PLANNING FOR THE FUTURE OF CYBER ATTACK ATTRIBUTION

HEARING

BEFORE THE

SUBCOMMITTEE ON TECHNOLOGY AND INNOVATION

COMMITTEE ON SCIENCE AND TECHNOLOGY

HOUSE OF REPRESENTATIVES

ONE HUNDRED ELEVENTH CONGRESS

SECOND SESSION

JULY 15, 2010

Serial No. 111–105

Printed for the use of the Committee on Science and Technology

Available via the World Wide Web: http://www.science.house.gov

U.S. GOVERNMENT PRINTING OFFICE
57–603PDF WASHINGTON : 2010

For sale by the Superintendent of Documents, U.S. Government Printing Office
Internet: bookstore.gpo.gov Phone: toll free (866) 512–1800; DC area (202) 512–1800
Fax: (202) 512–2104 Mail: Stop IDCC, Washington, DC 20402–0001

COMMITTEE ON SCIENCE AND TECHNOLOGY

HON. BART GORDON, Tennessee, *Chair*

JERRY F. COSTELLO, Illinois
EDDIE BERNICE JOHNSON, Texas
LYNN C. WOOLSEY, California
DAVID WU, Oregon
BRIAN BAIRD, Washington
BRAD MILLER, North Carolina
DANIEL LIPINSKI, Illinois
GABRIELLE GIFFORDS, Arizona
DONNA F. EDWARDS, Maryland
MARCIA L. FUDGE, Ohio
BEN R. LUJÁN, New Mexico
PAUL D. TONKO, New York
STEVEN R. ROTHMAN, New Jersey
JIM MATHESON, Utah
LINCOLN DAVIS, Tennessee
BEN CHANDLER, Kentucky
RUSS CARNAHAN, Missouri
BARON P. HILL, Indiana
HARRY E. MITCHELL, Arizona
CHARLES A. WILSON, Ohio
KATHLEEN DAHLKEMPER, Pennsylvania
ALAN GRAYSON, Florida
SUZANNE M. KOSMAS, Florida
GARY C. PETERS, Michigan
JOHN GARAMENDI, California
VACANCY

RALPH M. HALL, Texas
F. JAMES SENSENBRENNER JR., Wisconsin
LAMAR S. SMITH, Texas
DANA ROHRABACHER, California
ROSCOE G. BARTLETT, Maryland
VERNON J. EHLERS, Michigan
FRANK D. LUCAS, Oklahoma
JUDY BIGGERT, Illinois
W. TODD AKIN, Missouri
RANDY NEUGEBAUER, Texas
BOB INGLIS, South Carolina
MICHAEL T. McCAUL, Texas
MARIO DIAZ-BALART, Florida
BRIAN P. BILBRAY, California
ADRIAN SMITH, Nebraska
PAUL C. BROUN, Georgia
PETE OLSON, Texas

SUBCOMMITTEE ON TECHNOLOGY AND INNOVATION

HON. DAVID WU, Oregon, *Chair*

DONNA F. EDWARDS, Maryland
BEN R. LUJÁN, New Mexico
PAUL D. TONKO, New York
HARRY E. MITCHELL, Arizona
GARY C. PETERS, Michigan
JOHN GARAMENDI, California
BART GORDON, Tennessee

ADRIAN SMITH, Nebraska
JUDY BIGGERT, Illinois
W. TODD AKIN, Missouri
PAUL C. BROUN, Georgia
RALPH M. HALL, Texas

HILARY CAIN *Subcommittee Staff Director*
MEGHAN HOUSEWRIGHT *Democratic Professional Staff Member*
TRAVIS HITE *Democratic Professional Staff Member*
MELÉ WILLIAMS *Republican Professional Staff Member*
VICTORIA JOHNSTON *Research Assistant*

CONTENTS

July 15, 2010

Opening Statements

Witnesses:

Appendix: Answers to Post-Hearing Questions

PLANNING FOR THE FUTURE OF CYBER ATTACK ATTRIBUTION

THURSDAY, JULY 15, 2010

HOUSE OF REPRESENTATIVES,
SUBCOMMITTEE ON TECHNOLOGY AND INNOVATION,
COMMITTEE ON SCIENCE AND TECHNOLOGY,
Washington, DC.

The Subcommittee met, pursuant to call, at 10:04 a.m., in Room 2318 of the Rayburn House Office Building, Hon. David Wu [Chairman of the Subcommittee] presiding.

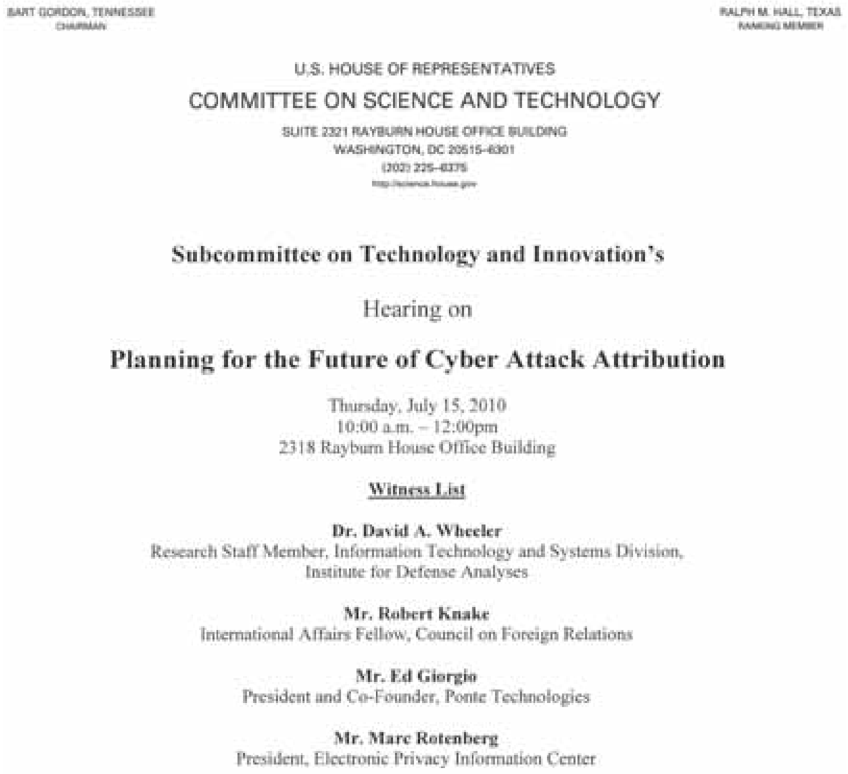

BART GORDON, TENNESSEE
CHAIRMAN

RALPH M. HALL, TEXAS
RANKING MEMBER

U.S. HOUSE OF REPRESENTATIVES

COMMITTEE ON SCIENCE AND TECHNOLOGY

SUITE 2321 RAYBURN HOUSE OFFICE BUILDING
WASHINGTON, DC 20515-6301
(202) 225-6375
http://science.house.gov

Subcommittee on Technology and Innovation's

Hearing on

Planning for the Future of Cyber Attack Attribution

Thursday, July 15, 2010
10:00 a.m. – 12:00pm
2318 Rayburn House Office Building

Witness List

Dr. David A. Wheeler
Research Staff Member, Information Technology and Systems Division,
Institute for Defense Analyses

Mr. Robert Knake
International Affairs Fellow, Council on Foreign Relations

Mr. Ed Giorgio
President and Co-Founder, Ponte Technologies

Mr. Marc Rotenberg
President, Electronic Privacy Information Center

HEARING CHARTER

COMMITTEE ON SCIENCE AND TECHNOLOGY
SUBCOMMITTEE ON TECHNOLOGY AND INNOVATION
U.S. HOUSE OF REPRESENTATIVES

Planning for the Future of Cyber Attack Attribution

THURSDAY, JULY 15, 2010
10:00 A.M.–12:00 P.M.
2318 RAYBURN HOUSE OFFICE BUILDING

I. Purpose

On Thursday, July 15, 2010, the Subcommittee on Technology and Innovation will hold a hearing to discuss attribution in cyber attacks, and how attribution technologies have the potential to affect the anonymity and privacy of internet users.

II. Witnesses

Dr. David Wheeler is a Research Staff Member of the Information Technology and Systems Division at the Institute for Defense Analyses.

Mr. Robert Knake is an International Affairs Fellow at the Council on Foreign Relations.

Mr. Ed Giorgio is the President and Co-Founder of Ponte Technologies.

Mr. Marc Rotenberg is the President of the Electronic Privacy Information Center.

III. Background

Cyber Attacks

Statistics clearly show that cyber attacks are common and costly. Following a recent survey of more than 2000 companies worldwide, Symantec reported that 42 percent rated cyber risk as their top concern, beating out other risks such as natural disasters, terrorism, and traditional crime. Symantec also reported that 75 percent of companies reported cyber attacks in the past twelve months and that 92 percent had seen significant monetary costs, averaging $2 million per year per company, as a result of those attacks.[1]

A 2004 Congressional Research Service report stated that "the stock price impact of cyber-attacks show that identified target firms suffer losses of 1%–5% in the days after an attack. For the average New York Stock Exchange corporation, price drops of these magnitudes translate into shareholder losses of between $50 million and $200 million".[2] According to a Market Wire article published in 2007, the economic impact from one comprehensive cyber attack on critical infrastructure could exceed $700 billion.[3]

Role of Attribution Technology

Being able to identify an attacker can be a strong deterrent against attack. During the Cold War, the Soviet Union and the United States remained in a nuclear standoff because either country would have been able to identify its attacker and stage a counter attack. In contrast, if a person, company, or government is attacked in cyberspace, it is often arduous—if not impossible—to determine the perpetrator of the attack.

[1] Symantec. (2010). *2010 State of Enterprise Security Global Results*. Retrieved from *http://www.slideshare.net/symantec/2010-state-of-enterprise-security*

[2] Congressional Research Service. (2004, April 1). *The Economic Impact of Cyber-Attacks*. (Order Code RL32331). Washington, D.C.: Congressional Research Service. Retrieved from *http://www.cisco.com/warp/public/779/govtaffairs/images/CRS_Cyber_Attacks.pdf*

[3] "New Research Shows Cyber Attack Could Cost U.S. 50 Times More Than Katrina". Market Wire. FindArticles.com. 09 Jul, 2010. *http://findarticles.com/p/articles/mi_pwwi/is_200707/ai_n19429846/*

Attribution technologies can be a useful tool in identifying and locating the assailant in a cyber attack. In terms of cyber attacks, attribution can be defined as "determining the identity or location of an attacker or an attacker's intermediary".[4] The attacker's identity can include a person's name, account information, or an alias. The location may include a geographical location or a virtual location, such as an IP address or Ethernet address.

In some cases, attribution technology may simply trace an attack back to an intermediary through which the attacker worked. For example, an attack can be transmitted via a fleet of 'zombies', or computers that can both delay and increase the severity of the attack. A sophisticated attacker may even be able to hide his or her identity so well that those looking for the attacker might falsely attribute the attack to an unrelated party. This can be done by an attacker who intentionally creates a false trail by sending incorrect data through any attribution process. To be effective and useful, new attribution technologies will need to have the ability to counter these, and future, methods of contravention.

The December 2009 attack on Google email accounts belonging to Chinese human rights activists in the United States, Europe, and China demonstrates the need for improvements in attribution technologies. Because the attacks showed a new level of sophistication, attributing their source has been a particularly difficult process. While the U.S. has been successful in tracing the attacks to two technical schools, it is still not known who was specifically behind these attacks.

In addition to helping to gain information about an isolated attack on a specific machine or network, successful attribution technologies can also be used to increase the security of the internet for people accessing personal information online—logging into a personal bank account, for example. If an online account required a recognizable IP range in addition to a pin code to retrieve account information, the ability of a hacker to access the account would be limited.

Anonymity and Privacy

Complete attribution may have negative ramifications for internet anonymity and privacy. For example, dissidents in countries where the government censures websites with firewalls may bypass or attack those firewalls to access prohibited information. If the government had attribution technology that allowed it to completely attribute the attack to its firewall, the government might use the information gained through attribution to punish dissidents for accessing the information. There is also the potential for attribution technologies to be used by a government, a company, or individual to identify the source of a posting or comment on the internet that is intended to be anonymous.

IV. Issues and Concerns

As more and more of the Nation's infrastructure becomes dependent on the internet, the potential impact of a successful cyber attack against the United States increases. Many of the tools we rely upon in our daily lives (traffic lights, restocking food supplies, millions of office jobs, etc.) have the potential to be rendered non-functional through a cyber attack. While attribution technologies may play an important role in limiting the effects of such crippling attacks, there may need to be clearly defined limits on when such technologies should be used. For example, proactively tracing interactions within a system may help determine where an attack originated after one occurs, but tracing every interaction is impractical and quite likely unconstitutional. It may be appropriate, therefore, to limit the use of attribution technology in most cases to post-attack.

A second area of interest is who is, or should be, responsible for the development, coordination, and implementation of attribution technologies. Even if some critical infrastructure is privately owned, the government arguably has a responsibility to its citizens to ensure that the infrastructure is protected. Given the interest in ensuring that government resources are utilized efficiently, there may be a need to strengthen coordination and collaboration between government and industry on the development of new attribution technologies in order to avoid redundancy and leverage resources.

There may also be a need to determine the appropriate role of the government in responding to cyber attacks on private companies and individuals. In general, if a company or individual is physically attacked by an outside government, a company, or an individual, it is quite likely that the government would step in and defend the attacked company or individual. If a company or individual is the victim

[4] David A. Wheeler and Gregory N. Larsen, *Techniques for Cyber Attack Attribution* (Institute for Defense Analysis, IDA Paper P–3792. October 2003), p.1

of a cyber attack, it is currently unclear what the government's role is, or should be, in responding to the attack.

Finally, the implications of attribution technologies for the anonymity and privacy of internet users should be considered. It may be necessary to consider ways to limit the use of attribution technologies to identifying the source of cyber attacks and in ways that do not suppress the freedom of speech or otherwise implicate the anonymity and privacy of people using the internet for legitimate purposes. There may also be a need to determine who (government or industry or both) should maintain responsibility for ensuring that attribution technologies are used consistent with any identified limits.

V. Overarching Questions

The following questions were asked of each witness:

- As has been stated by many experts, deterrence is a productive way to prevent physical attacks. How can attack attribution play a role in deterring cyber attacks?
- What are the proper roles of both the government and private industry in developing and improving attack attribution capabilities? What R&D is needed to address capability gaps in attack attribution and who should be responsible for completing that R&D?
- What are the distinguishing factors between anonymity and privacy? How should we account for both in the development and use of attribution technologies?
- Is there a need for standards in the development and implementation of attack attribution technologies? Is there a specific need for privacy standards and if so, what should be the government's role in the development of these standards?

Chairman WU. The hearing will come to order.

Good morning, and thank you very much for being at this cyber attribution hearing.

This cybersecurity hearing is one in a series that this Subcommittee has held on ways that we can protect our Nation's critical cyber infrastructure. Over the last two years, we have held hearings on cybersecurity activities at the National Institute of Standards and Technology and the Department of Homeland Security, as well as on the Administration's Cyberspace Policy Review. Just two weeks ago, we had an important hearing on the Smart Grid, and spent a great deal of time talking about the necessity of developing strong cybersecurity standards for our national energy infrastructure.

We are well aware of the critical role that IT [Information Technology] networks play in managing much of our day-to-day activity from online banking to systems that make sure there is food on our grocery shelves. This growing reliance on networks has made us more vulnerable to cyber attacks and has increased the potential for such attacks to have far-reaching and crippling effects. Now more than ever, we need to be focused on the development of tools and technologies to prevent, detect, and respond to cyber attacks.

History shows that one of the best deterrents to an attack is the ability to identify your attacker. The question is whether such deterrence methods are still relevant today. During the Cold War, the United States and the Soviet Union, each with quite expansive offensive capabilities, were held in check by the notion that an attack would result in retaliation. This was achieved because each country would have been able to precisely identify its attacker. This method of deterrence, the ability to attribute an attack to a particular person, party or system, can be equally vital to defending against cyber attack. While they are not the end-all solution to our cybersecurity challenges, the development of effective and reliable attribution technologies should be an essential part of our efforts to secure the Nation's cyberspace.

Given that the Internet is intended to be open and anonymous, the attribution of cyber attacks can be very, very difficult to achieve and should not be taken lightly. As co-chair of the Global Internet Freedom Caucus in the House, I am personally very concerned about the potential implications to privacy, anonymity and Internet freedom posed by attribution technologies. As a result, I believe that it is absolutely imperative that we define and implement clear restrictions on how attribution technologies are developed and used to ensure that they are not misused.

I look forward to today's discussion on attribution technologies and how they may help deter cyber attacks. I am interested in discussing the proper roles of the Federal Government and private industry in the development of these technologies, and the research and development that is needed to fill capability gaps. I am sure—and I am particularly eager to discuss ways to ensure that attribution technologies are not used to infringe upon the safety, privacy or individual liberties of Internet users.

I would like to thank the witnesses for appearing before us today, and I look forward to our discussion.

Now I recognize Mr. Hall, the Ranking Member of the Full Committee, for his opening statement.

[The prepared statement of Chairman Wu follows:]

PREPARED STATEMENT OF CHAIRMAN DAVID WU

Good morning and thank you for coming to today's hearing focused on interoperability in public safety communication equipment.

We've learned an important lesson from September 11th, Hurricane Katrina, and other disasters: interoperable communication is critical to effective emergency response. When time is of the essence and lives are at stake, a clear flow of information is essential. Unfortunately, it is not uncommon for police officers and firefighters from a single region, or even a single city, to be using incompatible communication systems. This lack of interoperability has contributed to the deaths of first responders and hindered the ability to rescue people in harm's way.

Enabling interoperable communication systems, where public safety personnel can talk with each other in real-time, takes planning and cooperation by all levels of government. However, interoperability also demands radios that are capable of communicating with one another. First responders on digital land mobile radio systems built to proprietary specifications cannot communicate. Ad-hoc solutions, like patching technologies or sharing radios, are less efficient than the seamless interoperability offered by systems based on open architecture.

The purpose of today's hearing is to examine the status of the standards development process for this open architecture. Since 1989, the public safety community and industry have been working together on Project-25, or P25, a suite of standards that will not only enable interoperability, but also promote competition in the marketplace for digital land mobile radio systems and provide other benefits. While there has been a lot of progress on the P25 standards since 1989, the entire set of standards remains incomplete. I would like to understand the implications of this for public safety agencies procuring systems sold as "P25 compliant" and get a better sense of when we realistically can expect all of the standards to be completed.

A second issue that we will discuss today is the lack of a formal compliance assessment process for the P25 standards. A compliance assessment process signals to the purchaser that a product meets all of the requirements of a standard. Any laptop with a Wi-Fi logo, or any toaster with an Underwriter's Laboratory sticker, had to go through testing and certification to be able to display those marks. P25 does not have an equivalent process. The Department of Homeland Security's Compliance Assessment Program fills this gap, but we must be sure it provides the highest possible level of assurance to the public safety community that systems sold as P25-complaint actually meet all of the requirements of the standards. It seems to me that there ought to be a formal, comprehensive system in place to ensure that it is not *caveat emptor* when first responders spend millions of dollars on complex communications technology.

The most important question for the first responders who rely on this equipment is "does it work?" In addition to being mission-critical technology, these systems represent *major expenditures* for government agencies across the country. Particularly at a time of uncertain and dwindling budgets, cost-effective procurement enabled by an open-architecture is essential.

