal logging of Internet traffic, since only the packet digests are stored at each router and not the actual packet contents. Note that, in general, a victim (or an intrusion detection system on behalf of a victim) submits a query by presenting the actual contents of the attack packet, and not the digest. However, for particularly sensitive cases, a victim will be able to submit a query without revealing the actual packet contents, at the cost of significant additional computational resources.[43]

An additional problem in single-packet traceback remains. During its transit through the Internet, a packet can be transformed in a number of ways, making it difficult to match transformed instances of a packet with the original packet. Common transformations include address translation,

[43] This is mathematically possible, but the engineering details to accomplish it have yet to be worked out.

fragmentation, and tunneling (such as when packets are encrypted to create a virtual private network as described in Part I of this report). Transformation information corresponding to the packet digests is stored in a *transformation lookup table,* which provides the information needed to track packets despite common transformations. The transformation information is retained by the router for the same amount of time as the packet digests are.

Hash-based IP traceback is accomplished using a system known as a *Source Path Isolation Engine (SPIE),* developed by Snoeren and his colleagues.[44] Hardware enhancements are necessary to support the hash-based traceback approach for very high speed routers, but these enhancements appear to be quite practical for current- or next-generation routers. Extensive simulations of this method have been run by Snoeren and his colleagues, and the results appear to validate the overall approach.

## 12.1 Comments on the Hash-Based Traceback Approach

This is arguably the most promising approach to packet tracking and tracing described in the research literature. It demonstrates the feasibility of tracking and tracing single packets, which has long been considered impractical.

A major drawback of this method is that the storage interval at each router is very short, in the order of a minute or a few minutes at best for high-speed routers. The high bandwidth and high traffic levels of today's Internet mean that hashes based on new traffic quickly fill storage and push out the old. The problem is most severe for routers in or near the high-speed core of the Internet. For routers in more "rural" lower speed locations, the traffic is significantly less, and packet hashes can be stored for a somewhat longer period of time. Snoeren and his colleagues have stated that probabilistic aging of packet digests (i.e., saving a sampling of packet digests for longer periods of time) can allow you a significantly longer window of time in which to trace attacks comprised of large packet flows. For example, if you decide to save every nth (say, $10{,}000^{th}$) packet for an extended period of time, then some attacks (i.e., those involving large packet flows) will still be traceable long after the storage interval for the complete set of packet digests would have expired.

The short-lived nature of stored packet hashes places a severe burden on the capability of the victim to quickly detect subtle attacks. The speed with which an attack can be detected is a crucial element in making this technique effective. With DDoS attacks, not only is there high-traffic volume (which makes traceback easier), but the detection of the attack is a "no-brainer"—there is no attempt by the attacker to disguise the attack, and the flood of traffic is obvious to the victim. For more subtle attacks, the burden is on the victim's intrusion detection system to quickly identify

[44] FreeBSD and Linux implementations of SPIE are publicly available at <http://www.ir.bbn.com/projects/SPIE>.

(say, in under a minute) a possible attack and to isolate the corresponding small number of packets so the routing/tracking infrastructure can be immediately queried to determine the path of the attack packets. In many cases, forensic analysis will reveal an attack long after the packet digest cache on most of the routers along the attack path has expired (i.e., the cache has been overwritten by new traffic). All this being said, the reduction of storage requirements by the hash-based traceback method to 0.5% of link capacity per unit of time is an impressive achievement that, combined with the SPIE architecture, demonstrates the feasibility of tracing single IP packets.

## 12.2 Policy Implications

As stated above, this method overcomes many of the privacy concerns associated with maintaining logs of Internet traffic, because it stores only packet digests and not the complete packet contents. Although communications among the components of a SPIE system can be authenticated, the query results produced by a SPIE component under the control of a malicious individual or organization could be falsified. Hence, one must rely on the trustworthiness of the ostensibly cooperating entities to reliably trace a packet to its source. A query result indicating an impossible or unlikely attack path can suggest the duplicity of one of the reporters.

Traceback is only possible within a single administrative domain of control, unless cooperation and collaboration in the legal, policy, and technical realms permit cross-domain traceback. The need for hardware enhancements and the significant amount of technical skill required to support this traceback technique suggest the value of providing financial and/or technical assistance as part of any international agreement involving cooperation and collaboration in establishing and using the traceback infrastructure.

# 13 Policy Considerations

Much of the effect of the proposed technical solutions will be blunted without policy considerations to guide and empower the designers, developers, and users of that technology. This chapter discusses some of the policy issues and considerations that are relevant to tracking and tracing cyber-attacks.

## 13.1 Intense International Cooperation is Essential

It is clear that tracking and tracing attackers across a borderless cyber-world, and holding them accountable, requires multilateral actions that transcend jurisdictions and national boundaries. Tracking and tracing requires cooperation encompassing the legal, political, technical, and economic realms. Significant international efforts to achieve effective, cooperative, and collaborative approaches for dealing with cybercrime and cyber-terrorism are currently underway. For example, the *G8 Recommendations on Transnational Crime*[45] [G8 2002] and the *Council of Europe Convention on Cybercrime*[46] [COE 2001] provide high-level frameworks for international cooperation and collaboration in the legal, policy, and technical realms.

One of the most significant policy implications of the technical approaches to tracking and tracing, described earlier in this report, is the need for intense international cooperation at a deeply technical level. This cooperation must go well beyond simple agreements in principle to share tracking data. The hacker community exchanges vulnerability information, detailed exploits, and attacker toolkits on a continual basis. As a result, defenders will be completely outmatched unless they routinely and extensively share highly technical information and resources, such as vulnerability information, incident data, new tracking methodologies and techniques, recommendations on hardware and software tools to support tracking, recommendations on best practices, and intelligence on the latest hacker capabilities and trends, including means of evading attempts to track and trace malicious activities. The development of interoperation standards for track and trace technology is an essential part of the required technical cooperation.

---

[45] Part IV, Section D of the *G8 Recommendations on Transnational Crime* covers "High-Tech and Computer-Related Crimes."

[46] As of this writing, the treaty has been signed by 34 nations, but has not yet been ratified by the 5 nations (at least 3 of which must be member States of the Council of Europe) required to enter the treaty into force.

A variety of organizational arrangements for exchanging highly technical information is possible. One of the most effective arrangements could be in the form of a shared multilateral technical organization, somewhat akin to the successful (domestic) Bellcore model. After the breakup of AT&T in 1984, the newly-formed regional Bell operating companies (also known as the "Baby Bells") created (and supplied ongoing funding for) a shared research organization called Bellcore. It was Bellcore's mission to research new telecommunications technologies for the benefit of all the regional Bells. Similarly, a shared multilateral technical organization would be far superior to ad-hoc arrangements for information exchange. In particular, due to the continual, rapid evolution of technology and attacker-defender capabilities, much of the technical information in the cyber realm has a very short "shelf life." A technical team needs stability and continuity to develop and maintain world-class expertise in this highly volatile environment.