I'd like to thank our witnesses for being here today. Project 25 is unique in the world of standards development in that the users of the technology—in this case, our public safety officials—are integral to, and directly involved in, the standards development process. It is important that this process move forward, and that the public safety community and industry continue to work together to make further advances in first responder technology.

Mr. HALL. Thank you, Mr. Chairman, and since you have made an excellent opening statement and covered almost everything, I can be brief, and I am filling in for the Ranking Member, Mr. Smith, and I thank you for calling the hearing on cyber attack attribution technologies. I also want to thank our very distinguished panel. We rely on you to tell us what the facts are, and from that we glean legislation, and don't be disturbed by the empty chairs here because they will all receive copies of your testimony, and many have received copies ahead of time. I have scanned through your testimony. I want to thank the panel for being here and ask you to remember that we are not technical experts, so keep it as simple as you possibly can. I have read some of your testimony and understood a lot of it. Ranking Smith is going to be here shortly. In the event it takes him longer than expected, I ask unanimous consent that his statement be made a part of the record, Mr. Chairman.

Otherwise I will yield the remainder of my time to him when he arrives. Thank you, sir.

[The prepared statement of Mr. Smith follows:]

Prepared Statement of Representative Adrian Smith

Thank you, Chairman Wu, for calling today's hearing on cyber attack attribution. Once again this subcommittee will have the opportunity to hear from an outstanding panel of expert witnesses, and I thank them for taking the time to be with us today.

With the integration of computing technology into nearly every aspect of our professional and private lives—from growing our food to managing our electrical grid to tracking every financial transaction no matter how small—the threat of a catastrophic attack on the networks which manage every sector of our economic and security infrastructure has also grown exponentially.

As we search for effective ways to prevent such an attack, one widely discussed means is deterrence through attribution—ensuring would-be attackers know any activities would be traced back to them with reciprocal action in return.

The work of tracing such attacks, particularly in the United States where the presumption of innocence is sacrosanct and where privacy for the innocent is respected, this is easier said than done. This raises a number of questions I hope we can address in today's hearing:

- What are the best methods for tracing attacks?
- What harriers exist, aside from technological ones, to tracing attacks inside and outside our borders?
- If we can trace attacks, what is an effective deterrent to prevent them?
- And if we can answer the first three questions effectively, what is the role for standards-setting bodies in assisting government and the private sector in reaching those conclusions?

I hope we can also consider the consequences of traceability on the overwhelming majority who use computer systems lawfully and whose privacy we should respect.

Before we move on to hearing from our witness, I would like to briefly note it is my understanding a follow-up hearing in which we hear from NIST, National Science Foundation, and other applicable Federal agencies is under consideration, and I would like to offer my support for holding such a hearing.

Thank you again, Chairman Wu and witnesses. I expect we will learn a lot today, and I yield back the balance of my time.

Chairman Wu. Thank you very much, Mr. Hall.

If there are Members who wish to submit opening statements, your statements will be added to the record at this point. And I also want to recognize the Chairman of the Full Committee, who is in attendance, and Chairman Gordon—very good. Thank you.

Now it is my pleasure to introduce our witnesses. Dr. David A. Wheeler is a Research Staff Member of the Information Technology and Systems Division at the Institute for Defense Analyses. Mr. Robert Knake is International Affairs Fellow at the Council on Foreign Relations. Mr. Ed Giorgio is the President and Co-Founder of Ponte Technologies. He also has over 30 years of security experience at the National Security Agency, or NSA, and is a leading authority on security and cryptography, and I want to recognize that Mr. Giorgio is also wearing a Distinguished Service Medal awarded by the NSA. And our final witness is Mr. Marc Rotenberg, who is the President of the Electronic Privacy Information Center, or

EPIC, and at our prior hearing on grid security, one of your vice presidents provided very, very interesting, elucidating comments.

You will each have five minutes for your spoken testimony, and your written testimony will be included in the record of this hearing. When you all complete your testimony, we will begin with questions, and each Member will have five minutes to question the witnesses.

Dr. Wheeler, please proceed.

STATEMENT OF DAVID A. WHEELER, RESEARCH STAFF MEMBER, INFORMATION TECHNOLOGY AND SYSTEMS DIVISION, INSTITUTE FOR DEFENSE ANALYSES

Dr. WHEELER. Mr. Chairman, distinguished Members of the House Subcommittee on Technology and Innovation and the Committee on Science and Technology, I am delighted to speak with you today. As noted, my name is Dr. David A. Wheeler. I work at the Institute for Defense Analyses, also known as IDA. IDA is, and I quote, "a nonprofit corporation that operates three federally funded research and development centers," or FFRDCs. These FFRDCs provide objective analyses of national security issues, particularly those requiring scientific and technical expertise, and they conduct related research on other national challenges.

In 2002 and 2003, I developed a survey of cyber attack attribution technologies on behalf of the Department of Defense, DoD. This survey has been provided to this Subcommittee and is also available to the public from the Defense Technical Information Center as IDA paper P–3792, Techniques on Cyber Attribution. Attribution in this context is determining the identity or location of an attacker or an attacker's intermediary. Since writing that paper, I have worked on improving the security and assurance of systems, lowering supply chain risks, improving open standards and eliminating barriers to the use and development of open source software.

It is good that this Subcommittee is examining the relationship between attribution, privacy and anonymity. As I noted in my paper, we should be concerned if attribution technologies developed in democracies are acquired and redeployed by governments with abusive human rights records to suppress freedom of speech and democracy movements.

Apart from any concern of abuse by foreign governments, the use of these techniques by our government requires consideration of the Fourth Amendment's guarantee that people must be secure against unreasonable searches and seizures. Section 3.13 of my paper specifically discusses the need to protect privacy and freedom of speech. With that as context, I will address the overarching questions in this hearing's charter.

The first question asked about the role of attack attribution in deterring cyber attacks. It noted that deterrence is a productive way to prevent physical attacks. In a similar way, cyber attack attribution can play an important role in deterring cyber attacks by enabling many deterrence measures. While there is great need to harden U.S. infrastructure from cyber attacks, passive computer network defenses cannot be and never will be perfect. This means that in some cases we may need to be able to respond to an attack.

Unfortunately, many other countermeasures such as computer network counterattack, legal action and kinetic energy counterattack can only be deployed if the source of the attack can be attributed with high confidence.

The second question asked what roles that government and private industry should play. As of 2003, there was little evidence that the commercial sector was willing to shoulder the costs to develop attribution capabilities. Most commercial companies appear to view identifying attackers as a law enforcement or military task, not a commercial one. If the government wants the ability to attribute attacks, in many cases the government may need to pay for it directly. One approach is to fund development and deployment of these abilities for widely used applications both proprietary and open source software. More than one product in each category should be funded, so that the government is not locked into a single supplier.

The third question asked for the distinguishing factors between anonymity and privacy and how to account for both in the development and use of attribution technologies. As I noted in my paper, if the United States is to develop attribution technology, it should encourage the development or implementation of those attribution technologies that pose less danger to privacy. For example, logging systems could store message hashes, also known as message fingerprints, instead of the messages themselves. Since the data isn't stored, hashing only supports attribution of data the requester has already seen. A key part of implementing attribution technologies with few risks to privacy and anonymity is to ensure that any standards development related to attribution should include efforts to address these privacy and anonymity concerns.

This brings me to the issue of standards, the focus of the fourth question. Standards are critically necessary for some attribution technologies, and the standards development process should work to address these privacy and anonymity concerns through public development and review. Such standards should be open standards to permit competition; in particular, they should be publicly defined and held and shouldn't be patent-encumbered. This suggests that the U.S. government should be involved in the development of such standards to ensure that its needs and concerns are met, just as the government is already involved in the development of standards where there are specific government needs and concerns.

I will be happy to address your questions.

[The prepared statement of Dr. Wheeler follows:]

PREPARED STATEMENT OF DAVID A. WHEELER

It is an honor to provide testimony to you. Please consider the attached paper, "Techniques for Cyber Attack Attribution" (IDA Paper P–3792) as my written testimony. This paper discusses techniques for cyber attack attribution, including notes about the relationship of attribution to privacy.

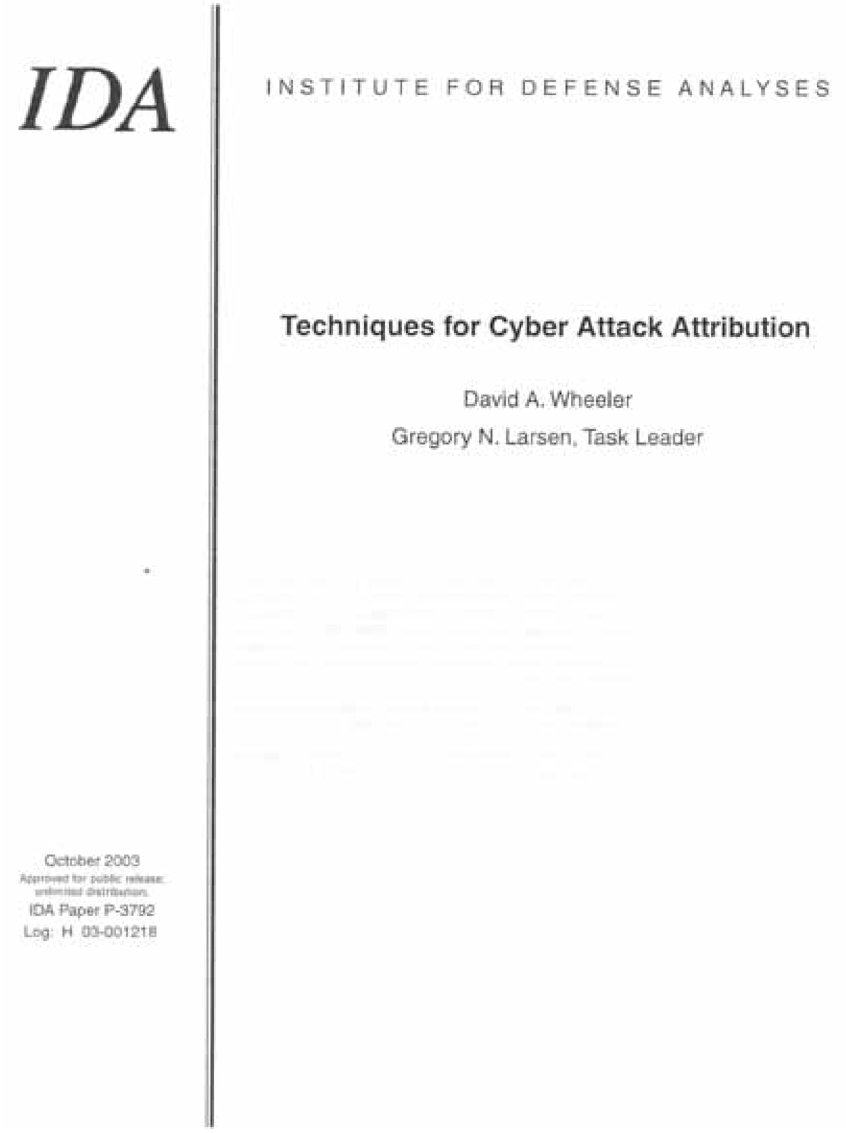

IDA

INSTITUTE FOR DEFENSE ANALYSES

Techniques for Cyber Attack Attribution

David A. Wheeler
Gregory N. Larsen, Task Leader

October 2003
Approved for public release;
unlimited distribution.
IDA Paper P-3792
Log: H 03-001218

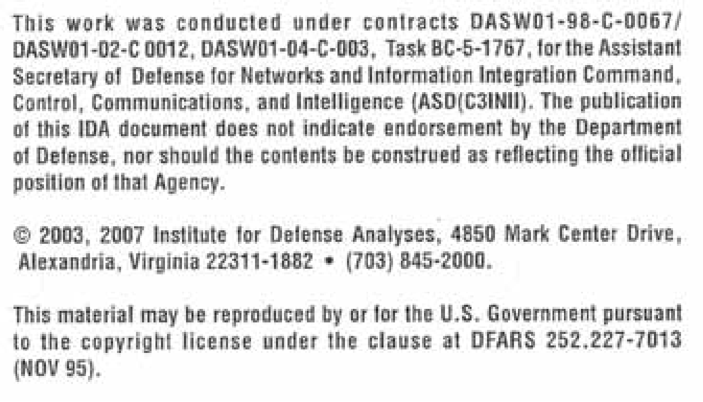

This work was conducted under contracts DASW01-98-C-0067/ DASW01-02-C 0012, DASW01-04-C-003, Task BC-5-1767, for the Assistant Secretary of Defense for Networks and Information Integration Command, Control, Communications, and Intelligence (ASD(C3INII). The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

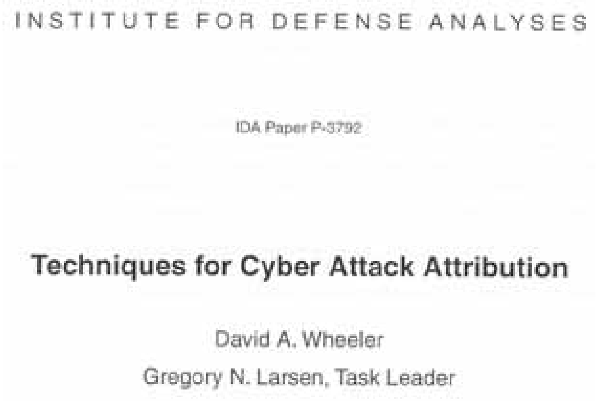

INSTITUTE FOR DEFENSE ANALYSES

IDA Paper P-3792

Techniques for Cyber Attack Attribution

David A. Wheeler
Gregory N. Larsen, Task Leader

Preface

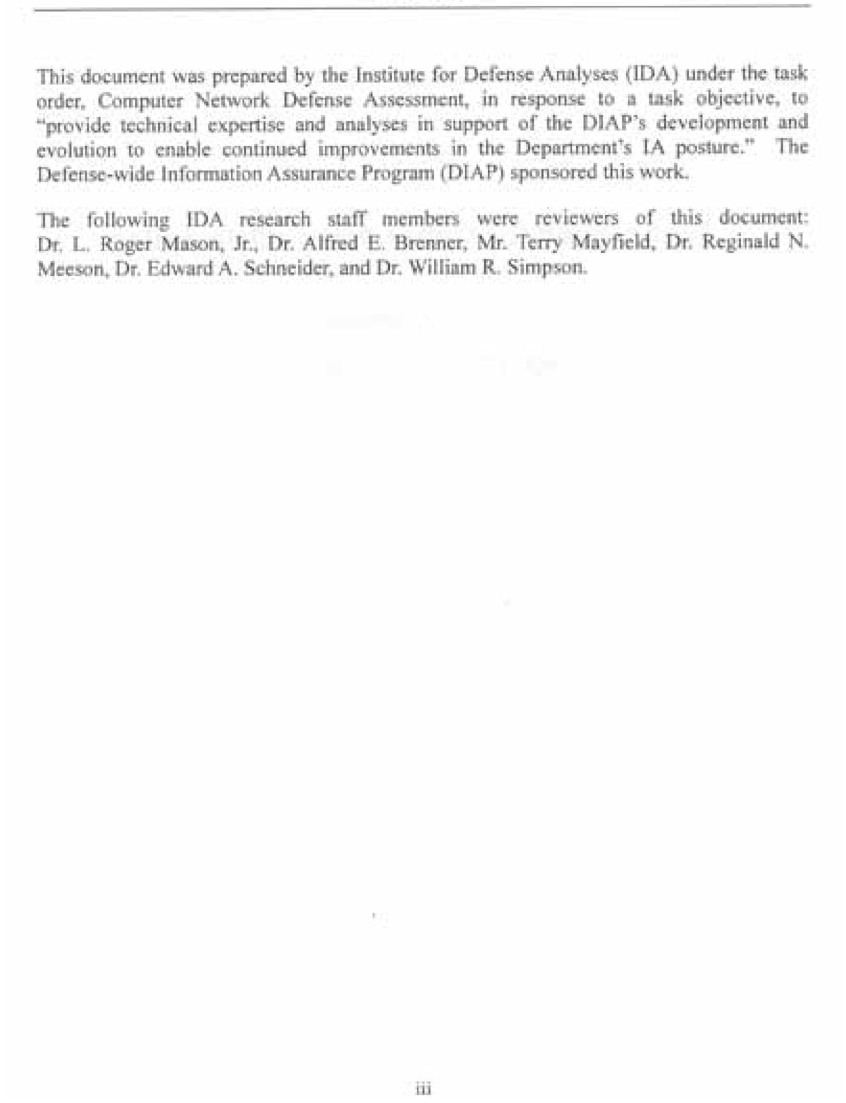

This document was prepared by the Institute for Defense Analyses (IDA) under the task order, Computer Network Defense Assessment, in response to a task objective, to "provide technical expertise and analyses in support of the DIAP's development and evolution to enable continued improvements in the Department's IA posture." The Defense-wide Information Assurance Program (DIAP) sponsored this work.

The following IDA research staff members were reviewers of this document: Dr. L. Roger Mason, Jr., Dr. Alfred E. Brenner, Mr. Terry Mayfield, Dr. Reginald N. Meeson, Dr. Edward A. Schneider, and Dr. William R. Simpson.