Regardless of the precise organizational structure, a *multilateral technical research, engineering, and advisory capability*[47] is essential to (*a*) research and recommend the best tracking and tracing techniques and practices, (*b*) provide ongoing support for a multilateral tracking and tracing capability, (*c*) provide ongoing training and awareness for cooperating incident response and investigatory teams world-wide, (*d*) make recommendations to international engineering bodies, such as the Internet Engineering Task Force (IETF), for protocol improvements and standards creation in support of member states' requirements for tracking and tracing attackers, (*e*) interact with those creating cyber-law and policy to ensure that the technical and non-technical approaches complement and support each other, (*f*) help assure that the tracking and tracing infrastructures and technologies of cooperating entities can interoperate, and (*g*) assess the results of cooperation already undertaken by technical and law enforcement agencies, in order to provide feedback for continual improvement.

Although a multilateral technical organization could provide resources (such as tools, procedures, and expert advice) to support forensic analysis by incident response teams and investigatory groups, it is less likely that its services would be extended to provide global incident response (i.e., to serve as a multilateral incident response team, or "meta-CERT"). In contrast to providing technical advisory, engineering, and research services, a full-blown involvement in day-to-day incident response operations would be more likely to place the team at the focal point of conflicts of interest among member states.

[47] A successful example of intense international cooperation at the technical level is the Financial Action Task Force (FATF), established by the G-7 in 1989 to help deal with the problem of money laundering. Money laundering exploits the mismatch between laws based on national boundaries and international money flows that cross those boundaries. The FATF has assessed the effectiveness of existing cooperative efforts "to prevent the utilization of the banking system and financial institutions for the purpose of money laundering ..." and has conducted many technical studies to improve the knowledge and awareness of member nations and make multilateral cooperation more effective [Reinicke 1998].

## 13.2 Migrating Critical Applications to the Internet

Although promising, research on tracking and tracing cyber-attacks is in a nascent state. The lack of proven techniques for effectively and consistently tracking sophisticated cyber-attacks to their source (and rarely to the individuals or entities responsible) severely diminishes any deterrent effect. Perpetrators feel free to act with nearly total anonymity. Moreover, the general state of computer security is arguably worse than it was thirty years ago. In their paper, "Thirty Years Later: Lessons from the Multics Security Evaluation" [Karger 2002], Karger and Schell argue that the decades-old Multics operating system (which was used in a relatively benign closed environment) is more secure than most operating systems of today, which must function in a wide-open Internet environment.

> *Given the understanding of system vulnerabilities that existed nearly thirty years ago, today's "security enhanced" or "trusted" systems would not be considered suitable for processing even in the benign closed environment. Also, considering the extent of network interconnectivity and the amount of commercial off-the-shelf (COTS) and other software without pedigree (e.g., libraries mined from the Web), today's military and commercial environments would be considered very much "open." To make matters worse, everyone is now expected to be a system administrator of their own desktop system, but without any training!…Thus, systems that are weaker than Multics are considered for use in environments in excess of what even Multics could deliver without restructuring around a security kernel.*

The logical policy implication of poor security combined with an extremely limited track and trace capability is to move with extreme caution in migrating critical applications and infrastructures to the public Internet environment. This is particularly important in those cases where there are significant risks to public health and safety or risks of major economic loss or disruption of essential societal services.

Caution may be expressed in the form of delaying the migration of critical applications to the Internet, or by incorporating Internet technologies—but only for use on a private network. However, these approaches are often impractical, as the economic benefits of being on the Internet are typically essential for corporate survival. A more general approach is to design systems for survivability [Lipson 1999] to keep mission-critical services running despite attacks, accidents, or subsystem failures. This survivability engineering approach reduces risk by providing redundant and diverse means of delivering the essential services that support a system's mission despite the compromise of individual components by cyber-attack.

## 13.3 Privacy

Privacy is a central issue in all tracking and tracing technology. An important policy consideration in choosing among alternative technologies is to select a technology that not only minimizes the

disclosure of information under normal operation, but also does not create a target of opportunity for attackers. For example, technologies that concentrate lots of tracking information in a single location create the opportunity for massive disclosure (unless the information is highly sanitized).

Another critical part of the privacy equation is the *period of retention* for data that may be used for tracking evidence of a crime, evidence of a treaty violation, and so forth. International agreements on how long data must be retained and when it must be expunged are important components of a comprehensive policy of cooperation on tracking and tracing. Note that those formulating policy in this area must take into account the technical limitations on retaining large amounts of tracking data for very long, since the high bandwidth of the contemporary Internet can quickly overwhelm even large storage systems.

There has been considerable concern that the *Council of Europe Convention on Cybercrime* too broadly defines the tracking data to be collected and retained by ISPs, thereby threatening privacy [IWGDPT 2000]. Here is a situation in which the multilateral technical team described earlier in this chapter could help interact with policy and legal groups to implement a technical solution that collects, retains, and exchanges only data that is essential for tracking and tracing, thereby improving privacy. The particular solution will vary or evolve over time, but the selection of tracking methodologies that seek to lessen the privacy impact (such as retaining hashes of content instead of the full packet), could be one of the contributions of the experts on the multilateral technical team.

## 13.4 Incorporate Policy into Automated Tracking, Recovery, and Response Procedures

The increasing criticality and time-sensitive nature of today's highly distributed applications is continually increasing the need for rapid tracking, response, and recovery. The Council of Europe's Convention on Cybercrime calls for each State to establish a 24/7 point of contact to provide assistance during an incident. However, due to the increasingly severe timeliness requirements of many mission-critical Internet applications, it will soon be essential for many aspects of tracking, response, and recovery to be automated. As a result, many policy decisions will have to be incorporated into the automated procedures. For example, the amount of information to be shared in support of tracking may depend on the severity of the incident involved. Nations or other entities may agree to share a larger quantity of data, or more sensitive data, in the event of a severe incident. After an incident has occurred, it will be important to be able to audit the means by which the severity of the incident was determined, to verify that the sharing of additional data was justified.

## 13.5 Imprecise Jurisdictional Boundaries

The borderless cyber-world has created legal and political gray areas, where existing laws and agreements intended for the physical world (including the jurisdictional boundaries of law enforcement) are open to differing interpretations in the virtual world.

For example, in a well-publicized "sting operation" [Carter 2001], the FBI lured two Russian hackers to Seattle in November 2000 by offering them consulting jobs in a firm that appeared to be genuine but had only been created to serve as bait for the hackers. The FBI believed that the two Russians had stolen tens of thousands of credit card numbers through unauthorized computer accesses and had also attempted extortion by threatening businesses with the disclosure of sensitive information. As part of the "job interview," the Russians were asked to demonstrate their hacking talent. During the demonstration, the hackers accessed their machines in Russia, using workstations at the offices of the FBI's phony firm. Unbeknownst to the hackers, the workstations had been set up with keystroke monitoring software, so that the hackers login IDs and passwords could be captured. The FBI later used the captured IDs and passwords to download information from the Russian computers to use as evidence. Both men were arrested and indicted. One has since been convicted and the other is awaiting trial. The most recent and ironic twist to this story is that in August 2002, Russia's Federal Security Service charged one of the FBI agents with illegally accessing the hackers' computers, which were physically located in Chelyabinsk, Russia [Brunker 2002]. However, it is more likely that this particular case will be resolved in a political forum, rather than a legal one.