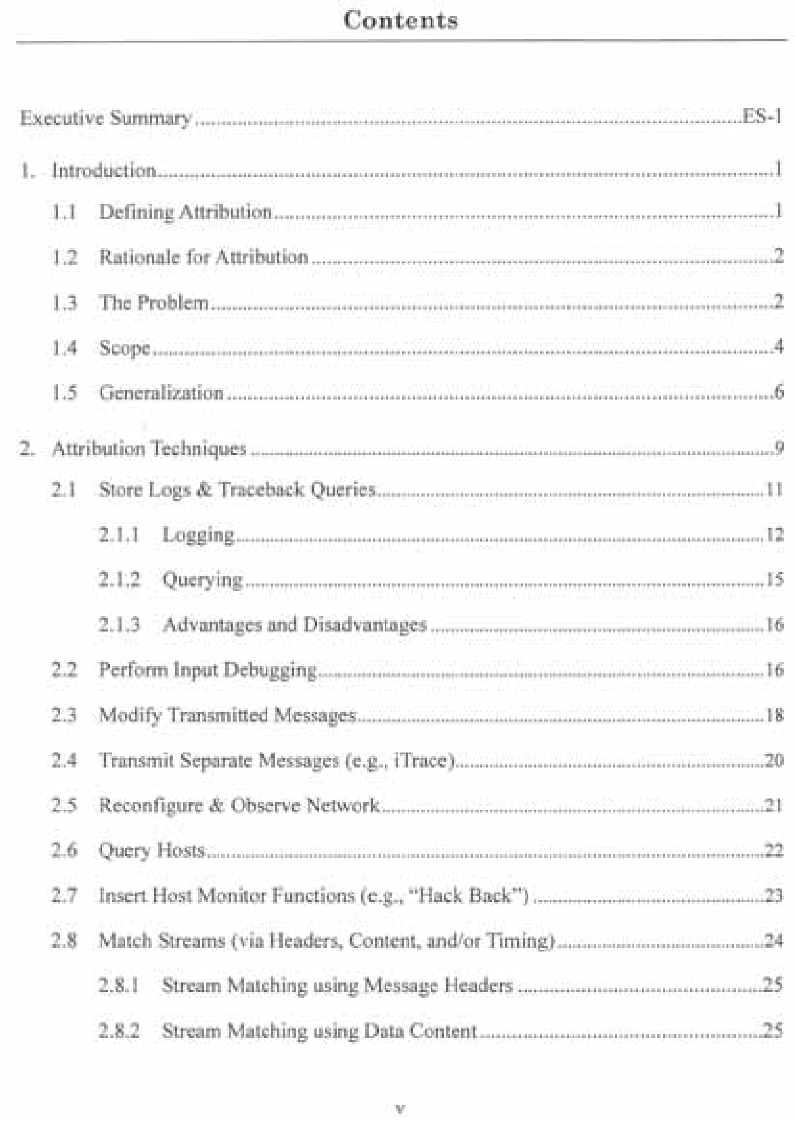

Contents

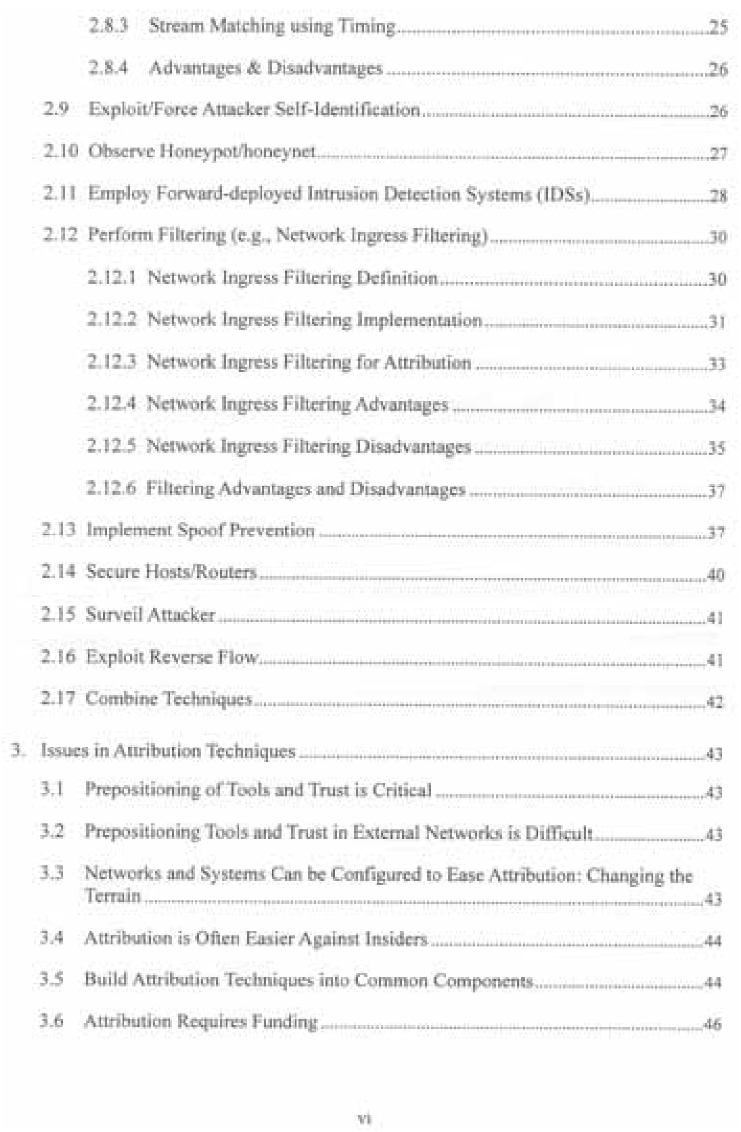

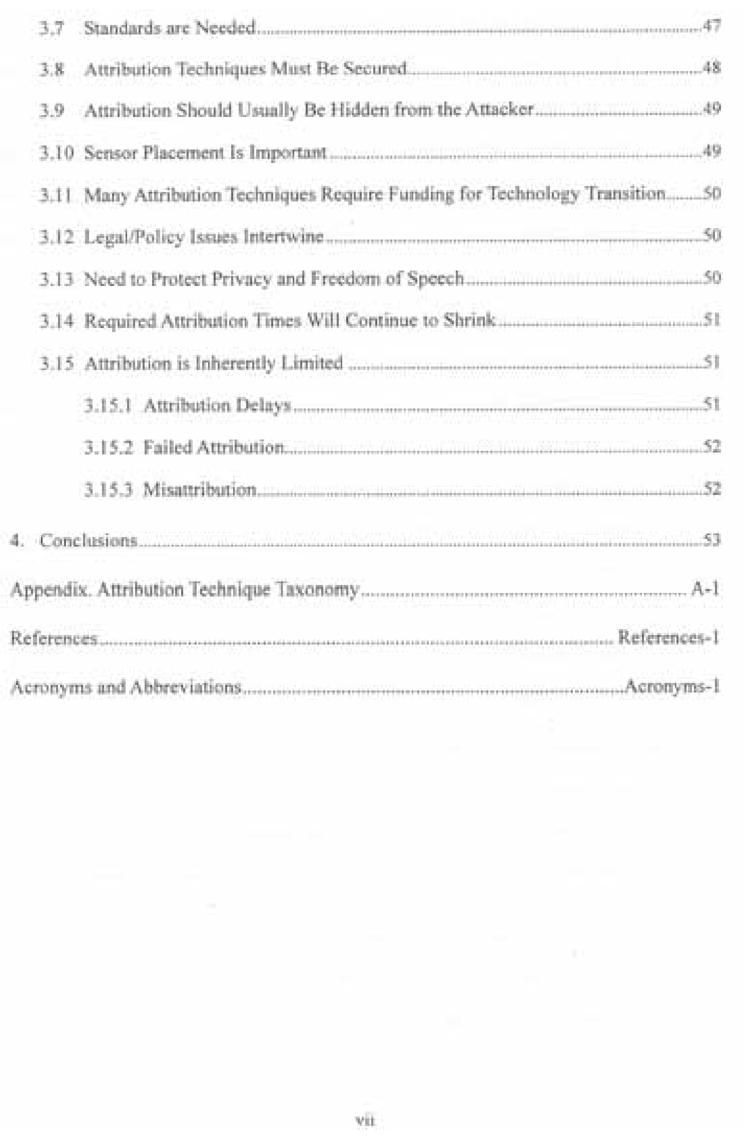

Figures

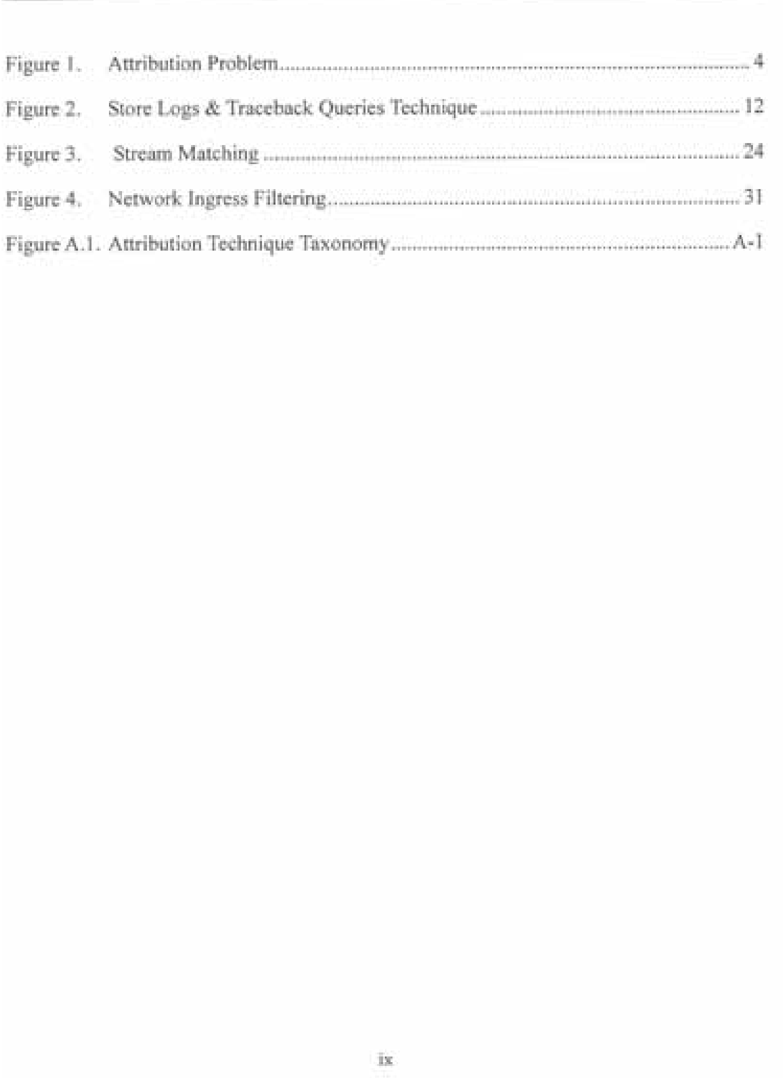

Tables

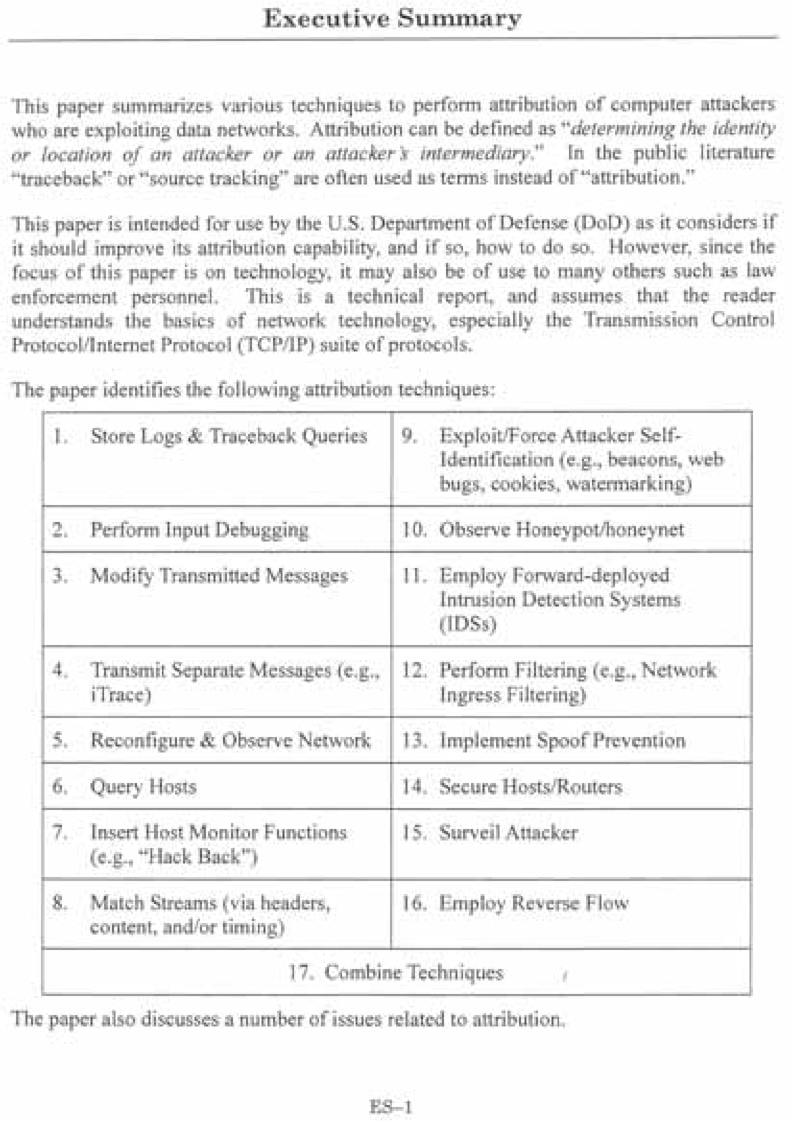

Executive Summary

This paper summarizes various techniques to perform attribution of computer attackers who are exploiting data networks. Attribution can be defined as "*determining the identity or location of an attacker or an attacker's intermediary*." In the public literature "traceback" or "source tracking" are often used as terms instead of "attribution."

This paper is intended for use by the U.S. Department of Defense (DoD) as it considers if it should improve its attribution capability, and if so, how to do so. However, since the focus of this paper is on technology, it may also be of use to many others such as law enforcement personnel. This is a technical report, and assumes that the reader understands the basics of network technology, especially the Transmission Control Protocol/Internet Protocol (TCP/IP) suite of protocols.

The paper identifies the following attribution techniques:

1. Store Logs & Traceback Queries	9. Exploit/Force Attacker Self-Identification (e.g., beacons, web bugs, cookies, watermarking)
2. Perform Input Debugging	10. Observe Honeypot/honeynet
3. Modify Transmitted Messages	11. Employ Forward-deployed Intrusion Detection Systems (IDSs)
4. Transmit Separate Messages (e.g., iTrace)	12. Perform Filtering (e.g., Network Ingress Filtering)
5. Reconfigure & Observe Network	13. Implement Spoof Prevention
6. Query Hosts	14. Secure Hosts/Routers
7. Insert Host Monitor Functions (e.g., "Hack Back")	15. Surveil Attacker
8. Match Streams (via headers, content, and/or timing)	16. Employ Reverse Flow
17. Combine Techniques	

The paper also discusses a number of issues related to attribution.

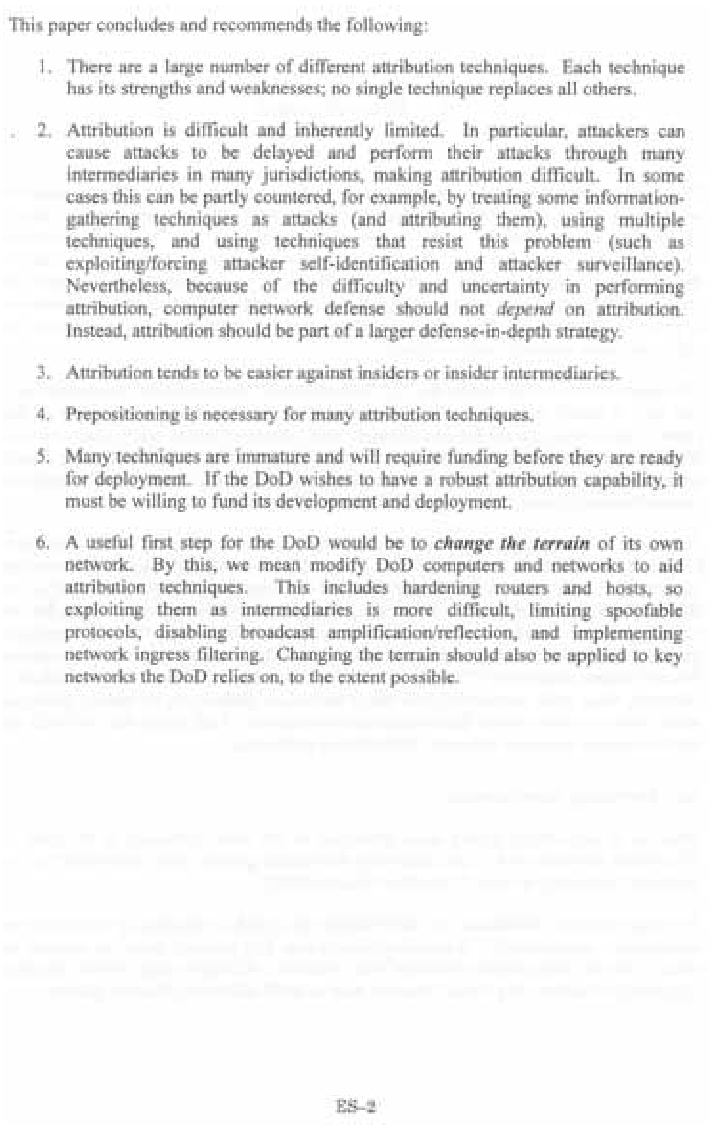

This paper concludes and recommends the following:

1. There are a large number of different attribution techniques. Each technique has its strengths and weaknesses; no single technique replaces all others.
2. Attribution is difficult and inherently limited. In particular, attackers can cause attacks to be delayed and perform their attacks through many intermediaries in many jurisdictions, making attribution difficult. In some cases this can be partly countered, for example, by treating some information-gathering techniques as attacks (and attributing them), using multiple techniques, and using techniques that resist this problem (such as exploiting/forcing attacker self-identification and attacker surveillance). Nevertheless, because of the difficulty and uncertainty in performing attribution, computer network defense should not *depend* on attribution. Instead, attribution should be part of a larger defense-in-depth strategy.
3. Attribution tends to be easier against insiders or insider intermediaries.
4. Prepositioning is necessary for many attribution techniques.
5. Many techniques are immature and will require funding before they are ready for deployment. If the DoD wishes to have a robust attribution capability, it must be willing to fund its development and deployment.
6. A useful first step for the DoD would be to ***change the terrain*** of its own network. By this, we mean modify DoD computers and networks to aid attribution techniques. This includes hardening routers and hosts, so exploiting them as intermediaries is more difficult, limiting spoofable protocols, disabling broadcast amplification/reflection, and implementing network ingress filtering. Changing the terrain should also be applied to key networks the DoD relies on, to the extent possible.

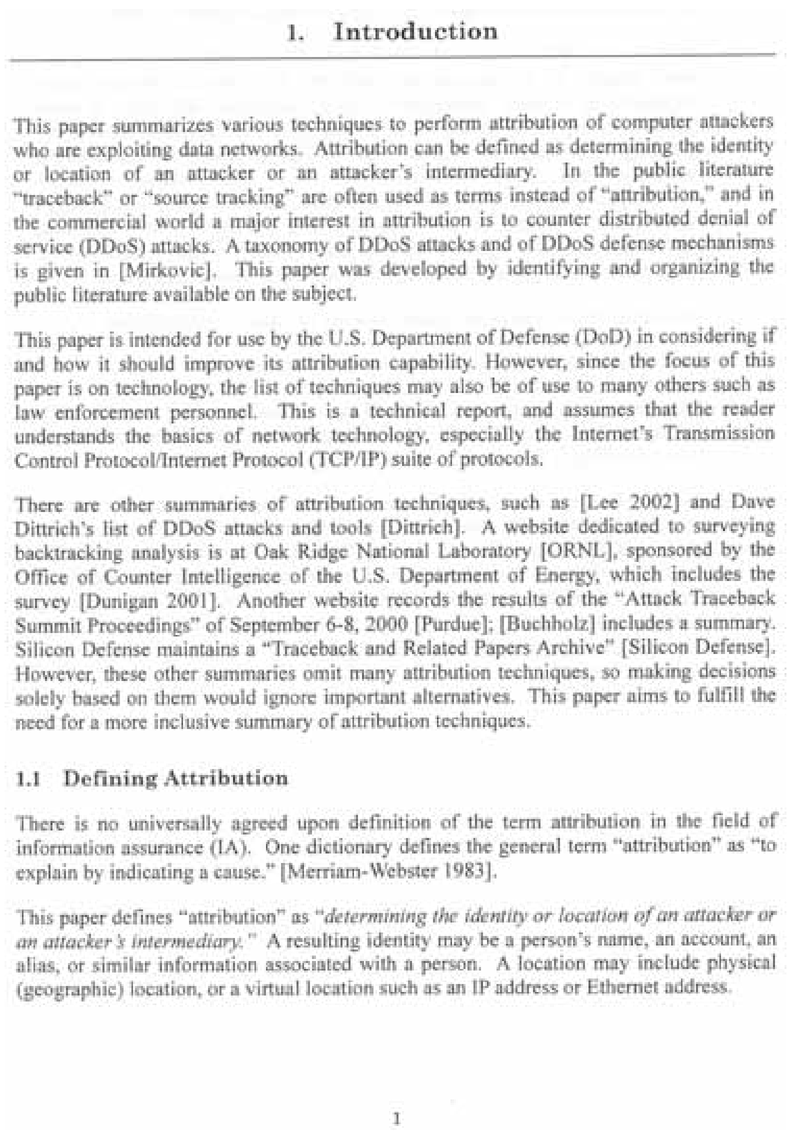

1. Introduction

This paper summarizes various techniques to perform attribution of computer attackers who are exploiting data networks. Attribution can be defined as determining the identity or location of an attacker or an attacker's intermediary. In the public literature "traceback" or "source tracking" are often used as terms instead of "attribution," and in the commercial world a major interest in attribution is to counter distributed denial of service (DDoS) attacks. A taxonomy of DDoS attacks and of DDoS defense mechanisms is given in [Mirkovic]. This paper was developed by identifying and organizing the public literature available on the subject.

This paper is intended for use by the U.S. Department of Defense (DoD) in considering if and how it should improve its attribution capability. However, since the focus of this paper is on technology, the list of techniques may also be of use to many others such as law enforcement personnel. This is a technical report, and assumes that the reader understands the basics of network technology, especially the Internet's Transmission Control Protocol/Internet Protocol (TCP/IP) suite of protocols.

There are other summaries of attribution techniques, such as [Lee 2002] and Dave Dittrich's list of DDoS attacks and tools [Dittrich]. A website dedicated to surveying backtracking analysis is at Oak Ridge National Laboratory [ORNL], sponsored by the Office of Counter Intelligence of the U.S. Department of Energy, which includes the survey [Dunigan 2001]. Another website records the results of the "Attack Traceback Summit Proceedings" of September 6-8, 2000 [Purdue]; [Buchholz] includes a summary. Silicon Defense maintains a "Traceback and Related Papers Archive" [Silicon Defense]. However, these other summaries omit many attribution techniques, so making decisions solely based on them would ignore important alternatives. This paper aims to fulfill the need for a more inclusive summary of attribution techniques.

1.1 Defining Attribution

There is no universally agreed upon definition of the term attribution in the field of information assurance (IA). One dictionary defines the general term "attribution" as "to explain by indicating a cause." [Merriam-Webster 1983].

This paper defines "attribution" as "*determining the identity or location of an attacker or an attacker's intermediary.*" A resulting identity may be a person's name, an account, an alias, or similar information associated with a person. A location may include physical (geographic) location, or a virtual location such as an IP address or Ethernet address.

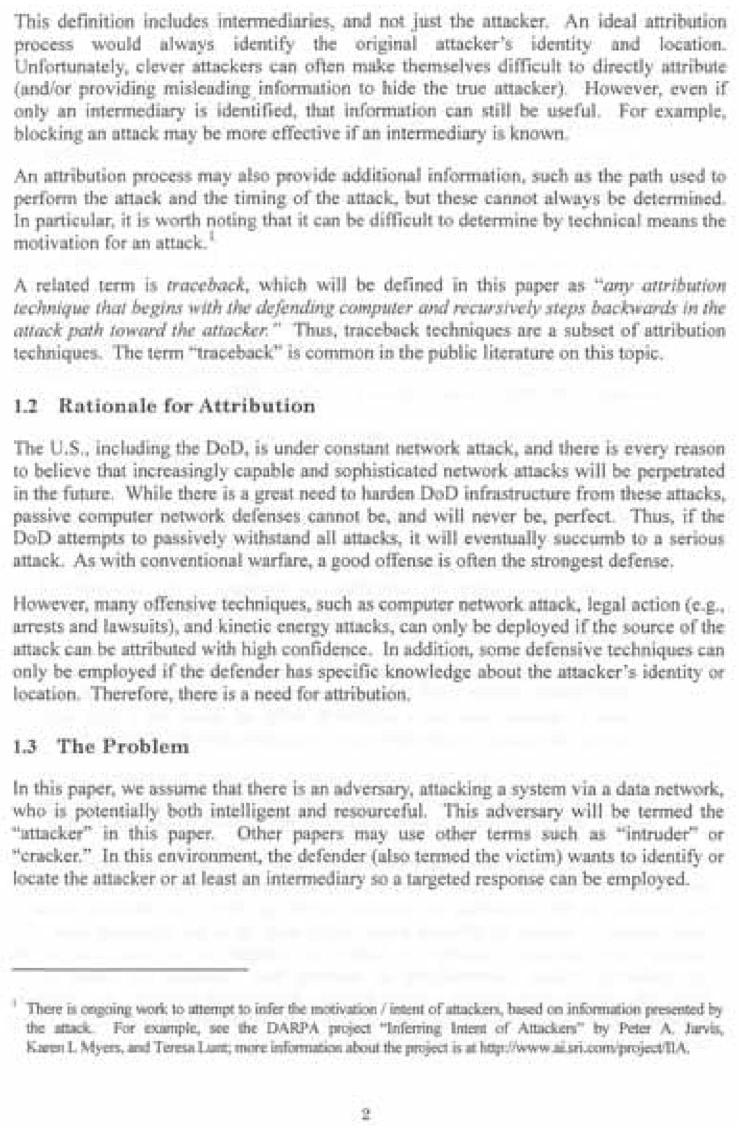

This definition includes intermediaries, and not just the attacker. An ideal attribution process would always identify the original attacker's identity and location. Unfortunately, clever attackers can often make themselves difficult to directly attribute (and/or providing misleading information to hide the true attacker). However, even if only an intermediary is identified, that information can still be useful. For example, blocking an attack may be more effective if an intermediary is known.

An attribution process may also provide additional information, such as the path used to perform the attack and the timing of the attack, but these cannot always be determined. In particular, it is worth noting that it can be difficult to determine by technical means the motivation for an attack.[1]

A related term is *traceback*, which will be defined in this paper as "*any attribution technique that begins with the defending computer and recursively steps backwards in the attack path toward the attacker.*" Thus, traceback techniques are a subset of attribution techniques. The term "traceback" is common in the public literature on this topic.

1.2 Rationale for Attribution

The U.S., including the DoD, is under constant network attack, and there is every reason to believe that increasingly capable and sophisticated network attacks will be perpetrated in the future. While there is a great need to harden DoD infrastructure from these attacks, passive computer network defenses cannot be, and will never be, perfect. Thus, if the DoD attempts to passively withstand all attacks, it will eventually succumb to a serious attack. As with conventional warfare, a good offense is often the strongest defense.

However, many offensive techniques, such as computer network attack, legal action (e.g., arrests and lawsuits), and kinetic energy attacks, can only be deployed if the source of the attack can be attributed with high confidence. In addition, some defensive techniques can only be employed if the defender has specific knowledge about the attacker's identity or location. Therefore, there is a need for attribution.

1.3 The Problem

In this paper, we assume that there is an adversary, attacking a system via a data network, who is potentially both intelligent and resourceful. This adversary will be termed the "attacker" in this paper. Other papers may use other terms such as "intruder" or "cracker." In this environment, the defender (also termed the victim) wants to identify or locate the attacker or at least an intermediary so a targeted response can be employed.

[1] There is ongoing work to attempt to infer the motivation / intent of attackers, based on information presented by the attack. For example, see the DARPA project "Inferring Intent of Attackers" by Peter A. Jarvis, Karen L Myers, and Teresa Lunt; more information about the project is at http://www.ai.sri.com/project/IIA.

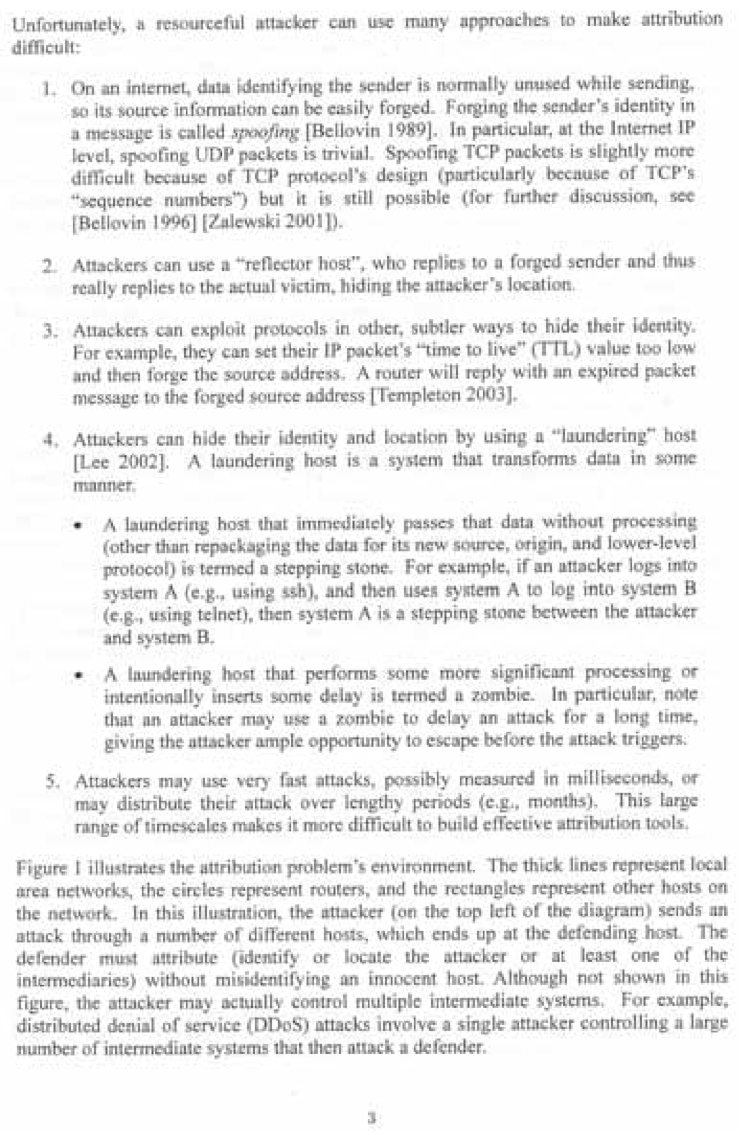

Unfortunately, a resourceful attacker can use many approaches to make attribution difficult:

1. On an internet, data identifying the sender is normally unused while sending, so its source information can be easily forged. Forging the sender's identity in a message is called *spoofing* [Bellovin 1989]. In particular, at the Internet IP level, spoofing UDP packets is trivial. Spoofing TCP packets is slightly more difficult because of TCP protocol's design (particularly because of TCP's "sequence numbers") but it is still possible (for further discussion, see [Bellovin 1996] [Zalewski 2001]).

2. Attackers can use a "reflector host", who replies to a forged sender and thus really replies to the actual victim, hiding the attacker's location.

3. Attackers can exploit protocols in other, subtler ways to hide their identity. For example, they can set their IP packet's "time to live" (TTL) value too low and then forge the source address. A router will reply with an expired packet message to the forged source address [Templeton 2003].

4. Attackers can hide their identity and location by using a "laundering" host [Lee 2002]. A laundering host is a system that transforms data in some manner.

   - A laundering host that immediately passes that data without processing (other than repackaging the data for its new source, origin, and lower-level protocol) is termed a stepping stone. For example, if an attacker logs into system A (e.g., using ssh), and then uses system A to log into system B (e.g., using telnet), then system A is a stepping stone between the attacker and system B.

   - A laundering host that performs some more significant processing or intentionally inserts some delay is termed a zombie. In particular, note that an attacker may use a zombie to delay an attack for a long time, giving the attacker ample opportunity to escape before the attack triggers.

5. Attackers may use very fast attacks, possibly measured in milliseconds, or may distribute their attack over lengthy periods (e.g., months). This large range of timescales makes it more difficult to build effective attribution tools.

Figure 1 illustrates the attribution problem's environment. The thick lines represent local area networks, the circles represent routers, and the rectangles represent other hosts on the network. In this illustration, the attacker (on the top left of the diagram) sends an attack through a number of different hosts, which ends up at the defending host. The defender must attribute (identify or locate the attacker or at least one of the intermediaries) without misidentifying an innocent host. Although not shown in this figure, the attacker may actually control multiple intermediate systems. For example, distributed denial of service (DDoS) attacks involve a single attacker controlling a large number of intermediate systems that then attack a defender.

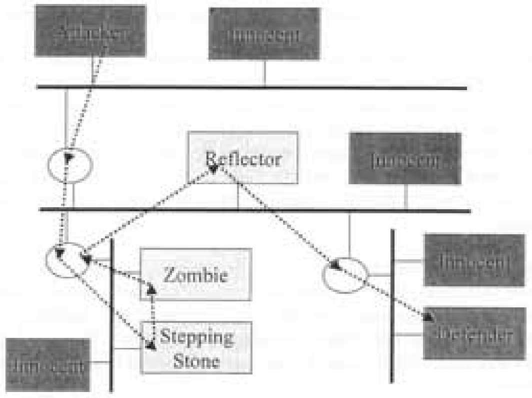

Figure 1. Attribution Problem

Modern environments often make attribution quite difficult. Typical computer network environments are *not* designed to support attribution of attackers. There are often many components in a network, making it easier for attackers to hide. Data paths may go through many systems in many countries or may be controlled by many different administrative domains, including those who may be hostile or noncooperative. Many networking capabilities unintentionally create complications for attribution, such as network address translation (NAT) that can change the sender and receiver address.

1.4 Scope

This paper is focused solely on identifying different techniques that could be used for attribution of attackers. This paper only examines attribution techniques for attackers attacking via an electronic data network (usually an Internet standards based network). Other attacks, such as physical attacks, social engineering attacks, or trusted programmers inserting malicious code into their own programs during development, are concerns but are outside the scope of this paper. This paper concentrates on approaches based on technology; non-technical approaches such as various human intelligence techniques are not the focus of this paper.

This paper does not cover identifying or locating people who are not directly attacking the defender. In particular, identifying or locating people voluntarily cooperating with each other is not covered in this paper, although some attribution techniques may also be useful in that case. [Wright 2002] describes some attribution-like techniques for anonymous peer-to-peer (P2P) networks. It also does not cover the general issue of

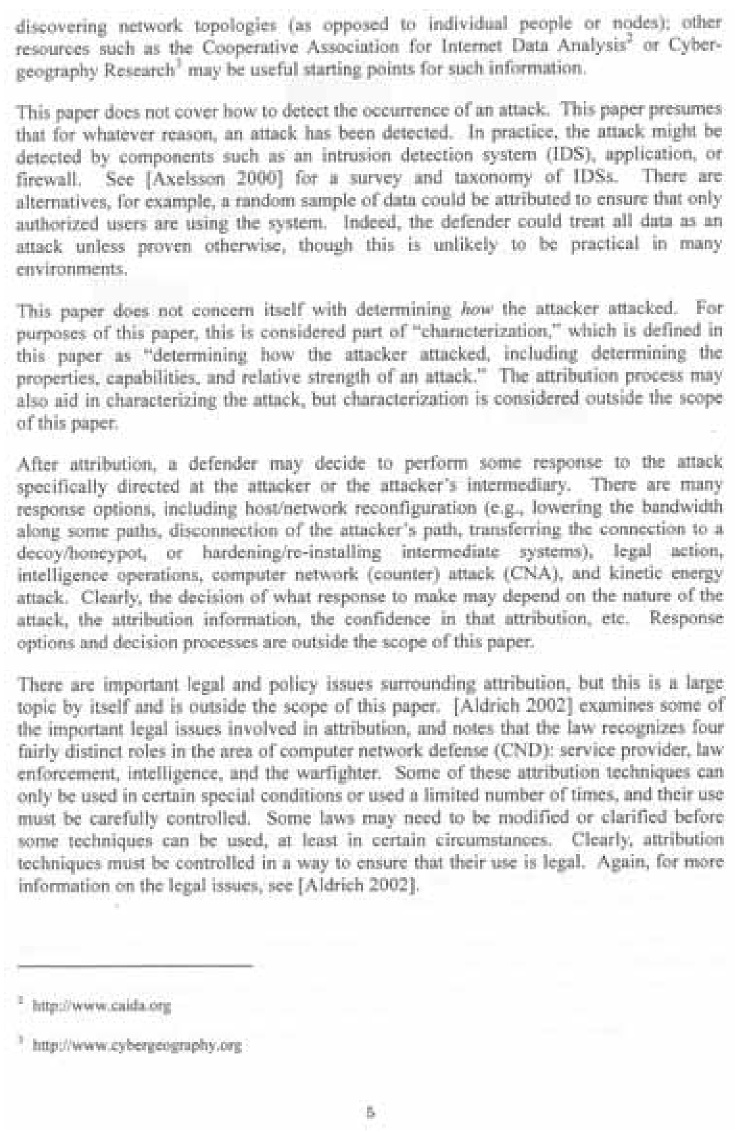

discovering network topologies (as opposed to individual people or nodes); other resources such as the Cooperative Association for Internet Data Analysis[2] or Cyber-geography Research[3] may be useful starting points for such information.

This paper does not cover how to detect the occurrence of an attack. This paper presumes that for whatever reason, an attack has been detected. In practice, the attack might be detected by components such as an intrusion detection system (IDS), application, or firewall. See [Axelsson 2000] for a survey and taxonomy of IDSs. There are alternatives, for example, a random sample of data could be attributed to ensure that only authorized users are using the system. Indeed, the defender could treat all data as an attack unless proven otherwise, though this is unlikely to be practical in many environments.

This paper does not concern itself with determining *how* the attacker attacked. For purposes of this paper, this is considered part of "characterization," which is defined in this paper as "determining how the attacker attacked, including determining the properties, capabilities, and relative strength of an attack." The attribution process may also aid in characterizing the attack, but characterization is considered outside the scope of this paper.

After attribution, a defender may decide to perform some response to the attack specifically directed at the attacker or the attacker's intermediary. There are many response options, including host/network reconfiguration (e.g., lowering the bandwidth along some paths, disconnection of the attacker's path, transferring the connection to a decoy/honeypot, or hardening/re-installing intermediate systems), legal action, intelligence operations, computer network (counter) attack (CNA), and kinetic energy attack. Clearly, the decision of what response to make may depend on the nature of the attack, the attribution information, the confidence in that attribution, etc. Response options and decision processes are outside the scope of this paper.

There are important legal and policy issues surrounding attribution, but this is a large topic by itself and is outside the scope of this paper. [Aldrich 2002] examines some of the important legal issues involved in attribution, and notes that the law recognizes four fairly distinct roles in the area of computer network defense (CND): service provider, law enforcement, intelligence, and the warfighter. Some of these attribution techniques can only be used in certain special conditions or used a limited number of times, and their use must be carefully controlled. Some laws may need to be modified or clarified before some techniques can be used, at least in certain circumstances. Clearly, attribution techniques must be controlled in a way to ensure that their use is legal. Again, for more information on the legal issues, see [Aldrich 2002].