From the FBI's perspective, immediate access to the evidence on the Russian hackers' computers was crucial, since the electronic evidence could quickly be destroyed at any time. The FBI called this the first case in their history to "utilize the technique of extra-territorial seizure." From the Russian perspective, traditional national boundaries and legal jurisdiction need to be respected, and this case sets a troublesome precedent for future actions.

The very ephemeral nature of evidence in the cyber-world means that time is of the essence, much more so than in the physical world. In the physical world, law-enforcement cooperation across international boundaries depends on treaties and other international agreements that are often arduously slow in practice.[48]

One can envision international agreements that permit electronic "extra-territorial seizure" under explicitly specified circumstances, but with protections that help prevent misuse. Such protections may include placing the seized evidence in escrow, where it will remain locked away until permission (e.g., a search warrant) is granted for its use by the nation on whose territory the elec-

[48] For example, Mutual Legal Assistance Treaties (MLATs) allow you to serve a subpoena in a foreign country by making a request through the U.S. State Department. A six-month delay by a foreign country in honoring the request is not that unusual. Moreover, each such request expends a certain amount of "political capital," so requests are reserved for major cases only.

tronic seizure took place. If permission is denied, the international agreement may specify that the seized data must be destroyed.[49]

Thus, since cyber-attacks transcend our traditional notions of physical borders, our technical ability to track and trace attackers must be complemented by well-crafted international agreements that formalize the cooperation and collaboration necessary to track attacks (and collect evidence) across administrative, jurisdictional, and national boundaries. General agreements regarding cooperation and collaboration should include specifics such as the degree of information sharing, the speed with which a request for tracking information will be granted (which could vary according to attack severity), the equipment and procedures to be used, who will be responsible for the day-to-day fixed costs, and who will be responsible for the variable costs associated with a big incident. What financial and technical assistance should be provided to make the international tracking infrastructure more effective? What issues will arise if an attack packet crosses multiple international boundaries and is considered a criminal act in one country but not others? International agreements may specify cooperation in tracking only in the case of a criminal act. Who decides, and how can agreement be reached to resolve this prior to a major event?

Additional issues arise when one considers extra-territorial actions that extend beyond the collection of evidence in response to a cyber-attack. Extra-territorial defensive or retaliatory actions by the victim of an electronic attack might be justified by Article 51 of the United Nations charter, which confirms a nation's inherent right of self-defense. One could defend the right of a nation to take extra-territorial actions that interrupt a serious cyber-attack in progress, but it is less clear what responses are justified when an attack appears to be over. The appropriate response may depend on how severe the consequences of the cyber-attack were, the time scale over which multiple attacks occurred, and the severity of the threat of further attacks. Rather than leave extra-territorial action in response to a cyber-attack as a solely unilateral decision, international agreements are sorely needed to create a framework for acceptable behavior under a variety of complex circumstances.

## 13.6 Levels of Evidence

International agreements will be needed to specify what quality or level of forensic evidence (and the assurance that it was well preserved) will constitute an internationally recognized justification for sanctions, an information warfare response, or even a military response. Issues that come into play include the ability to demonstrate high assurance of the tamper-resistance of hardware and software systems that support tracking. A provable degree of certainty of correctness (i.e., accuracy, precision, and tamper-resistance) in tracking perpetrators is needed if serious responses to cyber-attack are to be justified.

[49] Open questions include: how is the integrity of the seized evidence assured, and, in the event that the evidence must be destroyed, how is it guaranteed that this occurs and no copies are made?

## 13.7 Levels of Damage and Threat

Once tracking of attacks becomes highly effective, it will be more important than ever to establish international agreements on liability (such as for intermediate systems that inadvertently help to amplify an attack), on defining levels of damage or threat (such as what constitutes an act of war in cyber-space), and on the appropriate response to each level of damage. Agreements in these areas are crucial, since prosecution, retaliatory information war, or even military retaliation can occur. Level of damage must be combined with the level of evidence to provide an internationally recognized justification for a response to cyber-attack.

## 13.8 Liability Issues

Here are some liability issues that could become the subject of international agreements once tracking and tracing makes it easier to identify entities and individuals who may bear some responsibility for the results of a cyber-attack:

- What are the liability exposures for the following?
  - the perpetrators of the attack
  - vendors of software whose vulnerabilities made the attack possible
  - owners and administrators of the intermediate (e.g., zombie or anonymizer) systems that participated in generating attack packets or obscuring the original source of the attack
  - transmitting ISPs and networks that did not squelch the attack when notified or did not help trace in accordance with international policy agreements
- Would there be waivers of certain kinds of liability exposure for those who participate in tracking and tracing? For example, a corporation may share computer security incident data with other organizations to aid in tracking and tracing attackers. However, by sharing this information the corporation may reveal information involving their own system vulnerabilities, which may provide evidence of liability, thereby increasing their risk of being successfully sued. As another example, if the system vulnerabilities are exposed, can an insurer avoid paying a business interruption claim to the corporation?
- Should liability extend to the owners of systems that participated in an attack without the owners' knowledge or permission? What are the standards of "due care" that system owners must comply with to avoid responsibility for use of their systems by attackers?
- Are those who provide anonymizer services liable for providing anonymity to attackers, even inadvertently? On the other hand, the privacy protection provided by anonymizers is a valuable social function. Should anonymizers receive special protection from liability claims?
- Should liability extend to system designers or owners who participate in placing an overly critical system on the Internet, despite knowledge of the fragility and poor security of the system, and without taking adequate backup measures (e.g., employing appropriate survivability strategies) in the event of attack? Placing highly critical, yet fragile, systems on the Internet could be an international destabilizing factor, and could allow mischief-makers to force international confrontations.

## 13.9 Honeypots and Honeynets

Honeypots and honeynets are decoy hosts and networks used to observe and document attacker behavior. Although the technical details are beyond the scope of this report, honeypots and honeynets are controlled environments that employ deception to masquerade as vulnerable and valuable targets of interest [Honeynet Project 2001, Spitzner 2002].[50] Unless an attacker goes beyond the confines of a honeynet or honeypot system, no real harm is done, although the actions taken by the attacker can be fully recorded for use as evidence of the attack.

Domestic laws and international agreements need to specify whether breaking into (or otherwise directing an attack against) a honeynet or honeypot would be sufficient cause for criminal prosecution (or a retaliatory response), or would instead be considered "entrapment."[51] At the very least, from the perspective of tracking and tracing attackers, international agreements should specify whether multilateral track and trace procedures should be triggered by attacks against decoy systems or whether they should be reserved solely for "real" attacks. It would be beneficial for such agreements to support multilateral cooperation for tracing attacks against decoy systems, even though many requests for assistance in tracking and tracing an attacker would likely be made without revealing that the victim of the attack was, in fact, a decoy system. The bottom line is that honeypots and honeynets can play a significant role in gaining early warning of, and thereby preempting, high-consequence cyber-attacks. Hence, there can be considerable justification for using a multilateral track and trace infrastructure to help identify perpetrators of attacks against decoy systems, and for having international agreements that support this.

[50] Honeypots and honeynets are useful as supplements to good Internet security. They should never be used as a replacement.