[2] http://www.caida.org

[3] http://www.cybergeography.org

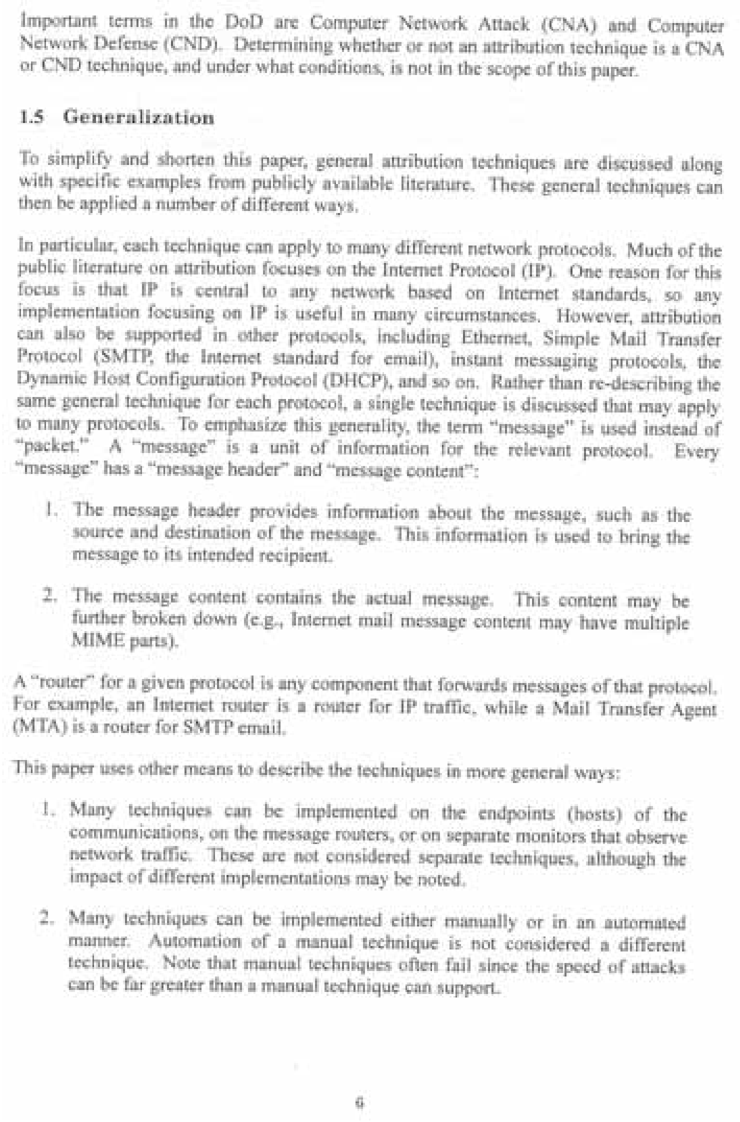

Important terms in the DoD are Computer Network Attack (CNA) and Computer Network Defense (CND). Determining whether or not an attribution technique is a CNA or CND technique, and under what conditions, is not in the scope of this paper.

1.5 Generalization

To simplify and shorten this paper, general attribution techniques are discussed along with specific examples from publicly available literature. These general techniques can then be applied a number of different ways.

In particular, each technique can apply to many different network protocols. Much of the public literature on attribution focuses on the Internet Protocol (IP). One reason for this focus is that IP is central to any network based on Internet standards, so any implementation focusing on IP is useful in many circumstances. However, attribution can also be supported in other protocols, including Ethernet, Simple Mail Transfer Protocol (SMTP, the Internet standard for email), instant messaging protocols, the Dynamic Host Configuration Protocol (DHCP), and so on. Rather than re-describing the same general technique for each protocol, a single technique is discussed that may apply to many protocols. To emphasize this generality, the term "message" is used instead of "packet." A "message" is a unit of information for the relevant protocol. Every "message" has a "message header" and "message content":

1. The message header provides information about the message, such as the source and destination of the message. This information is used to bring the message to its intended recipient.

2. The message content contains the actual message. This content may be further broken down (e.g., Internet mail message content may have multiple MIME parts).

A "router" for a given protocol is any component that forwards messages of that protocol. For example, an Internet router is a router for IP traffic, while a Mail Transfer Agent (MTA) is a router for SMTP email.

This paper uses other means to describe the techniques in more general ways:

1. Many techniques can be implemented on the endpoints (hosts) of the communications, on the message routers, or on separate monitors that observe network traffic. These are not considered separate techniques, although the impact of different implementations may be noted.

2. Many techniques can be implemented either manually or in an automated manner. Automation of a manual technique is not considered a different technique. Note that manual techniques often fail since the speed of attacks can be far greater than a manual technique can support.

3. Many techniques that involve querying can respond with either the information being requested, or simply store the response and respond with an index to that response. The advantage of the latter approach is that the information is stored, but authentication and authorization of the person requesting the attribution information can happen separately. Since such authentication and authorization may take a long time, but the data may disappear if not stored quickly, this approach can be valuable.

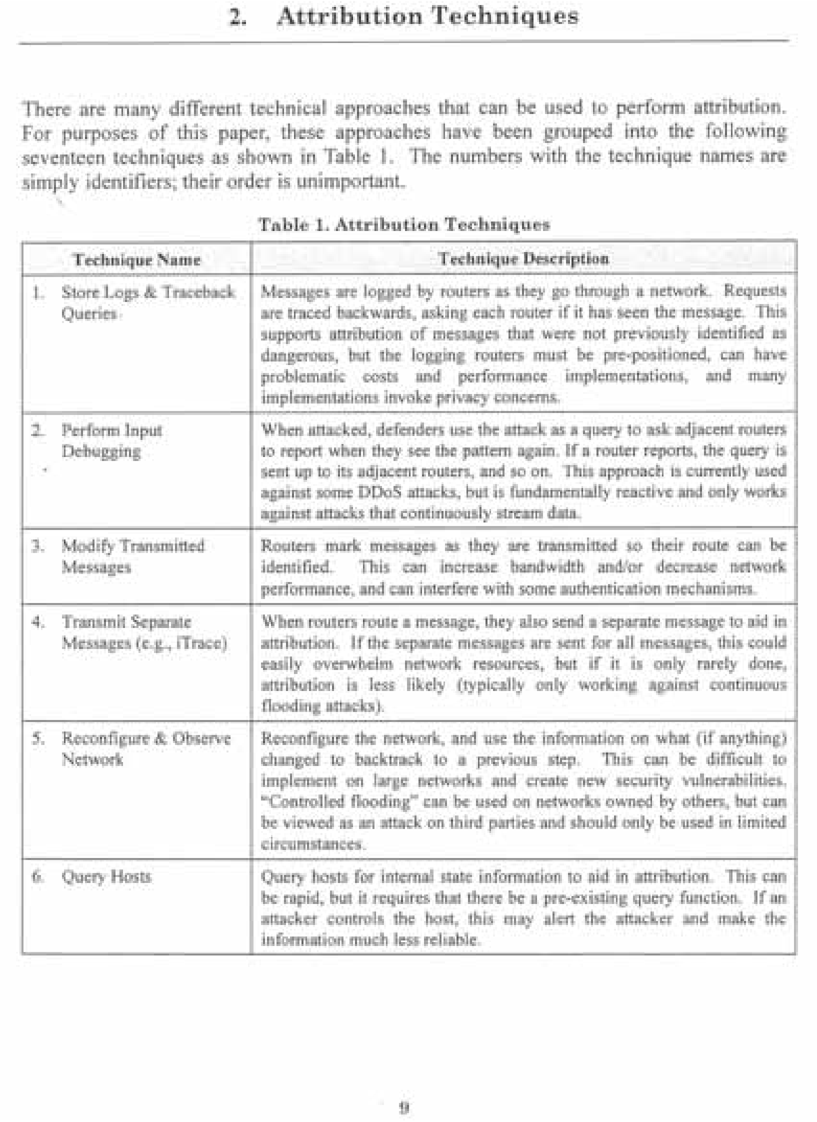

2. Attribution Techniques

There are many different technical approaches that can be used to perform attribution. For purposes of this paper, these approaches have been grouped into the following seventeen techniques as shown in Table 1. The numbers with the technique names are simply identifiers; their order is unimportant.

Table 1. Attribution Techniques

Technique Name	Technique Description
1. Store Logs & Traceback Queries	Messages are logged by routers as they go through a network. Requests are traced backwards, asking each router if it has seen the message. This supports attribution of messages that were not previously identified as dangerous, but the logging routers must be pre-positioned, can have problematic costs and performance implementations, and many implementations invoke privacy concerns.
2. Perform Input Debugging	When attacked, defenders use the attack as a query to ask adjacent routers to report when they see the pattern again. If a router reports, the query is sent up to its adjacent routers, and so on. This approach is currently used against some DDoS attacks, but is fundamentally reactive and only works against attacks that continuously stream data.
3. Modify Transmitted Messages	Routers mark messages as they are transmitted so their route can be identified. This can increase bandwidth and/or decrease network performance, and can interfere with some authentication mechanisms.
4. Transmit Separate Messages (e.g., iTrace)	When routers route a message, they also send a separate message to aid in attribution. If the separate messages are sent for all messages, this could easily overwhelm network resources, but if it is only rarely done, attribution is less likely (typically only working against continuous flooding attacks).
5. Reconfigure & Observe Network	Reconfigure the network, and use the information on what (if anything) changed to backtrack to a previous step. This can be difficult to implement on large networks and create new security vulnerabilities. "Controlled flooding" can be used on networks owned by others, but can be viewed as an attack on third parties and should only be used in limited circumstances.
6. Query Hosts	Query hosts for internal state information to aid in attribution. This can be rapid, but it requires that there be a pre-existing query function. If an attacker controls the host, this may alert the attacker and make the information much less reliable.

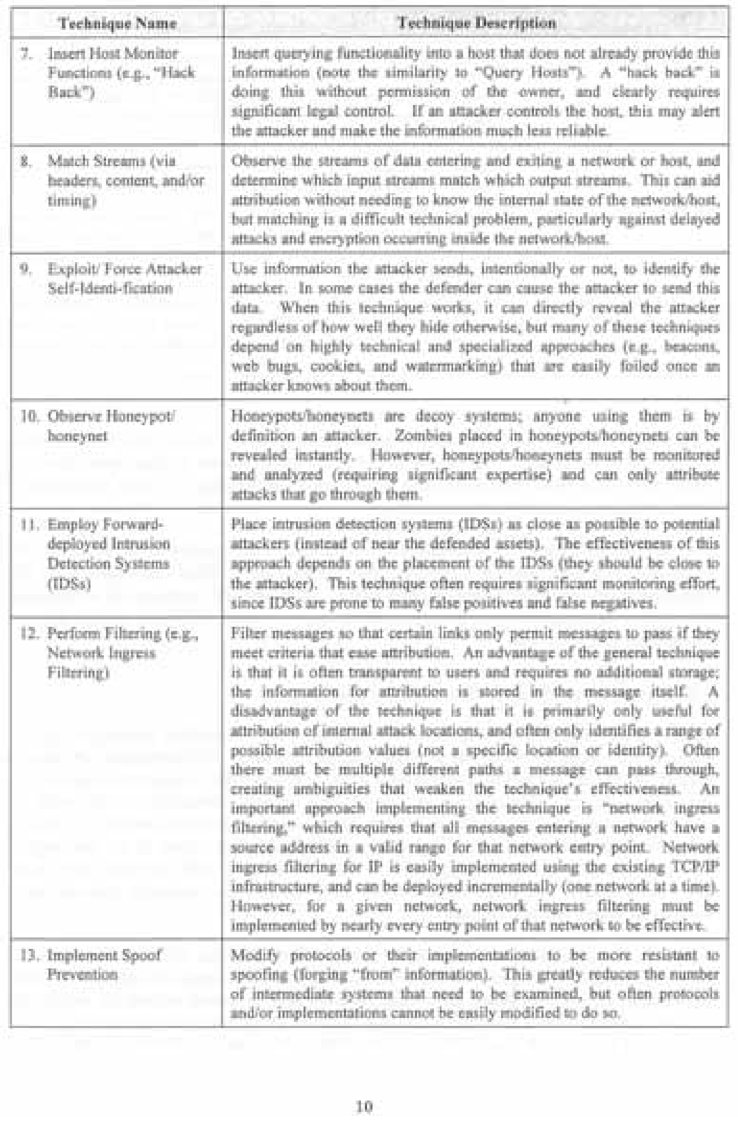

Technique Name	Technique Description
7. Insert Host Monitor Functions (e.g., "Hack Back")	Insert querying functionality into a host that does not already provide this information (note the similarity to "Query Hosts"). A "hack back" is doing this without permission of the owner, and clearly requires significant legal control. If an attacker controls the host, this may alert the attacker and make the information much less reliable.
8. Match Streams (via headers, content, and/or timing)	Observe the streams of data entering and exiting a network or host, and determine which input streams match which output streams. This can aid attribution without needing to know the internal state of the network/host, but matching is a difficult technical problem, particularly against delayed attacks and encryption occurring inside the network/host.
9. Exploit/ Force Attacker Self-Identi-fication	Use information the attacker sends, intentionally or not, to identify the attacker. In some cases the defender can cause the attacker to send this data. When this technique works, it can directly reveal the attacker regardless of how well they hide otherwise, but many of these techniques depend on highly technical and specialized approaches (e.g., beacons, web bugs, cookies, and watermarking) that are easily foiled once an attacker knows about them.
10. Observe Honeypot/ honeynet	Honeypots/honeynets are decoy systems; anyone using them is by definition an attacker. Zombies placed in honeypots/honeynets can be revealed instantly. However, honeypots/honeynets must be monitored and analyzed (requiring significant expertise) and can only attribute attacks that go through them.
11. Employ Forward-deployed Intrusion Detection Systems (IDSs)	Place intrusion detection systems (IDSs) as close as possible to potential attackers (instead of near the defended assets). The effectiveness of this approach depends on the placement of the IDSs (they should be close to the attacker). This technique often requires significant monitoring effort, since IDSs are prone to many false positives and false negatives.
12. Perform Filtering (e.g., Network Ingress Filtering)	Filter messages so that certain links only permit messages to pass if they meet criteria that ease attribution. An advantage of the general technique is that it is often transparent to users and requires no additional storage; the information for attribution is stored in the message itself. A disadvantage of the technique is that it is primarily only useful for attribution of internal attack locations, and often only identifies a range of possible attribution values (not a specific location or identity). Often there must be multiple different paths a message can pass through, creating ambiguities that weaken the technique's effectiveness. An important approach implementing the technique is "network ingress filtering," which requires that all messages entering a network have a source address in a valid range for that network entry point. Network ingress filtering for IP is easily implemented using the existing TCP/IP infrastructure, and can be deployed incrementally (one network at a time). However, for a given network, network ingress filtering must be implemented by nearly every entry point of that network to be effective.
13. Implement Spoof Prevention	Modify protocols or their implementations to be more resistant to spoofing (forging "from" information). This greatly reduces the number of intermediate systems that need to be examined, but often protocols and/or implementations cannot be easily modified to do so.

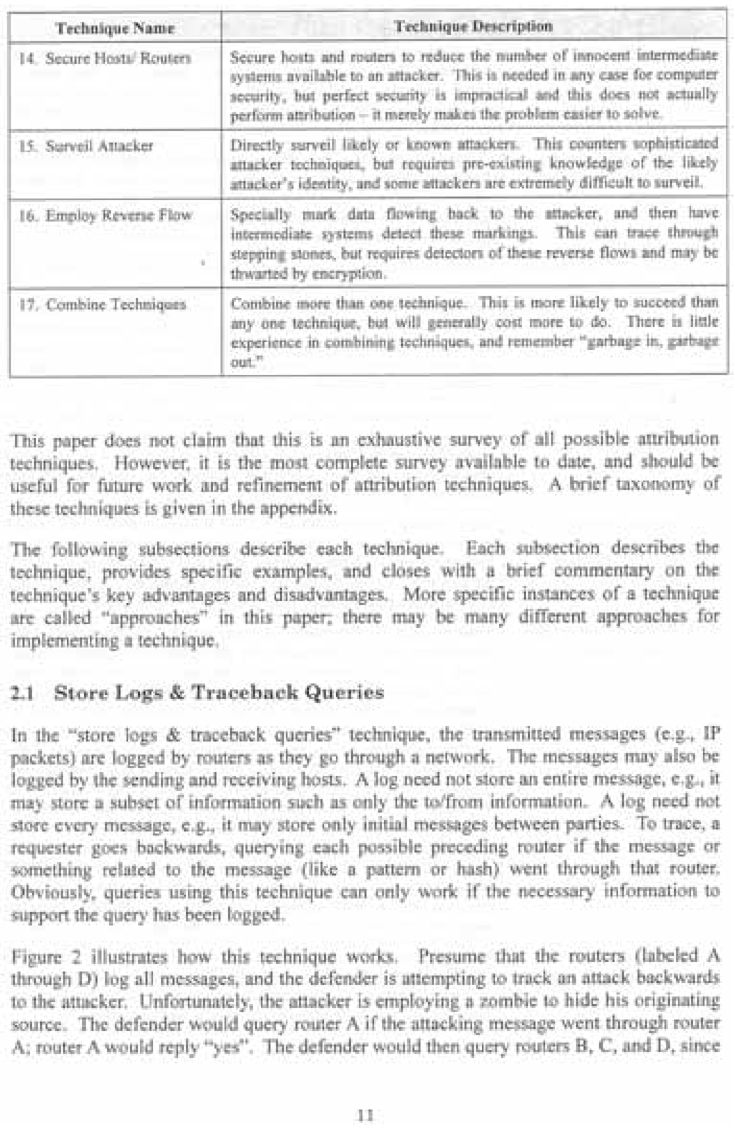

Technique Name	Technique Description
14. Secure Hosts/ Routers	Secure hosts and routers to reduce the number of innocent intermediate systems available to an attacker. This is needed in any case for computer security, but perfect security is impractical and this does not actually perform attribution – it merely makes the problem easier to solve.
15. Surveil Attacker	Directly surveil likely or known attackers. This counters sophisticated attacker techniques, but requires pre-existing knowledge of the likely attacker's identity, and some attackers are extremely difficult to surveil.
16. Employ Reverse Flow	Specially mark data flowing back to the attacker, and then have intermediate systems detect these markings. This can trace through stepping stones, but requires detectors of these reverse flows and may be thwarted by encryption.
17. Combine Techniques	Combine more than one technique. This is more likely to succeed than any one technique, but will generally cost more to do. There is little experience in combining techniques, and remember "garbage in, garbage out."

This paper does not claim that this is an exhaustive survey of all possible attribution techniques. However, it is the most complete survey available to date, and should be useful for future work and refinement of attribution techniques. A brief taxonomy of these techniques is given in the appendix.

The following subsections describe each technique. Each subsection describes the technique, provides specific examples, and closes with a brief commentary on the technique's key advantages and disadvantages. More specific instances of a technique are called "approaches" in this paper; there may be many different approaches for implementing a technique.

2.1 Store Logs & Traceback Queries

In the "store logs & traceback queries" technique, the transmitted messages (e.g., IP packets) are logged by routers as they go through a network. The messages may also be logged by the sending and receiving hosts. A log need not store an entire message, e.g., it may store a subset of information such as only the to/from information. A log need not store every message, e.g., it may store only initial messages between parties. To trace, a requester goes backwards, querying each possible preceding router if the message or something related to the message (like a pattern or hash) went through that router. Obviously, queries using this technique can only work if the necessary information to support the query has been logged.

Figure 2 illustrates how this technique works. Presume that the routers (labeled A through D) log all messages, and the defender is attempting to track an attack backwards to the attacker. Unfortunately, the attacker is employing a zombie to hide his originating source. The defender would query router A if the attacking message went through router A; router A would reply "yes". The defender would then query routers B, C, and D, since those routers are connected to the next network. Router B would reply "no", suggesting that the attack did not go through that route. Router D would also probably reply "no", since in most cases a zombie would change the message so that the connection could not be easily determined from a log. Router C would reply "yes", suggesting that the defender query further at that point. At this point, the defender has at least identified an intermediate network, and possibly the intermediate node on the network. The defender may even be able to backtrack further through the zombie, depending on the logged information.

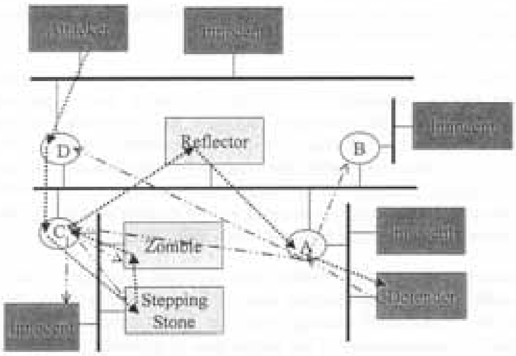

Figure 2. Store Logs & Traceback Queries Technique

This technique can be subdivided into two parts: logging and querying.

2.1.1 Logging

From the point of view of attribution, ideally every router would log every message and keep that log in perpetuity. In practice, this is undesirable: such logging may have unacceptable performance, storage space, or privacy implications. There are three ways to rectify this situation:

1. *Limit the number of messages logged.* For example, store only the data destined for an especially sensitive destination, or only packets that appear "suspicious."

2. *Limit the amount of data stored about each message.* For example, store only such as connection information (e.g., to/from information for only the initial message of a session), only to/from information from each message, only a

subset of the message (e.g., an initial fragment), only hashes of the message, or only hashes of a subset of the message.

3. *Accept the undesirable implications.* For example, buy large disk arrays or massive memory arrays to store the log data, buy faster processors (or more of them), and accept privacy risks. Clearly, limiting the number of messages logged and the amount of data stored about each message is more desirable where possible.

The criteria for selecting messages to be logged, or the amount of data to be logged in each message, could be changed dynamically. Dynamically changing these values could be accomplished by connecting these values with intrusion detection systems; see the forward-deployed IDS discussion below.

What data is stored, and how it can be retrieved, also has legal and privacy ramifications. In some circumstances, recording or retrieving information about a message (such as from/to information) may be considered different than recording or retrieving the message itself. In traditional telephone systems, it's possible to obtain (with a warrant) information on who a suspect has called (pen register) or who has received calls from the suspect (trap and trace). Warrants for pen register/trap and trace (PR/TT) information (as it's called by the law enforcement community) are often easier to acquire than recordings of the actual message traffic, which enjoys stronger legal protection.

It may be easier to store message hashes than messages themselves, since hashes are usually far smaller than the messages they hash. This is especially true in higher protocol levels, where one higher-level message may be implemented by many lower-level messages. Storing hashes instead of the actual data is probably more palatable legally and socially as well. Since the data itself isn't stored, hashing only supports attribution for data the requestor has already seen, and does not reveal the data itself. Storing hashes (instead of actual data) appears to be more akin to PR/TT data than to a recording of the message, though it is unclear if the courts will agree to this viewpoint.

One useful approach to logging events is logging authentication records for every authentication event (e.g., host login, ftp/http login, etc.), along with information such as the network address of the requestor. Reservations of resources (e.g., DHCP) are also useful for attribution. A log could be kept of every email received and sent (including the IP addresses of the other party exchanging mail, the from/to addresses, and a hash of the contents).

One problem with logging is protecting the logs themselves, particularly if an attacker gains administrative privileges over the system generating or storing the logs. A partial solution is to store logs on a separate machine from the machine performing the activities. As long as the logging system itself is secure, an attacker may be able to use other systems to append incorrect data but not remove correct data.

[Sager 1998] describes capturing flow information from Cisco routers that may be useful for attribution. This information includes (reported) source and destination IP addresses and ports, number of packets and bytes, IP protocol, and TCP flags.

One unusual solution for high volume IP packet-level logging is named Source Path Isolation Engine (SPIE). In the DARPA-sponsored SPIE approach, IP packet hashes are stored using a Bloom filter to store the information efficiently (see [Snoeren]). The SPIE approach dramatically increases the amount of data that can be logged, but even then, on high-rate routers (e.g., IP routers on an Internet backbone) this is still difficult and expensive for software-only implementations to perform in terms of memory size and memory performance. However, on lower-speed software routers (or where only a subset of packets are logged) this is not as difficult. [Sanchez 2001] describes the design of a hardware implementation of SPIE, and suggests that this would make SPIE practical at high speeds for relatively small amounts of money.[4] An Internet draft discussing traceback protocols, including SPIE, is available [Partridge 2001]. The developers of SPIE are working with the Internet Engineering Task Force (IETF) working group on traceback protocols and appear to be actively continuing the work.

Most of these approaches presume that logging is decentralized, and queries are made against the logs later. Logs are often decentralized, because the overhead of transmitting logs to a central location, and then performing analysis at a central location, does not scale well to very large networks. Nevertheless, there are those who have established centralized logging facilities for at least a portion of logging data, and then use the resulting information for attribution. For example, the Distributed Intrusion Detection System (DIDS) is a host-based approach that attempts to track all users in a network. Each monitored host sends abstracts of the audit trail to a centralized DIDS director for further analysis. Note that DIDS establishes a "Network-user ID" (NID), and audit records of session starts (logins) are sent to the central DIDS director. As a result, DIDS is able to track users moving through the network using normal logins when inside the DIDS-covered network [Snapp 1991a, 1991b, 1992, Ko].

A far more extreme version of this approach is the first approach suggested in [Arkin 2002]. In this approach, all public network traffic is logged for a period of time, and later the log of all network traffic is searched. First, a query for the network-visible pattern is requested, and then queries to find any transmission of the source code that caused the pattern are made. The presumption is that attacks tend to be tested at first in smaller regions, and so by identifying early attack tests it may be easier to identify the attacker. Searching for early versions of the source code may also aid in identifying the real attacker. Finally, profiles of attackers are built up (e.g., based on unusual approaches or rootkits), so attacks on different targets can be correlated to help identify the attacker. This approach depends on a massive storage system monitoring and logging data from

[4] Software that implements SPIE, and related papers, are available at http://www.ir.bbn.com/projects/SPIE. The software is open source software/free software under an MIT-style license. This software was developed for FreeBSD-4.3 and Linux-2.4.2, but it can monitor packets from arbitrary operating systems and the approach should easily apply to other operating systems.

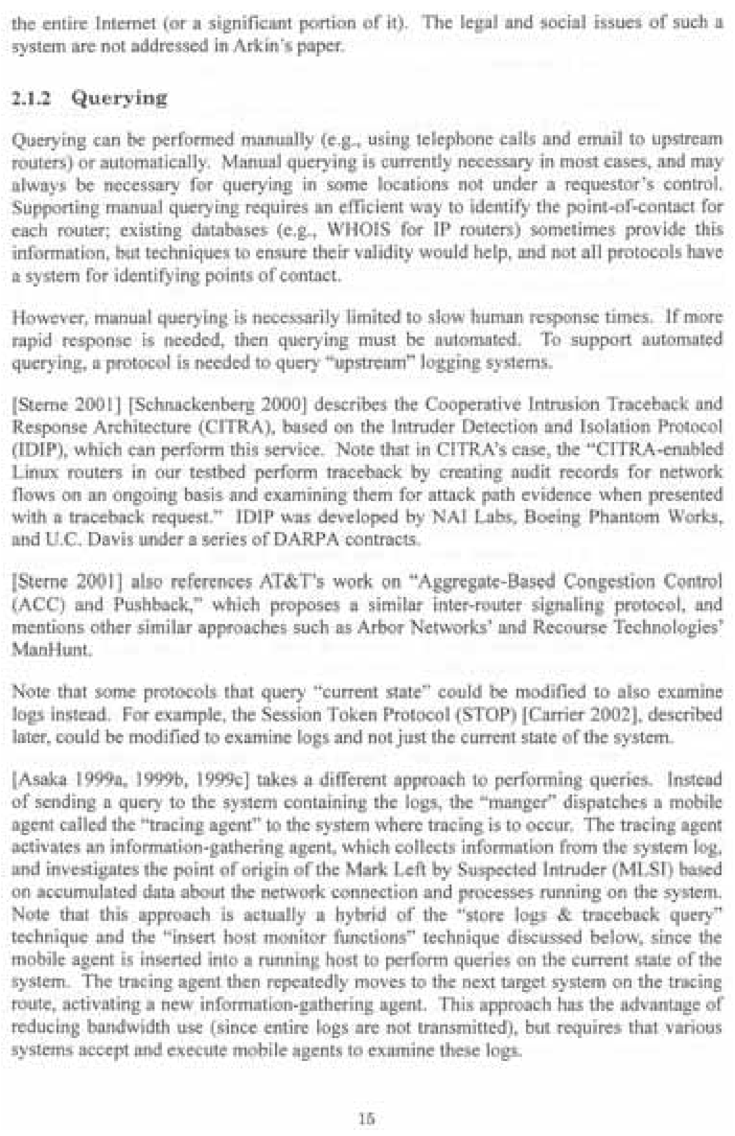

the entire Internet (or a significant portion of it). The legal and social issues of such a system are not addressed in Arkin's paper.

2.1.2 Querying

Querying can be performed manually (e.g., using telephone calls and email to upstream routers) or automatically. Manual querying is currently necessary in most cases, and may always be necessary for querying in some locations not under a requestor's control. Supporting manual querying requires an efficient way to identify the point-of-contact for each router; existing databases (e.g., WHOIS for IP routers) sometimes provide this information, but techniques to ensure their validity would help, and not all protocols have a system for identifying points of contact.

However, manual querying is necessarily limited to slow human response times. If more rapid response is needed, then querying must be automated. To support automated querying, a protocol is needed to query "upstream" logging systems.

[Sterne 2001] [Schnackenberg 2000] describes the Cooperative Intrusion Traceback and Response Architecture (CITRA), based on the Intruder Detection and Isolation Protocol (IDIP), which can perform this service. Note that in CITRA's case, the "CITRA-enabled Linux routers in our testbed perform traceback by creating audit records for network flows on an ongoing basis and examining them for attack path evidence when presented with a traceback request." IDIP was developed by NAI Labs, Boeing Phantom Works, and U.C. Davis under a series of DARPA contracts.

[Sterne 2001] also references AT&T's work on "Aggregate-Based Congestion Control (ACC) and Pushback," which proposes a similar inter-router signaling protocol, and mentions other similar approaches such as Arbor Networks' and Recourse Technologies' ManHunt.

Note that some protocols that query "current state" could be modified to also examine logs instead. For example, the Session Token Protocol (STOP) [Carrier 2002], described later, could be modified to examine logs and not just the current state of the system.

[Asaka 1999a, 1999b, 1999c] takes a different approach to performing queries. Instead of sending a query to the system containing the logs, the "manger" dispatches a mobile agent called the "tracing agent" to the system where tracing is to occur. The tracing agent activates an information-gathering agent, which collects information from the system log, and investigates the point of origin of the Mark Left by Suspected Intruder (MLSI) based on accumulated data about the network connection and processes running on the system. Note that this approach is actually a hybrid of the "store logs & traceback query" technique and the "insert host monitor functions" technique discussed below, since the mobile agent is inserted into a running host to perform queries on the current state of the system. The tracing agent then repeatedly moves to the next target system on the tracing route, activating a new information-gathering agent. This approach has the advantage of reducing bandwidth use (since entire logs are not transmitted), but requires that various systems accept and execute mobile agents to examine these logs.

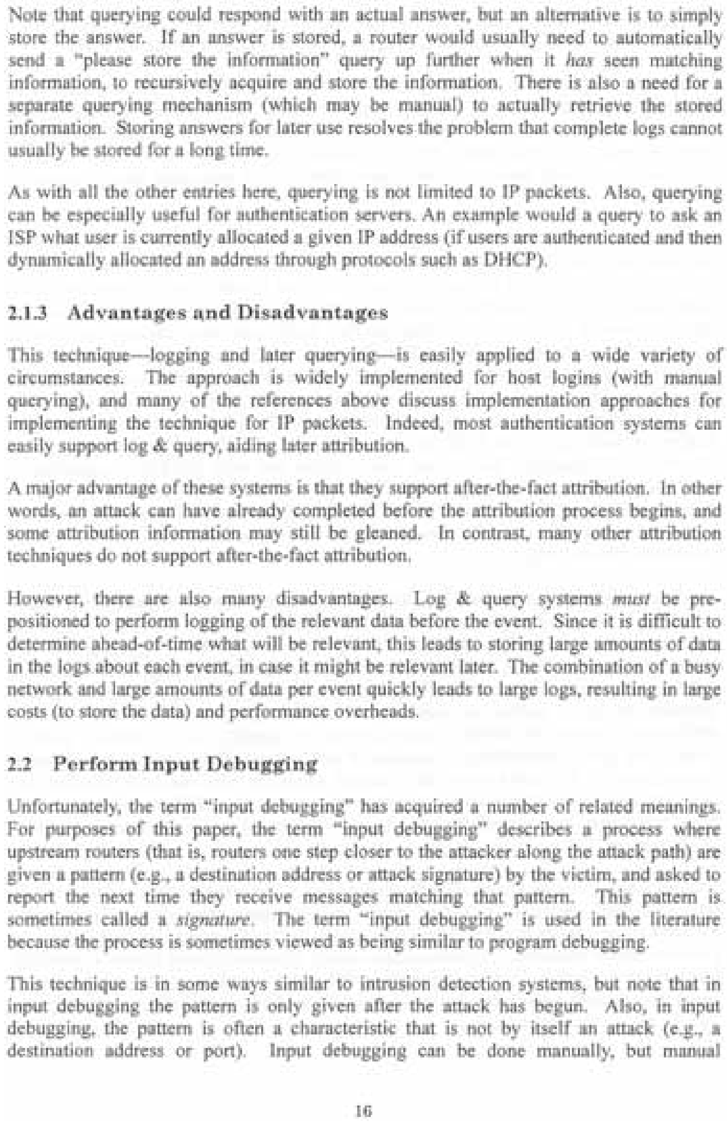

Note that querying could respond with an actual answer, but an alternative is to simply store the answer. If an answer is stored, a router would usually need to automatically send a "please store the information" query up further when it *has* seen matching information, to recursively acquire and store the information. There is also a need for a separate querying mechanism (which may be manual) to actually retrieve the stored information. Storing answers for later use resolves the problem that complete logs cannot usually be stored for a long time.

As with all the other entries here, querying is not limited to IP packets. Also, querying can be especially useful for authentication servers. An example would a query to ask an ISP what user is currently allocated a given IP address (if users are authenticated and then dynamically allocated an address through protocols such as DHCP).

2.1.3 Advantages and Disadvantages

This technique—logging and later querying—is easily applied to a wide variety of circumstances. The approach is widely implemented for host logins (with manual querying), and many of the references above discuss implementation approaches for implementing the technique for IP packets. Indeed, most authentication systems can easily support log & query, aiding later attribution.

A major advantage of these systems is that they support after-the-fact attribution. In other words, an attack can have already completed before the attribution process begins, and some attribution information may still be gleaned. In contrast, many other attribution techniques do not support after-the-fact attribution.

However, there are also many disadvantages. Log & query systems *must* be pre-positioned to perform logging of the relevant data before the event. Since it is difficult to determine ahead-of-time what will be relevant, this leads to storing large amounts of data in the logs about each event, in case it might be relevant later. The combination of a busy network and large amounts of data per event quickly leads to large logs, resulting in large costs (to store the data) and performance overheads.

2.2 Perform Input Debugging

Unfortunately, the term "input debugging" has acquired a number of related meanings. For purposes of this paper, the term "input debugging" describes a process where upstream routers (that is, routers one step closer to the attacker along the attack path) are given a pattern (e.g., a destination address or attack signature) by the victim, and asked to report the next time they receive messages matching that pattern. This pattern is sometimes called a *signature*. The term "input debugging" is used in the literature because the process is sometimes viewed as being similar to program debugging.

This technique is in some ways similar to intrusion detection systems, but note that in input debugging the pattern is only given after the attack has begun. Also, in input debugging, the pattern is often a characteristic that is not by itself an attack (e.g., a destination address or port). Input debugging can be done manually, but manual

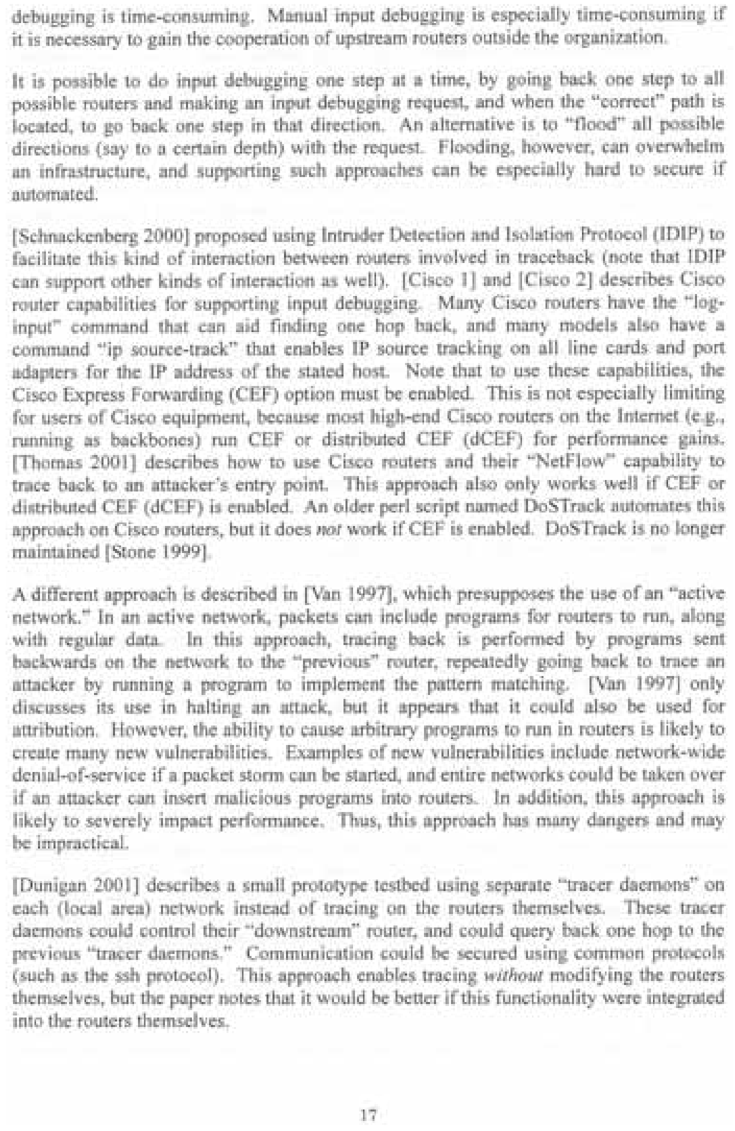

debugging is time-consuming. Manual input debugging is especially time-consuming if it is necessary to gain the cooperation of upstream routers outside the organization.