[51] Lance Spitzner, author of *Honeypots: Tracking Hackers* [Spitzner 2002], addressed the entrapment issue in a message posted to the honeypots@securityfocus.com mailing list: "Surprisingly, most legal professionals feel that even for law enforcement, honeypots are not an entrapment issue. The attacker was going to hack into some boxes either way, [so] your honeypot did not change the attackers behavior; at most it just changed his intended target." For the full message see: <http://online.securityfocus.com/archive/119/291711/2002-09-22/2002-09-28/0> (September 13, 2002). Chapter 15 of *Honeypots* discusses entrapment and other legal issues (e.g., liability and privacy) associated with honeypots, from the perspective of U.S. federal law.

# 14 Technical and Policy Requirements for Next-Generation Internet Protocols

## 14.1 Background

As described in Part I of this report, the standard Internet protocols in widespread use today were designed for a network environment in which the users of the network were considered to be trustworthy and the motivations and rewards for malicious activity were negligible. A high degree of anonymity for users was more the exception than the rule. As late as the early 1990s the "whois" databases maintained by the various Regional Internet Registries around the world painted a reasonably accurate picture of who was responsible for any allocated range of IP addresses, and provided administrative and technical contact information. These contacts could often be expected to be able to trace misuse of a given IP address in their range to an individual user or group (such as a university lab).

As the Internet population continued its exponential growth, anonymity increased.[52] Moreover, the kinds of applications that were placed on the Internet grew more critical in nature and were more tempting targets for malicious attackers with financial, political, or terrorist motives. Early on, traditional security approaches, such as firewalls, were seen as the solution to the problem of malicious users. However, fortress models of security were largely ineffective in protecting highly distributed applications. Such applications typically have a diverse and highly distributed user community, with varying degrees of trustworthiness (e.g., employees, collaborators, customers, suppliers, the general public, etc.) Fortress models are binary in nature: the attacker is outside and everything is fine, or the attacker is inside and all is lost. Highly distributed applications have indistinct physical and logical perimeters, so it difficult to distinguish inside from outside in a binary fashion. Rather, there are varying shades of gray, varying degrees of trust. In fact, in modern highly distributed systems, trust boundaries are essentially replacing the physical boundaries of earlier computer systems. Next-generation Internet protocols will be required to deal with trust not on a binary basis but at very fine levels of granularity.

On the Internet today, traditional computer security is being replaced by a new security paradigm called *survivability.* Survivability is the ability of a system to continue to fulfill its mission despite attacks, accidents, or subsystem failures. This approach is a blend of computer security and busi-

[52] For example, a shortage of IP addresses has led to the increased use of dynamic IP addresses instead of fixed IP addresses, and the use of network address translation (NAT) allowing multiple machines to share a single globally routable IP address.

ness risk management [Lipson 1999]. *Trust management* is a central part of survivability, and it is trust that has been gradually eroding throughout the Internet. Ultimately, trust is a human quality, based on individual and organizational attributes. It is particularly difficult to establish trust in an environment in which the users of a network have near-total anonymity (e.g., no history), and where IP source addresses can be readily spoofed. Can next-generation Internet protocols be designed to support the establishment of trust?

## 14.2 The "Entry-Point Anonymity" Problem

This report has described the extreme technical obstacles to tracing an attack packet to the IP address of its point of origin. Even when this difficult feat can be accomplished, it is becoming increasingly difficult, if not impossible, to link that IP address with the actual perpetrator of the attack. As an example, consider the emergence of prepaid Internet access cards, for which no personal identification of any kind is required by the ISP. The ISP assigns a dynamic IP address from its pool of IP addresses whenever the user dials in to a local or toll-free access number. The only clues to the identity of the user might be the phone number used to access the ISP and the hardware (MAC) address[53] of the computer that makes use of the dynamic IP address (if logged and retained). However, prepaid cellular phones (and public phones) can be untraceable,[54] and hardware addresses can be readily cloned (i.e., forged or spoofed).[55] The existence of cyber-cafés (wired and wireless), as well as public Internet access in libraries, makes it possible for a user to be totally anonymous, and untraceable, at his or her point of entry to the Internet.[56] I will use the term *entry-point anonymity* to refer the inability to link an Internet IP address to any human actor or organization. In response to an Internet attack, it is often a human actor or organization one wishes to track or trace, not solely an IP address. Moreover, an IP address is a poor surrogate on which to establish a basis of trustworthiness.

---

[53] The Media Access Control (MAC) address, is a unique physical hardware address encoded within a computer (e.g., on a network interface card) or other network device. The conversion between MAC addresses and IP addresses is performed by the Address Resolution Protocol (ARP), which maintains a correspondence table known as an ARP cache.

[54] Even if the individual who owns or is using a particular cellular phone is unknown, the general regional location of the call can still be traced, since cell towers are aware of the active cell phones in their vicinity. Phone-resident GPS would offer the ability to determine the precise location of a cell phone, but could be subject to tampering by malicious users.

[55] Not only can MAC addresses be cloned, but also the hardware addresses of cell phones. The IMEI (International Mobile station Equipment Identity) is the cell phone equivalent of a MAC address for a GSM phone, whereas the ESN (Electronic Serial Number) provides the hardware identity for many non-GSM cell phones.

[56] As another example, an attacker within range of an insecure wireless corporate LAN could launch attacks with near total anonymity from within the corporate site.

## 14.3 Vigilant Resource Consumption

Many denial-of-service attacks work by overwhelming the limited resources of a host or network device. In such attacks, the computational or other resource cost to the defender may be orders of magnitude higher than the cost to the attacker. Can next-generation protocols be designed so as to increase the cost to the attacker and decrease the cost to the defender?

Over the past few years, researchers have been investigating protocol design from the perspective of relative costs to the attacker versus the defender [Kent 1996, Meadows 1999, Meadows 2000]. The basic idea is to be very vigilant in the commitment of system resources to the processing of packets. One approach is to gradually ramp up the commitment of resources as trust in the validity of a packet is increasingly established. For example, if part of the processing of a packet involves the verification of the packet using cryptographically strong authentication, a protocol could be designed that first allows authentication using a weaker, less computationally expensive algorithm. If the packet fails this first step, then it can be rejected at a lower computational cost to the defender. If the packet succeeds, then authentication proceeds to the second step with a greater commitment of resources, but with a higher level of confidence that the resource allocation is justified. This is not a guaranteed defense against an intelligent adversary, but it would significantly increase the costs to an attacker who tried to overcome this defensive measure.

Supporting this kind of vigilant resource consumption[57] would be a most desirable capability of next-generation Internet security protocols. This risk-mitigating approach is much more in line with the new survivability paradigm than it is with traditional computer security. The only policy implication of this technical approach might be an attempt to ensure that this technical issue is placed in front of an appropriate standards-making body, such as the IETF.