It is possible to do input debugging one step at a time, by going back one step to all possible routers and making an input debugging request, and when the "correct" path is located, to go back one step in that direction. An alternative is to "flood" all possible directions (say to a certain depth) with the request. Flooding, however, can overwhelm an infrastructure, and supporting such approaches can be especially hard to secure if automated.

[Schnackenberg 2000] proposed using Intruder Detection and Isolation Protocol (IDIP) to facilitate this kind of interaction between routers involved in traceback (note that IDIP can support other kinds of interaction as well). [Cisco 1] and [Cisco 2] describes Cisco router capabilities for supporting input debugging. Many Cisco routers have the "log-input" command that can aid finding one hop back, and many models also have a command "ip source-track" that enables IP source tracking on all line cards and port adapters for the IP address of the stated host. Note that to use these capabilities, the Cisco Express Forwarding (CEF) option must be enabled. This is not especially limiting for users of Cisco equipment, because most high-end Cisco routers on the Internet (e.g., running as backbones) run CEF or distributed CEF (dCEF) for performance gains. [Thomas 2001] describes how to use Cisco routers and their "NetFlow" capability to trace back to an attacker's entry point. This approach also only works well if CEF or distributed CEF (dCEF) is enabled. An older perl script named DoSTrack automates this approach on Cisco routers, but it does *not* work if CEF is enabled. DoSTrack is no longer maintained [Stone 1999].

A different approach is described in [Van 1997], which presupposes the use of an "active network." In an active network, packets can include programs for routers to run, along with regular data. In this approach, tracing back is performed by programs sent backwards on the network to the "previous" router, repeatedly going back to trace an attacker by running a program to implement the pattern matching. [Van 1997] only discusses its use in halting an attack, but it appears that it could also be used for attribution. However, the ability to cause arbitrary programs to run in routers is likely to create many new vulnerabilities. Examples of new vulnerabilities include network-wide denial-of-service if a packet storm can be started, and entire networks could be taken over if an attacker can insert malicious programs into routers. In addition, this approach is likely to severely impact performance. Thus, this approach has many dangers and may be impractical.

[Dunigan 2001] describes a small prototype testbed using separate "tracer daemons" on each (local area) network instead of tracing on the routers themselves. These tracer daemons could control their "downstream" router, and could query back one hop to the previous "tracer daemons." Communication could be secured using common protocols (such as the ssh protocol). This approach enables tracing *without* modifying the routers themselves, but the paper notes that it would be better if this functionality were integrated into the routers themselves.

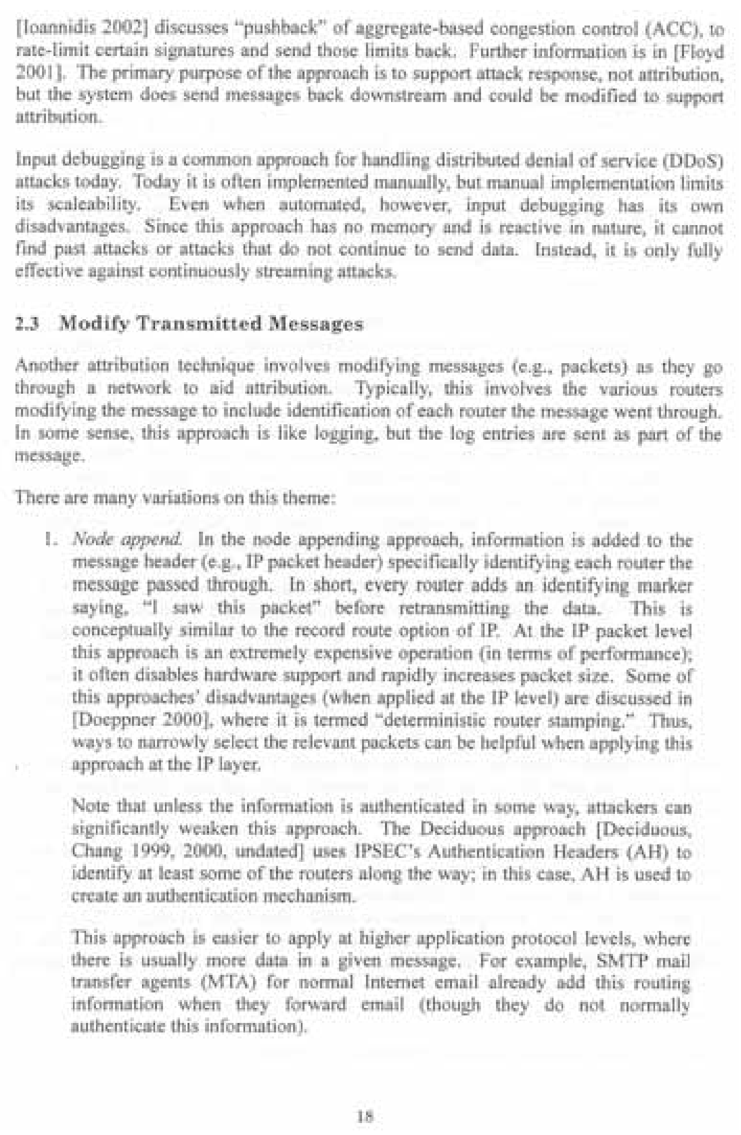

[Ioannidis 2002] discusses "pushback" of aggregate-based congestion control (ACC), to rate-limit certain signatures and send those limits back. Further information is in [Floyd 2001]. The primary purpose of the approach is to support attack response, not attribution, but the system does send messages back downstream and could be modified to support attribution.

Input debugging is a common approach for handling distributed denial of service (DDoS) attacks today. Today it is often implemented manually, but manual implementation limits its scaleability. Even when automated, however, input debugging has its own disadvantages. Since this approach has no memory and is reactive in nature, it cannot find past attacks or attacks that do not continue to send data. Instead, it is only fully effective against continuously streaming attacks.

2.3 Modify Transmitted Messages

Another attribution technique involves modifying messages (e.g., packets) as they go through a network to aid attribution. Typically, this involves the various routers modifying the message to include identification of each router the message went through. In some sense, this approach is like logging, but the log entries are sent as part of the message.

There are many variations on this theme:

1. *Node append.* In the node appending approach, information is added to the message header (e.g., IP packet header) specifically identifying each router the message passed through. In short, every router adds an identifying marker saying, "I saw this packet" before retransmitting the data. This is conceptually similar to the record route option of IP. At the IP packet level this approach is an extremely expensive operation (in terms of performance); it often disables hardware support and rapidly increases packet size. Some of this approaches' disadvantages (when applied at the IP level) are discussed in [Doeppner 2000], where it is termed "deterministic router stamping." Thus, ways to narrowly select the relevant packets can be helpful when applying this approach at the IP layer.

   Note that unless the information is authenticated in some way, attackers can significantly weaken this approach. The Deciduous approach [Deciduous, Chang 1999, 2000, undated] uses IPSEC's Authentication Headers (AH) to identify at least some of the routers along the way; in this case, AH is used to create an authentication mechanism.

   This approach is easier to apply at higher application protocol levels, where there is usually more data in a given message. For example, SMTP mail transfer agents (MTA) for normal Internet email already add this routing information when they forward email (though they do not normally authenticate this information).

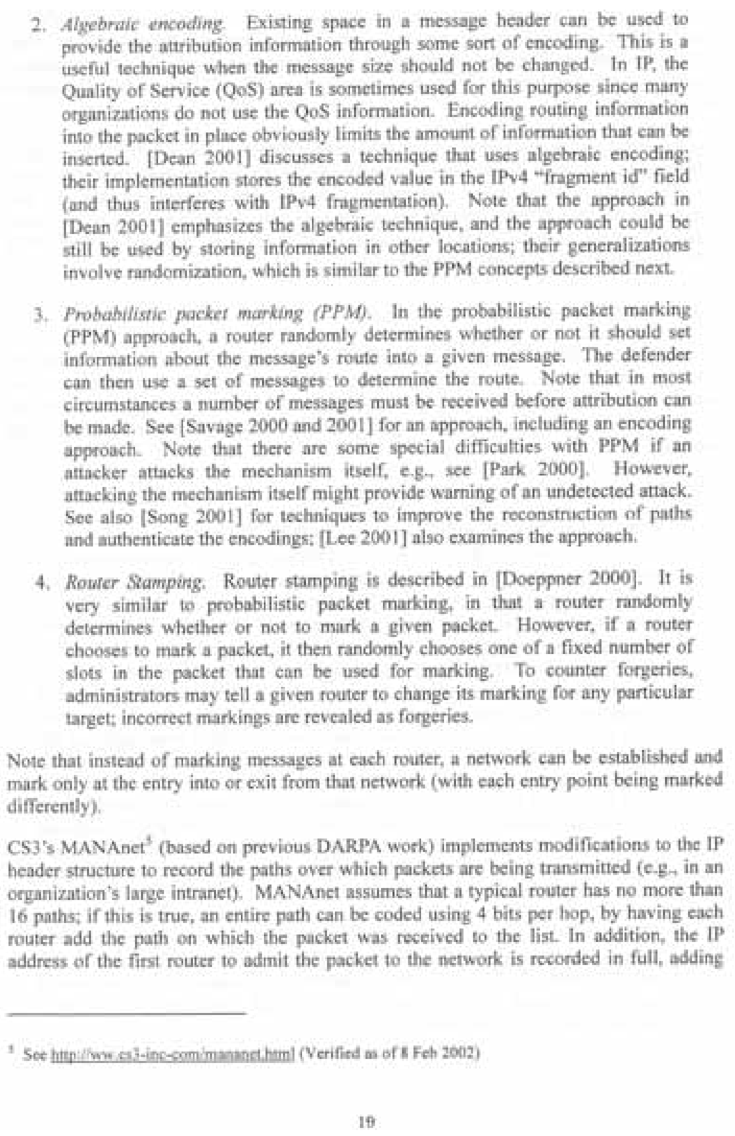

2. *Algebraic encoding.* Existing space in a message header can be used to provide the attribution information through some sort of encoding. This is a useful technique when the message size should not be changed. In IP, the Quality of Service (QoS) area is sometimes used for this purpose since many organizations do not use the QoS information. Encoding routing information into the packet in place obviously limits the amount of information that can be inserted. [Dean 2001] discusses a technique that uses algebraic encoding; their implementation stores the encoded value in the IPv4 "fragment id" field (and thus interferes with IPv4 fragmentation). Note that the approach in [Dean 2001] emphasizes the algebraic technique, and the approach could be still be used by storing information in other locations; their generalizations involve randomization, which is similar to the PPM concepts described next.

3. *Probabilistic packet marking (PPM).* In the probabilistic packet marking (PPM) approach, a router randomly determines whether or not it should set information about the message's route into a given message. The defender can then use a set of messages to determine the route. Note that in most circumstances a number of messages must be received before attribution can be made. See [Savage 2000 and 2001] for an approach, including an encoding approach. Note that there are some special difficulties with PPM if an attacker attacks the mechanism itself, e.g., see [Park 2000]. However, attacking the mechanism itself might provide warning of an undetected attack. See also [Song 2001] for techniques to improve the reconstruction of paths and authenticate the encodings; [Lee 2001] also examines the approach.

4. *Router Stamping.* Router stamping is described in [Doeppner 2000]. It is very similar to probabilistic packet marking, in that a router randomly determines whether or not to mark a given packet. However, if a router chooses to mark a packet, it then randomly chooses one of a fixed number of slots in the packet that can be used for marking. To counter forgeries, administrators may tell a given router to change its marking for any particular target; incorrect markings are revealed as forgeries.

Note that instead of marking messages at each router, a network can be established and mark only at the entry into or exit from that network (with each entry point being marked differently).

CS3's MANAnet[5] (based on previous DARPA work) implements modifications to the IP header structure to record the paths over which packets are being transmitted (e.g., in an organization's large intranet). MANAnet assumes that a typical router has no more than 16 paths; if this is true, an entire path can be coded using 4 bits per hop, by having each router add the path on which the packet was received to the list. In addition, the IP address of the first router to admit the packet to the network is recorded in full, adding

[5] See http://ww.cs3-inc-com/mananet.html (Verified as of 8 Feb 2002)

another 4 bytes (in IPv4). If the maximum number of hops to be expected to be less than 32, the entire path can be recorded as a 20 byte expansion of the header. For the purposes of attribution, the main component is the "MANAnet Router" which implements a proprietary modified form of IP called Path Enhanced IP (PEIP); each MANAnet Router adds the place from where it got a packet to the packet before it forwards it. The approach is reviewed by [Dietrich], who notes a number of problems with the approach. This includes near-universal imposition of ingress filtering before it can even be deployed, modification of IP headers with support by all routers, a resource forecasting or reservation mechanism, and a means for determining the real source of packets. In particular, it questions whether the computational techniques can be implemented in the fast paths of backbone routers, as well as concern for a proprietary modified form of the open IP standard.

This technique has the advantage of not requiring separate logs and a separate query system, and thus attribution information is immediately available. Also, there is no need to store and manage logs.

The technique also has many disadvantages. One obvious disadvantage is that this approach can greatly increase bandwidth requirements and/or severely reduce performance. This is not only due to the increased data and processing requirements; these approaches may disable hardware and software optimizations employed by some routers. There are many situations where changing a message is impractical, for example, it may defeat authentication mechanisms since those mechanisms may detect an "unauthorized" change to the message. Ensuring that this data is correct (and not from an attacker) is also difficult. Note that this approach requires the cooperation of remote routers to add the necessary information, implying the need for a standard way to insert this data.

2.4 Transmit Separate Messages (e.g., iTrace)

An alternative similar to modifying transmitted messages is to have routers send separate messages that can be used to support attribution.

The IETF's ICMP Traceback Working Group has been working on one such approach called "iTrace," which sends traceback messages a small amount of the time (e.g., 1/20,000 packets); their website is at http://www.ietf.org/html.charters/itrace-charter.html. See [Bellovin 2000] where the iTrace (ICMP traceback messages) are discussed.

[Mankin 2001] describes an "intention-driven" modification to iTrace; an iTrace message is only sent towards a path that has registered an "interest", and that interest information is shared using a minor modification of the standard Border Gateway Protocol (BGP). Indeed, the paper argues that the simple original iTrace proposal was essentially ineffective without this extension, although this modification requires more effort to implement. [Mankin 2001] also experimented with a heuristic that preferred longer paths, to increase the probability that more useful path information was sent, and the results appeared very encouraging and to improve the usefulness still further. See also [Wu 2001] for more information.

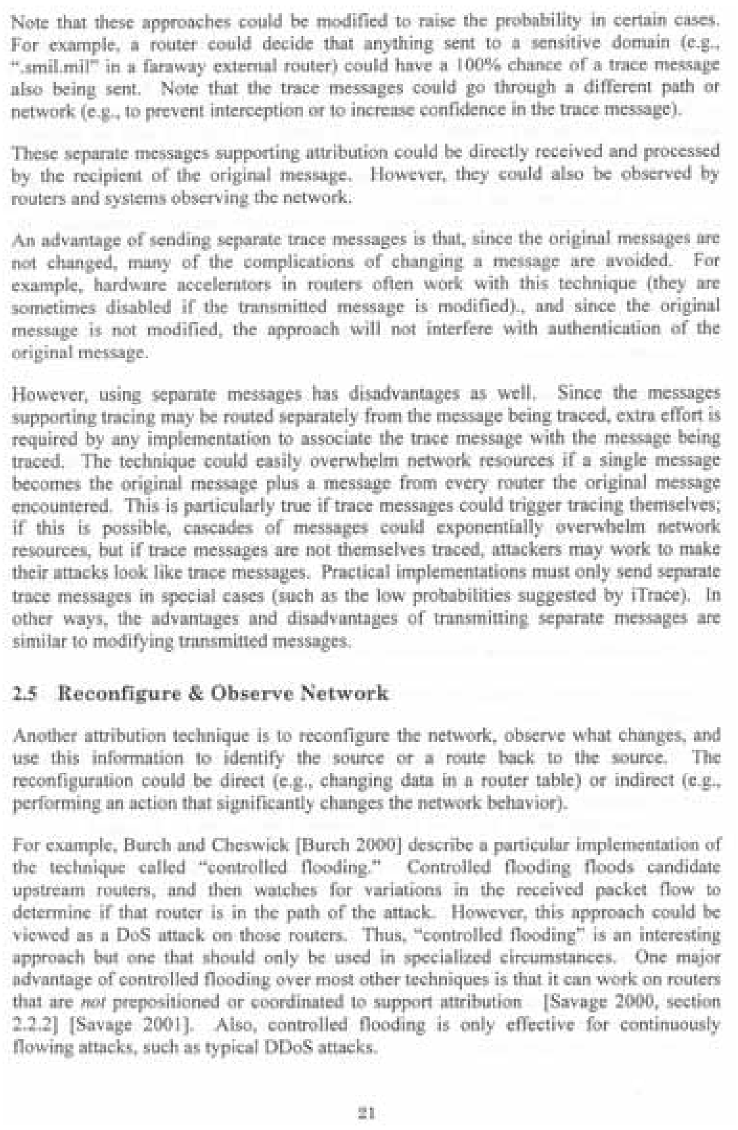

Note that these approaches could be modified to raise the probability in certain cases. For example, a router could decide that anything sent to a sensitive domain (e.g., ".smil.mil" in a faraway external router) could have a 100% chance of a trace message also being sent. Note that the trace messages could go through a different path or network (e.g., to prevent interception or to increase confidence in the trace message).

These separate messages supporting attribution could be directly received and processed by the recipient of the original message. However, they could also be observed by routers and systems observing the network.

An advantage of sending separate trace messages is that, since the original messages are not changed, many of the complications of changing a message are avoided. For example, hardware accelerators in routers often work with this technique (they are sometimes disabled if the transmitted message is modified)., and since the original message is not modified, the approach will not interfere with authentication of the original message.

However, using separate messages has disadvantages as well. Since the messages supporting tracing may be routed separately from the message being traced, extra effort is required by any implementation to associate the trace message with the message being traced. The technique could easily overwhelm network resources if a single message becomes the original message plus a message from every router the original message encountered. This is particularly true if trace messages could trigger tracing themselves; if this is possible, cascades of messages could exponentially overwhelm network resources, but if trace messages are not themselves traced, attackers may work to make their attacks look like trace messages. Practical implementations must only send separate trace messages in special cases (such as the low probabilities suggested by iTrace). In other ways, the advantages and disadvantages of transmitting separate messages are similar to modifying transmitted messages.

2.5 Reconfigure & Observe Network

Another attribution technique is to reconfigure the network, observe what changes, and use this information to identify the source or a route back to the source. The reconfiguration could be direct (e.g., changing data in a router table) or indirect (e.g., performing an action that significantly changes the network behavior).

For example, Burch and Cheswick [Burch 2000] describe a particular implementation of the technique called "controlled flooding." Controlled flooding floods candidate upstream routers, and then watches for variations in the received packet flow to determine if that router is in the path of the attack. However, this approach could be viewed as a DoS attack on those routers. Thus, "controlled flooding" is an interesting approach but one that should only be used in specialized circumstances. One major advantage of controlled flooding over most other techniques is that it can work on routers that are *not* prepositioned or coordinated to support attribution [Savage 2000, section 2.2.2] [Savage 2001]. Also, controlled flooding is only effective for continuously flowing attacks, such as typical DDoS attacks.