Some very practical examples of vigilant resource consumption have been provided by D. J. Bernstein [Bernstein 1996] and Schuba et al. [Schuba 1997], who describe techniques for preventing a host from over-committing resources in response to “session opening” requests by packets with spoofed IP addresses. Their techniques help to defeat the classic *SYN Flooding* denial-of-service attack, which attempts to drain the resources of a victim.[58]

Another interesting aspect of the resource consumption issue relates to the amount of resources consumed by the track and trace infrastructure. Hackers can do more than just attempt to evade track and trace technology. Hackers can exploit an overly reactive track and trace infrastructure by taking actions that trigger the infrastructure to consume inordinate amounts of resources, thereby creating a denial-of-service scenario that could cripple not only the track and trace infrastructure, but also the systems that the infrastructure was designed to protect. This is analogous to

[57] I use the term *vigilant resource consumption* to mean frugal until sufficient trust is established.

[58] In this case, the limited resource being consumed is a kernel data structure that stores information on all pending network connections. This data structure can be filled with illegitimate incoming connection requests by an attacker, thereby denying legitimate incoming connection requests.

the effects of diseases caused by an overly reactive immune system. Designers of the track and trace infrastructure must take measures to lessen the likelihood that such scenarios could occur.

## 14.4 Trust Management and Privacy

In any Internet transaction, trust ultimately depends not on IP addresses but on particular relationships among individuals and their roles within organizations and groups (which may be economic, political, educational, or social). Trust cannot be established while maintaining total anonymity of the actors involved.

It goes without saying that there is a great need for privacy on the Internet, and it must be carefully guarded. For example, oppressive regimes may take draconian measures to stop activities that would be considered perfectly justified in a democratic society.[59] However, trust and privacy trade-offs are a normal part of human social, political, and economic interactions, and such trade-offs can be resolved in a number of venues, particularly in the marketplace. Consider the telephone system, in particular the caller ID feature, which displays the phone number and name associated with incoming calls. Customers can pay extra to have additional privacy (for an unlisted number) and sometimes are given the option to be anonymous by default (i.e., have their phone number and their name blocked on the displays of customers who pay for caller ID). However, the caller ID customer can choose to block all incoming calls from anonymous callers. The anonymous caller is notified of this fact by an automated message. The only way to complete the call is to enter a key sequence to remove the anonymity for this particular call and redial. This choice is a form of negotiation between the caller and the intended recipient of the call, and it is a trade-off between anonymity (privacy) and trust that is supported by the technology of caller ID and the marketplace.[60] There is no government mandate that all calls must be anonymous or that no calls are allowed to be anonymous. The individual caller chooses whether or not to relinquish anonymity (i.e., some degree of privacy) in exchange for the perceived value of completing the call by increasing the degree of trust as seen by the recipient.

One can envision next-generation Internet protocols supporting this kind of marketplace negotiation of trust versus privacy trade-offs. This would involve third-party certifying authorities, which would serve as brokers of trust. These “trust brokers” would be commercial, non-profit, or governmental (e.g., governmental certifying authorities may be required to establish the trust necessary for accessing national critical infrastructure and defense systems, but would not be involved in the vast majority of commercial, political, social, or educational activities).

---

[59] Of course, even in democratic societies, privacy is an essential element in maintaining freedom.

[60] Other privacy options are sometimes available. In one approach, the originator of the call is prompted to state his or her name, which is then played to the recipient of the call. The recipient can then choose to accept or reject the call. Note that this option keeps the phone number (and hence the location) of the caller private.

These certifying authorities would provide mechanisms whereby packets would be cryptographically signed with very fine-grained authentication credentials of the sender, such as the fact that the sender works for a particular company, works in a certain industry, or lives in a certain zip code. (The source IP address of the packets would also be authenticated.) This is not the same as having an individual digitally sign a message—a digitally signed message may be too coarse-grained for a particular scenario and may reveal too much.

As an example of the use of fine-grained authentication, the owners of a web site may decide to only allow visitors who are authenticated as residents of Pennsylvania. The individual accessing the web site may only wish to access web sites that permit full anonymity, or he or she may choose (via an appropriate user interface) which personal credentials may be revealed to obtain whatever value the web site may provide.

A more controversial capability would be the escrowing, by these certifying authorities, of complete identifying information for a specified period of time, to be revealed in the event that one or more of a user's packets have been identified as participating in a confirmed attack.

Note that there would be no centralized authority for fine-grained trust management. This trust-management approach would be based on a multitude of individual arrangements and marketplace concerns, though standards for implementing the approach would emerge over time. This approach allows for the individual, the marketplace, and policy negotiations to determine the trade-offs between trust and privacy. Protocols supporting this approach would not dictate the particular trade-offs to be made, which could vary over time according to the social, political, and legal climate. Privacy legislation would restrict the information that could be asked for and would set specific limits on retention and on how the information could be used.

## 14.5 "Situation-Sensitive" Security and Trust Processing

Survivability demands the ability to gracefully degrade service and increase security in the event of attack or accident. Next-generation Internet protocols must allow for variable levels of trust under various attack states. Protocol elements that are not processed during ordinary conditions may be called into service during ongoing periods of attack or heightened threat of attack. It is essential that next-generation protocols support this capability.

## 14.6 Sufficient Header Space for Tracking Information

One of the most basic requirements for next-generation Internet protocols, in support of tracking and tracing, is to provide sufficient header space to hold authenticated tracking and other audit information that may be inserted into a packet by all cooperating routers along the packet's path.

As described in Chapter 11, under the current Internet protocol (IPv4), packet-marking schemes must squeeze tracking data into rarely-used header fields. Because of the severe space limitations, tracking data must be split across multiple packets. This provides the opportunity for attackers to insert false tracking information into a stream of packets [Waldvogel 2002].

There are, of course, serious policy implications associated with the technical capability to gather large amounts of tracking data. Use of the technology must be limited by international agreements on policy and law stating under what circumstances tracking data can be collected, retained, and used. Such policy and law may be situation-sensitive. For example, under severe circumstances, multilateral agreements may allow more data (or more sensitive data) to be collected and retained for a longer period of time.

## 14.7 Emerging Next-Generation Security Protocols

### 14.7.1 IPsec

IPsec is an emerging security standard for IP. It provides two new security protocols: an Authentication Header (AH) and an Encapsulating Security Payload (ESP). The goals of this security standard are described in Request for Comments (RFC) 2401 [Kent 1998a]:

> *IPsec is designed to provide interoperable, high quality, cryptographically-based security for IPv4 and IPv6. The set of security services offered includes access control, connectionless integrity, data origin authentication, protection against replays (a form of partial sequence integrity), confidentiality (encryption), and limited traffic flow confidentiality. These services are provided at the IP layer, offering protection for IP and/or upper layer protocols.*
>
> *These objectives are met through the use of two traffic security protocols, the Authentication Header (AH) and the Encapsulating Security Payload (ESP), and through the use of cryptographic key management procedures and protocols. The set of IPsec protocols employed in any context, and the ways in which they are employed, will be determined by the security and system requirements of users, applications, and/or sites/organizations.*

IPsec provides for packet authentication and confidentiality, and in particular it can be used to authenticate the source address of a packet.[61] However, the overall approach is much more akin to traditional security than it is to the new survivability paradigm. As a result, it does not provide explicit support for vigilant resource consumption, fine-grained authentication of trust, and situa-

[61] The IPsec Authentication Header (AH) provides packet integrity by authenticating all IP header fields (except those that may legitimately change in transit) and the data portion of the packet. The Encapsulating Security Payload (ESP) provides packet confidentiality and integrity.

tion-sensitive processing, which are three of the requirements proposed in this chapter for next-generation Internet protocols.[62] So there is much work to be done in designing and implementing protocols that effectively fulfill these requirements, thereby preempting many potential attacks and enhancing the technical ability to track and trace attacks to their point of origin and to the responsible individuals or organizations.

### 14.7.2 Internet Protocol Version 6 (IPv6)

*Internet Protocol Version 6* (IPv6)[63] is the next generation IETF standard protocol that is slowly replacing the current IPv4. The new protocol's most salient feature is an enormously expanded address space compared to IPv4. An IP address is 128 bits long in IPv6, as compared to only 32 bits in IPv4. Another key feature is that this next generation protocol has IPsec security built in.[64]

In addition, the IPv6 header design is particularly flexible and efficient,[65] providing for a sequence of extension headers to carry optional information. In contrast to IPv4, the IPv6 header structure can provide ample space for recording tracking or other audit data. Recall that the packet marking techniques (described in Chapter 11) had to make use of very limited space in rarely used IPv4 header fields to store tracking data. The need to distribute the tracking data over several packets allows an attacker to defeat the packet marking schemes by generating packets containing false tracking information [Waldvogel 2002]. The flexible header structure of IPv6 potentially offers considerable space in which to store tracking data along the path of the packet. However, if the tracking information causes the packet to grow beyond the maximum allowable packet size,[66] packets would have to be fragmented and performance would suffer.

The growing shortage of IP addresses under IPv4 has resulted in an increased sharing of global IP addresses through the use of network address translation (NAT) and dynamic IP assignment. This sharing of IP addresses is a detriment to tracking and tracing, since there is no long-term correspondence between a computer (or other network device) and an IP address. The transient correspondence may be temporarily recorded in an ISP's logs or in a network administrator's configuration documents, but it is certainly ephemeral. In contrast to IPv4, the IPv6 the address space is enormous, making it possible for every network device to be assigned a static (fixed) IP address within a multi-level hierarchy of addresses. This would allow IP assignments to be more stable over time, which would be a significant advantage from the perspective of tracking and tracing the individual or entity responsible for an attack originating from a particular IP address.

---

62 Fulfilling these requirements is technically feasible using IPsec, but practical approaches and implementations that provide effective support for these requirements need exploration and development.

63 IPv6 is sometimes called *IP Next Generation* (IPng).

64 IPsec is a mandatory part of every IPv6 implementation, but its use is optional.

65 For example, one can specify that the optional information is to be processed only at the destination, saving processing time at all intermediate nodes (routers), or one can specify that all nodes along the path of the packet must process the header information.

66 The Maximum Transmission Unit (MTU) specifies the maximum packet size.

# 15 Conclusion

Society continues to migrate increasingly critical applications and infrastructures onto the Internet, despite severe shortcomings in computer and network security and serious deficiencies in the design of the Internet itself. Internet protocols were designed for an environment of trustworthy academic and government users, with applications that were oriented primarily towards research and information exchange. The consequences of occasional interruptions in service were low, and the motivations for attacks other than pranks were few. Today's Internet environment supports a global, less trustworthy user population, but provides a broad range of social, legal, economic, political, and infrastructural services, and hence offers far more motivation for malicious cyber-attacks.

In this era of open, highly distributed, complex systems, vulnerabilities abound and adequate security, using defensive measures alone, can never be guaranteed. As with all other aspects of crime and conflict, deterrence plays an essential role in protecting society. The ability to track and trace attackers is crucial, because in an environment of total anonymity, deterrence is impossible, and an attacker can endlessly experiment with countless attack strategies and techniques until success is achieved.

Hence, accountability for cyber-attacks that cause serious damage is essential. The ability to accurately and precisely assign responsibility for cyber-attacks to entities or individuals (or to interrupt attacks in progress) would allow society's legal, political, and economic mechanisms to work both domestically and internationally, to deter future attacks and motivate evolutionary improvements in relevant laws, treaties, policies, and engineering technology.[67]

Although ongoing research on tracking and tracing cyber-attacks is promising, existing track and trace capabilities are primitive compared with the capabilities of attackers. Subtle attacks by sophisticated attackers can be extremely difficult to detect and virtually impossible to trace using current technology. However, improvements to current Internet technology, including improved protocols, cannot succeed without an in-depth understanding and inclusion of policy issues to specify what information can be collected, shared, or retained, and how cooperation across administrative, jurisdictional, and national boundaries is to be accomplished. Nor can policy alone, with only high-level agreements in principle, create an effective tracking and tracing infrastructure that would support multilateral technical cooperation in the face of attacks rapidly propagat-

[67] A successful track and trace may reveal vulnerabilities that were exploited in an attack, thereby motivating improvements in engineering technology. The information gained can also help improve track and trace technology.

ing across the global Internet. To be of value, the engineering design of tracking and tracing technology must be informed by policy considerations, and policy formulations must be guided by what is technically feasible and practical. International efforts to track and trace cyber-attacks must be supported by intense technical cooperation and collaboration in the form of a multilateral research, engineering, and technical advisory group that can provide the in-depth technical skill and training to significantly improve the capabilities of incident response teams and law enforcement.

This report has described the problems and shortfalls in the current Internet environment, and has described some of the promising research approaches for tracking and tracing intruders. Several relevant global policy issues were also considered and some requirements were proposed for the design of next-generation Internet protocols to improve support for tracking and tracing attackers. Discussions that proposed policies, technical solutions, or protocol requirements were meant to raise relevant issues and demonstrate the interactions between policy and technology rather than to denote definitive solutions.

An effective track and trace capability is not a substitute for robust, well-engineered, secure, and survivable systems.[68] But it is an indispensable adjunct, providing accountability, redress, and deterrence. There is much to be done to create a global track and trace infrastructure that also protects privacy rights and civil freedoms. Above all there is a need for policy experts in the legal, economic, and political realms to work closely with technical experts to create the technical and policy framework to support the required cooperation and collaboration. Such an effort will be arduous and challenging, but it is essential if we are to protect the Internet-based global information society, and all of the human endeavors that are increasingly dependent upon it.

[68] Our greatest need is for secure and survivable systems that effectively resist, recognize, and recover from attacks, and which cannot easily be used as stepping stones to hide the tracks of an attacker or to amplify an attack on others.

# Bibliography


**[ANSI 2001]** American National Standards Institute, Standards Committee T1 Telecommunications. *ANS T1.523-2001, Telecom Glossary 2000,* Alliance for Telecommunications Industry Solutions. <http://www.atis.org/tg2k/> (2001).

**[Bellovin 2001]** Bellovin, Steve; Leech, Marcus; & Taylor, Tom. *ICMP Traceback Messages*. Internet Engineering Task Force (IETF) Draft—Work in Progress. <ftp://ftp.ietf.org/internet-drafts/draft-ietf-itrace-01.txt> (October 2001).

**[Bernstein]** Bernstein, D. J. *SYN Cookies.* <http://cr.yp.to/syncookies.html>.

**[Bernstein 1996]** Bernstein, D. J. *Re: SYN Flooding [info]*. Newsgroup Posting: alt.security, comp.security.unix, comp.security.misc, comp.security, comp.protocols.tcp-ip. <http://cr.yp.to/syncookies/idea> (September 16, 1996).

**[Brunker 2002]** Brunker, Mike. "FBI Agent Charged With Hacking." *MSNBC.com Technology & Science.* <http://msnbc.com/news/563379.asp?cp1=1> (August 15, 2002).

**[Burch 2000]** Burch, Hal & Cheswick, Bill. (2000). "Tracing Anonymous Packets to Their Approximate Source." *USENIX LISA 2000.*

**[Carter 2001]** Carter, Mike. "E-sting nets 2 Russian hackers; FBI alleges pair stole credit info." *Seattle Times*, April 24, 2001.

**[CERT 1998]** CERT Coordination Center. *CERT Advisory CA-1998-01: Smurf IP Denial-of-Service Attacks.* <http://www.cert.org/advisories /CA-1998-01.html> (January 5, 1998; last revised March 13, 2000).

**[CERT 2000]** CERT Coordination Center. *Overview of Attack Trends.* <http://www.cert.org/archive/pdf/attack_trends.pdf> (February 2002).

**[Cisco 2001]** Cisco Systems. *Characterizing and Tracing Packet Floods Using Cisco Routers.* <http://www.cisco.com/warp/public/707/22.pdf> (April 9, 2001).

**[COE 2001]** Council of Europe. *Convention on Cybercrime*. <http://conventions.coe.int/Treaty/EN/CadreListeTraites.htm> (see ETS No. 185) (November 23, 2001).

**[Cohen]** Cohen, Donald; Narayanaswamy, K.; & Cohen, Fred. *Changes To IP To Eliminate Source Forgery.* <http://www.cs3-inc.com/sf.html> (2002).

**[Damgard 1990]** Damgård, I. "A Design Principle for Hash Functions." *Advances in Cryptology - Crypto '89*. Heidelberg: Springer-Verlag, 1990, 416–427.

**[Dean 2002]** Dean, Drew; Franklin, Matt; & Stubblefield, Adam. "An Algebraic Approach to IP Traceback." *ACM Transactions on Information and System Security (TISSEC) 5*, 2 (May 2002): 119–137.

**[Deering 1998]** Deering, S. & Hinden, R. *RFC 2460: Internet Protocol, Version 6 (IPv6) Specification*. Network Working Group, IETF. <http://www.ietf.org/rfc/rfc2460.txt> (December 1998).

**[Ferguson 2000]** Ferguson, P. & Senie, D. *RFC 2827: Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing*. Network Working Group, IETF. <http://www.ietf.org/rfc/rfc2827.txt> (May 2000).

**[G8 2002]** *G8 Recommendations on Transnational Crime.* Part IV: Transnational Crime, Section D: High-Tech and Computer-Related Crimes. Endorsed by G8 Ministers of Justice and the Interior. Mont-Tremblant, May 13–14, 2002. <http://www.g8j-i.ca/english/doc1.html> (2002).

**[Gemberling 2001]** Gemberling, Brian W.; Morrow, Christopher L.; & Greene, Barry R. *ISP Security – Real World Techniques*. Presentation, NANOG, October 2001. <http://www.nanog.org/mtg-0110/greene.html>.

**[Goldschlag 1999]** Goldschlag, David M.; Reed, Michael G..; & Syverson, Paul F. "Onion Routing for Anonymous and Private Internet Connections." *Communications of the ACM 42*, 2 (February 1999).

**[Honeynet Project]** The Honeynet Project. <http://project.honeynet.org/>.

**[Honeynet Project 2001]** The Honeynet Project (Editor). *Know Your Enemy: Revealing the Security Tools, Tactics, and Motives of the Blackhat Community*. Boston: Addison-Wesley, 2001.

**[Houle 2001]** Houle, Kevin J.; Weaver, George M.; Long, Neil; & Thomas, Rob. *Trends In Denial of Service Attack Technology*, CERT Coordination Center. <http://www.cert.org/archive/pdf/DoS_trends.pdf> (October 2001).

**[IWGDPT 2000]** International Working Group on Data Protection in Telecommunications. *Common Position on Data Protection Aspects in the Draft Convention on Cyber-Crime of the Council of Europe.* Adopted at the 28th Meeting of the Working Group on September 13-14, 2000 in Berlin. <http://www.datenschutz-berlin.de/doc/int/iwgdpt /cy_en.htm>.

**[Karger 2002]** Karger, Paul & Schell, Roger. *Thirty Years Later: Lessons from the Multics Security Evaluation*. Research Report RC22534 (W0207-134), IBM Research, July 31, 2002. Report contains two papers that will appear in *Proceedings of the 18th Annual Computer Security Applications Conference*, Las Vegas, Nevada, December 9-13, 2002, Applied Computer Security Associates, 2002. <http://www.acsac.org>.

**[Kent 1996]** Kent, S. T.; Ellis, D.; Helinek, P.; Sirois, K.; & Yuan, N. "Internet Routing Infrastructure Security Countermeasures." *BBN Report 8173*. BBN, January 1996.

**[Kent 1998a]** Kent, S. & Atkinson, R. *RFC 2401: Security Architecture for the Internet Protocol*. Network Working Group, IETF. <http://www.ietf.org/rfc/rfc2401.txt> (November 1998).

**[Kent 1998b]** Kent, S. & Atkinson, R. *RFC 2402: IP Authentication Header*. Network Working Group, IETF. <http://www.ietf.org/rfc/rfc2402.txt> (November 1998).

**[Kent 1998c]** Kent, S. & Atkinson, R. *RFC 2406: IP Encapsulating Security Payload (ESP)*. Network Working Group, IETF. <http://www.ietf.org/rfc/rfc2406.txt> (November 1998).

**[Lee 2001]** Lee, H. & Park, K. "On the Effectiveness of Probabilistic Packet Marking for IP Traceback Under Denial of Service Attack," 338–347. *Proceedings of IEEE INFOCOM 2001*. Anchorage, Alaska, April 22–26, 2001. New York: IEEE Computer Society Press, 2001. <http://www.ieee-infocom.org/2001/paper/721.ps> (2002).

**[Lipson 1999]** Lipson, Howard & Fisher, David. "Survivability—A New Technical and Business Perspective on Security," 33–39. *Proceedings of the 1999 New Security Paradigms Workshop*. Caledon Hills, Ontario, Canada, Sept. 22–24, 1999. New York: Association for Computing Machinery, 2000. <http://www.cert.org/research>.

**[Meadows 1999]** Meadows, Catherine. "A Formal Framework and Evaluation Method for Network Denial of Service," 4–13. *Proceedings of the 12th IEEE Computer Security Foundations Workshop*. Mordano, Italy, June 28–30, 1999. New York: IEEE Computer Society Press, 1999.

**[Meadows 2000]** Meadows, Catherine. "A Framework for Denial of Service Analysis." *Proceedings of the 2000 IEEE Information Survivability Workshop*. Boston, Massachusetts, October 24–26, 2000. New York: IEEE Computer Society Press, 2000. <http://www.cert.org/research/isw/isw2000/papers/37.pdf>.

**[Merkle 1990]** Merkle, R. C. "One Way Hash Functions and DES." *Advances in Cryptology - Crypto '89*. Heidelberg: Springer-Verlag, 1990, 428–446.

**[Mirkovic 2002]** Mirkovic, Jelena; Martin, Janice; & Reiher, Peter. *A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms* (Technical Report 020018). Los Angeles: UCLA Computer Science Department, 2002.

**[Morrow]** Morrow, Chris. *BlackHole Route Server and Tracking Traffic on an IP Network*. UUNET, WorldCom, Inc. <http://www.secsup.org/Tracking/>.

**[Onion Routing]** Onion Routing Research Project. <http://www.onion-router.net>.

**[Postel 1981]** Postel, J. *RFC 792: Internet Control Message Protocol*. Network Working Group, IETF. <http://www.ietf.org/rfc/rfc0792.txt> (September 1981).

**[Reinicke 1998]** Reinicke, Wolfgang H. *Global Public Policy: Governing Without Government?* Washington, D.C.: Brookings Institution Press. <http://brookings.nap.edu/books/0815773900/html/index.html> (1998): 157-172.

**[Savage 2000]** Savage, Stefan; Wetherall, David; Karlin, Anna R.; & Anderson, Tom. "Practical Network Support for IP Traceback," 295–306. *Proceedings of ACM SIGCOMM 2000*. Stockholm, Sweden, Aug. 28–Sept. 1, 2000. New York: Association for Computing Machinery, 2000. <http://www.acm.org/sigcomm/sigcomm2000/conf/abstract/8-4.htm> (2000).

**[Savage 2001]** Savage, Stefan; Wetherall, David; Karlin, Anna; & Anderson, Tom. "Network Support for IP Traceback." *IEEE/ACM Transactions on Networking 9*, 3 (June 2001): 226–239.

**[Schuba 1997]** Schuba, Christoph L.; Krsul, Ivan V.; Kuhn, Markus G.; Spafford, Eugene H.; Sundaram, Aurobindo; & Zamboni, Diego. "Analysis of a Denial of Service Attack on TCP," 208–223. *Proceedings of the 1997 IEEE Symposium on Security and Privacy*. Oakland, California, May 4–7, 1997. New York: IEEE Computer Society Press, 1997.

**[Snoeren 2001]** Snoeren, Alex; Partridge, Craig; Sanchez, Luis; Jones, Christine; Tchakountio, Fabrice; Kent, Stephen; & Strayer, Timothy. "Hash-Based IP Traceback," 3–14. *Proceedings of ACM SIGCOMM 2001*. San Diego, California, Aug. 27–31, 2001. New York: Association for Computing Machinery, 2001. <http://www.acm.org/sigcomm/sigcomm2001/p1.html> (2001).

**[Snoeren 2002]** Snoeren, Alex C.; Partridge, Craig; Sanchez, Luis A.; Jones, Christine E.; Tchakountio, Fabrice; Schwartz, Beverly; Kent, Stephen T.; & Strayer, W. Timothy. "Single-Packet IP Traceback." *IEEE/ACM Transactions on Networking 10*, 6 (December 2002), *to appear*.

**[Song 2001]** Song, Dawn & Perrig, Adrian. "Advanced and Authenticated Marking Schemes for IP Traceback," 878–886. *IEEE Infocom 2001*. <http://www.ieee-infocom.org/2001/> (2001).

**[Spitzner 2002]** Spitzner, Lance. *Honeypots: Tracking Hackers*. Boston: Addison Wesley, 2002.

**[Stone 2000]** Stone, Robert. "CenterTrack: An IP Overlay Network for Tracking DoS Floods," 199–212. *9th USENIX Security Symposium*. Denver, Colorado, August 14-17, 2000. <http://www.usenix.org/publications/library/proceedings/sec2000/full_papers/stone/stone.pdf> (2000).

**[Waldvogel 2002]** Waldvogel, Marcel. *GOSSIB vs. IP Traceback Rumors*. Research Report RZ3424 (#93671), IBM Research, Zurich Research Laboratory, June 17, 2002. Paper will appear in *Proceedings of the 18th Annual Computer Security Applications Conference*, Las Vegas, Nevada, December 9-13, 2002, Applied Computer Security Associates, 2002. <http://www.acsac.org>.

**[Zhang 2000]** Zhang, Yin & Paxson, Vern. "Detecting Stepping Stones," 171–183. *9th USENIX Security Symposium*. Denver, Colorado, August 14-17, 2000. <http://www.usenix.org/publications/library/proceedings/sec2000/full_papers/zhangstepping/zhangstepping.pdf> (2000).

# REPORT DOCUMENTATION PAGE

*Form Approved*
*OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

| 1. AGENCY USE ONLY (LEAVE BLANK) | 2. REPORT DATE | 3. REPORT TYPE AND DATES COVERED |
|---|---|---|
| | November 2002 | Final |

4. TITLE AND SUBTITLE
Tracking and Tracing Cyber-Attacks: Technical Challenges and Global Policy Issues

5. FUNDING NUMBERS
F19628-00-C-0003

6. AUTHOR(S)
Howard F. Lipson, Ph.D.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

8. PERFORMING ORGANIZATION REPORT NUMBER
CMU/SEI-2002-SR-009

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Department of State
2201 C Street NW
Washington, D.C. 20520

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
This special report was sponsored by, and written for, the U.S. Department of State.

12A DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified/Unlimited, DTIC, NTIS

12B DISTRIBUTION CODE

13. ABSTRACT (MAXIMUM 200 WORDS)

In the cyber world, the current state of the practice regarding the technical ability to track and trace Internet-based attacks is primitive at best. Sophisticated attacks can be almost impossible to trace to their true source using current practices. The anonymity enjoyed by today's cyber-attackers poses a grave threat to the global information society, the progress of an information-based international economy, and the advancement of global collaboration and cooperation in all areas of human endeavor.

Part I of this report examines the current state of the Internet environment and the reasons why tracking and tracing cyber-attackers is so difficult. Part II examines some promising research on technical approaches that may greatly improve the ability to track and trace cyber-attacks to their source. Also discussed are some policy considerations with regard to privacy, information sharing, liability, and other policy issues that would be faced by the U. S. State Department in negotiating international agreements for cooperation and collaboration in the tracking and tracing of cyber-attacks. The report concludes with a closer look at technical and policy considerations for next-generation Internet protocols to enhance track and trace capabilities.

14. SUBJECT TERMS
cyber-attack, Internet security, network security, computer security, denial-of-service attack, traceback, packet tracing, packet tracking, international policy, cybercrime

15. NUMBER OF PAGES
84

16. PRICE CODE

| 17. SECURITY CLASSIFICATION OF REPORT | 18. SECURITY CLASSIFICATION OF THIS PAGE | 19. SECURITY CLASSIFICATION OF ABSTRACT | 20. LIMITATION OF ABSTRACT |
|---|---|---|---|
| Unclassified | Unclassified | Unclassified | UL |