Another variation on this theme is Centertrack as described by [Stone 1999]. In this approach, an overlay network is created that links all ISP edge routers to a central tracking router (or network of tracking routers). Dynamic routing is then used to redirect all packets destined for the victim so that they will then go through the tracking routers; then hop-by-hop approaches are used to find the source. Centertrack is identified as a "defunct research project" in [Dittrich].

Active networks also fit into this category. Active networks permit programs to be sent to network infrastructure components (e.g., routers), which then run and change the network behavior. Once example is [Sterne 2002], which uses the secure ANTS execution environment. The paper also discusses many other active network activities and related papers.

An advantage of this technique is that, if the network being reconfigured is large, it may be possible to very rapidly identify the attacking source. However, the technique also has disadvantages. Direct control over a large network can be difficult to implement, and can create its own security vulnerabilities. Indirect approaches (such as controlled flooding) can be viewed as an attack, and thus should only be used in limited situations.

2.6 Query Hosts

Another attribution technique is to query hosts about their state. This includes using existing services to perform such queries, or adding new services to hosts to support such queries.

An authorized administrator can use existing host tools (such as standard administrative tools) to monitor an attacker in certain circumstances. These functions may not succeed if the attacker controls the host, because the attacker may subvert the tools used to do the monitoring. This is particularly a problem for administrative tools normally on the host; an intelligent adversary is likely to know of such tools and take steps to subvert those them.

[Carrier 2002] describes Session Token Protocol (STOP), an approach for traceback *through* a host toward its sender. The STOP approach extends the standard "ident" service to permit external entities to query about active processes. Once a query is sent, the system then examines all incoming connections to that set of processes, and can recursively trace back to previous hosts. STOP has the nice property of being an upward-compatible extension to an existing service and is able to rapidly backtrack to ordinary TCP/IP connections. STOP has been implemented, but only for Unix-like systems (there is no known reason the technique could not work for other systems such as Windows). STOP intentionally stores information about the traceback, and merely returns a hash sufficient to request the real data at a future time; this is an effective technique for supporting privacy, but this approach does limit its capabilities (the requestor cannot control the recursive traceback, so an attacker may thwart simple attempts to trace back, and an attacker might easily be alerted that a traceback is occurring). It should be easy to extend STOP to immediately report its data to a trusted third party (such as a CERT), presumably using authentication and encryption. When we contacted the developers of

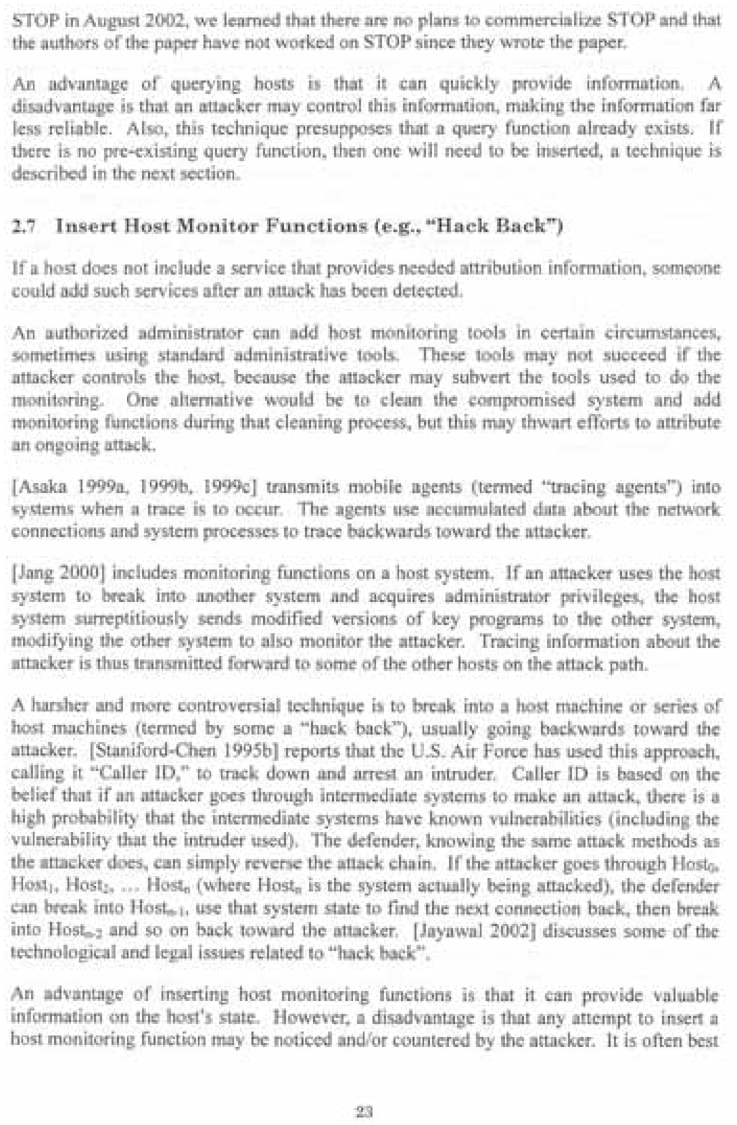

STOP in August 2002, we learned that there are no plans to commercialize STOP and that the authors of the paper have not worked on STOP since they wrote the paper.

An advantage of querying hosts is that it can quickly provide information. A disadvantage is that an attacker may control this information, making the information far less reliable. Also, this technique presupposes that a query function already exists. If there is no pre-existing query function, then one will need to be inserted, a technique is described in the next section.

2.7 Insert Host Monitor Functions (e.g., "Hack Back")

If a host does not include a service that provides needed attribution information, someone could add such services after an attack has been detected.

An authorized administrator can add host monitoring tools in certain circumstances, sometimes using standard administrative tools. These tools may not succeed if the attacker controls the host, because the attacker may subvert the tools used to do the monitoring. One alternative would be to clean the compromised system and add monitoring functions during that cleaning process, but this may thwart efforts to attribute an ongoing attack.

[Asaka 1999a, 1999b, 1999c] transmits mobile agents (termed "tracing agents") into systems when a trace is to occur. The agents use accumulated data about the network connections and system processes to trace backwards toward the attacker.

[Jang 2000] includes monitoring functions on a host system. If an attacker uses the host system to break into another system and acquires administrator privileges, the host system surreptitiously sends modified versions of key programs to the other system, modifying the other system to also monitor the attacker. Tracing information about the attacker is thus transmitted forward to some of the other hosts on the attack path.

A harsher and more controversial technique is to break into a host machine or series of host machines (termed by some a "hack back"), usually going backwards toward the attacker. [Staniford-Chen 1995b] reports that the U.S. Air Force has used this approach, calling it "Caller ID," to track down and arrest an intruder. Caller ID is based on the belief that if an attacker goes through intermediate systems to make an attack, there is a high probability that the intermediate systems have known vulnerabilities (including the vulnerability that the intruder used). The defender, knowing the same attack methods as the attacker does, can simply reverse the attack chain. If the attacker goes through $Host_0$, $Host_1$, $Host_2$, … $Host_n$ (where $Host_n$ is the system actually being attacked), the defender can break into $Host_{n-1}$, use that system state to find the next connection back, then break into $Host_{n-2}$ and so on back toward the attacker. [Jayawal 2002] discusses some of the technological and legal issues related to "hack back".

An advantage of inserting host monitoring functions is that it can provide valuable information on the host's state. However, a disadvantage is that any attempt to insert a host monitoring function may be noticed and/or countered by the attacker. It is often best if the attacker would not notice the additional monitors when they are inserted, but it can be a challenging task to add monitors without revealing them to a capable attacker.

The "hack back" approach has many additional disadvantages. Fundamentally this involves a number of complex legal issues. [Staniford-Chen 1995b] reports that performing this activity required special permission from the Department of Justice. It is an extreme measure with many social issues, such as privacy concerns. This is especially true if the counter-attack is performed by anyone other than the host owner or authorized administrator. There is also the possibility that an intermediate system cannot be broken into (perhaps the attacker used an attack not known to the defender, or the attacker improved the security of the systems they broke into). There is also the danger of accidentally damaging intermediate systems. If hack back is implemented manually, the attacker may be gone long before the attribution can succeed. If hack back is automated, there is the danger of going awry (either by programming error or by malicious misleading data provided by an attacker). In short, hack back is an approach with a large number of important disadvantages [Staniford-Chen 1995b].

2.8 Match Streams (via Headers, Content, and/or Timing)

Instead of trying to trace through every router and host in a network, it may be possible to observe the set of streams entering and exiting some network or system and determine, externally, which input streams match which output streams. This technique is often referred to as "stream matching."

Figure 3 illustrates this technique. In this example, there are a number of data flows (A through F) that enter or leave through a few specific links or ports of a network or host. The goal is to use externally-visible information about those flows to determine which incoming flows match which outgoing flows. In this example, flow A enters the network or host and re-emerges as flow E; flow D enters and re-emerges as flow F. Not all flows need match; flow B enters but does not leave, while flow C originates from the host or network (and has no corresponding entering flow).

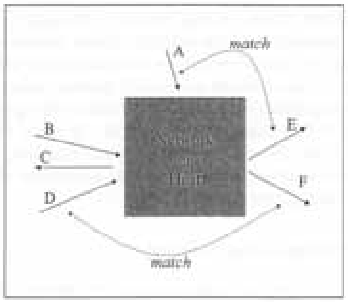

Figure 3. Stream Matching

Stream matching techniques can be further divided into techniques that examine message headers, message content, and message timing. Combinations of all three approaches are possible as well.

2.8.1 Stream Matching using Message Headers

In this approach, the headers of messages entering and exiting a network or host are examined to determine which incoming streams match which outgoing streams, and thus determine the source of a stream being examined.

[Yoda 2000] uses time stamps and headers of the packets, in particular the increase in sequence codes, to determine if one flow "matches" another flow. This approach is more effective if the intruder is manually inputting commands by hand (instead of transferring large files). [Yoda 2000] Only describes how to use the approach to match telnet or rlogin flows. The TCP sequence numbers are determined by data content length, so Yoda's approach could be considered to involve examining message content as well. Yoda's approach will not work if compression of the data stream occurs differently on each part of the chain, nor will it work well if link encryption is used (because the TCP headers will be encrypted).

2.8.2 Stream Matching using Data Content

In this approach, input and output data content of streams are examined to see if they match. Note that this approach will usually fail if the stream is encrypted "inside" the network/host, since the data content will be encrypted into a different value.

[Staniford-Chen 1995a, 1995b] developed a technique called "thumbprinting" that divides the stream data into discrete time intervals, creates digests of the packets within the interval, and computes the similarity of stream digests to determine if one stream matches another. [Buchholtz] reimplemented the approach and comments on it. Initial results appeared very promising.

[Mansfield 2000] discusses characterizing intrusion attempts by first identifying some specific noise or other indications of the attempt. These characterizations might include TCP-RESET packets, ICMP echo-response, or destination/port unreachable packets. The observed traffic pattern is then matched with traffic patterns collected from the connected network links. The communication between the various managers and agents is carried out using the standard SNMP management protocol. It could be argued that these are forward-deployed IDSs, but the approach emphasizes matching a particular stream of actions to an attack closer to the defender. Based the definitions used in this paper, this is considered a stream matching approach.

2.8.3 Stream Matching using Timing

In this approach, timing is used to match streams, by determining if the timing of the streams suggests a causal relation.

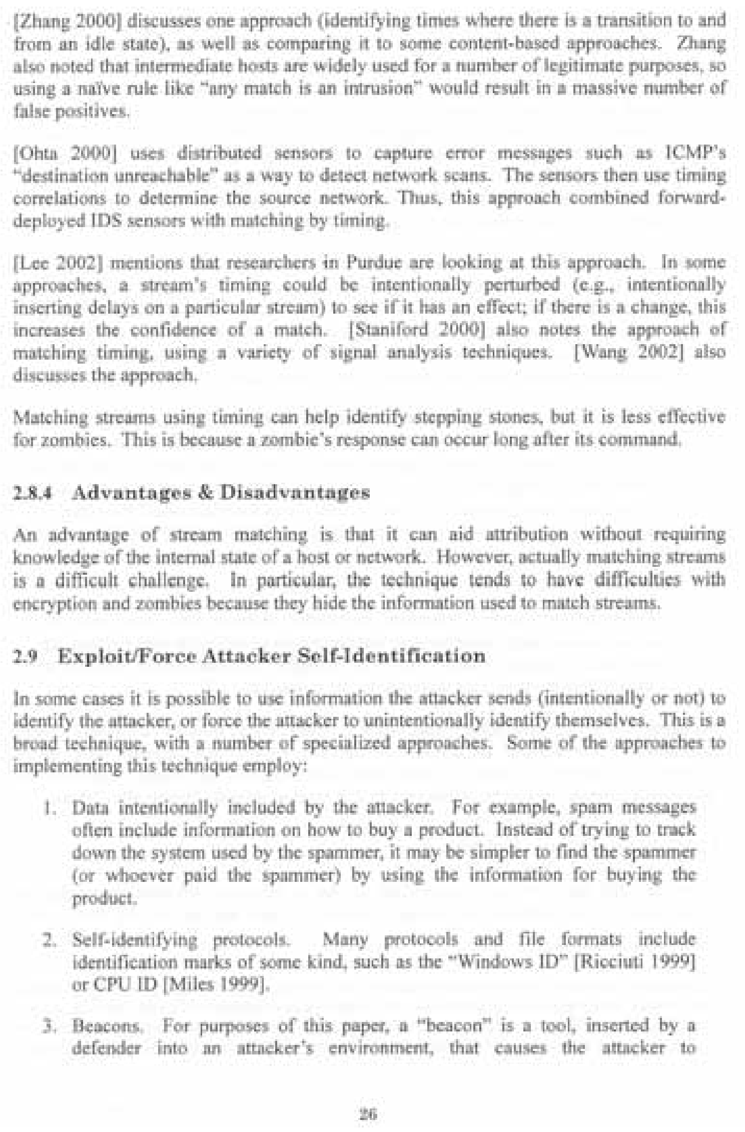

[Zhang 2000] discusses one approach (identifying times where there is a transition to and from an idle state), as well as comparing it to some content-based approaches. Zhang also noted that intermediate hosts are widely used for a number of legitimate purposes, so using a naïve rule like "any match is an intrusion" would result in a massive number of false positives.

[Ohta 2000] uses distributed sensors to capture error messages such as ICMP's "destination unreachable" as a way to detect network scans. The sensors then use timing correlations to determine the source network. Thus, this approach combined forward-deployed IDS sensors with matching by timing.

[Lee 2002] mentions that researchers in Purdue are looking at this approach. In some approaches, a stream's timing could be intentionally perturbed (e.g., intentionally inserting delays on a particular stream) to see if it has an effect; if there is a change, this increases the confidence of a match. [Staniford 2000] also notes the approach of matching timing, using a variety of signal analysis techniques. [Wang 2002] also discusses the approach.

Matching streams using timing can help identify stepping stones, but it is less effective for zombies. This is because a zombie's response can occur long after its command.

2.8.4 Advantages & Disadvantages

An advantage of stream matching is that it can aid attribution without requiring knowledge of the internal state of a host or network. However, actually matching streams is a difficult challenge. In particular, the technique tends to have difficulties with encryption and zombies because they hide the information used to match streams.

2.9 Exploit/Force Attacker Self-Identification

In some cases it is possible to use information the attacker sends (intentionally or not) to identify the attacker, or force the attacker to unintentionally identify themselves. This is a broad technique, with a number of specialized approaches. Some of the approaches to implementing this technique employ:

1. Data intentionally included by the attacker. For example, spam messages often include information on how to buy a product. Instead of trying to track down the system used by the spammer, it may be simpler to find the spammer (or whoever paid the spammer) by using the information for buying the product.

2. Self-identifying protocols. Many protocols and file formats include identification marks of some kind, such as the "Windows ID" [Ricciuti 1999] or CPU ID [Miles 1999].

3. Beacons. For purposes of this paper, a "beacon" is a tool, inserted by a defender into an attacker's environment, that causes the attacker to

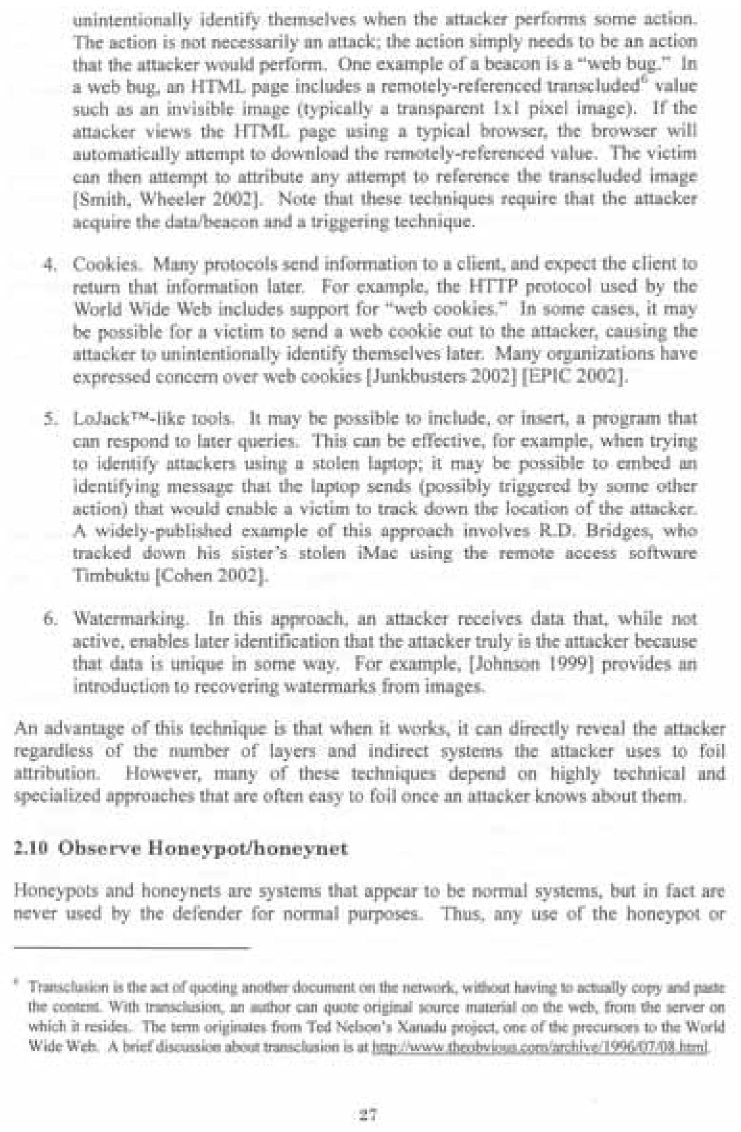

unintentionally identify themselves when the attacker performs some action. The action is not necessarily an attack; the action simply needs to be an action that the attacker would perform. One example of a beacon is a "web bug." In a web bug, an HTML page includes a remotely-referenced transcluded[6] value such as an invisible image (typically a transparent 1x1 pixel image). If the attacker views the HTML page using a typical browser, the browser will automatically attempt to download the remotely-referenced value. The victim can then attempt to attribute any attempt to reference the transcluded image [Smith, Wheeler 2002]. Note that these techniques require that the attacker acquire the data/beacon and a triggering technique.

4. Cookies. Many protocols send information to a client, and expect the client to return that information later. For example, the HTTP protocol used by the World Wide Web includes support for "web cookies." In some cases, it may be possible for a victim to send a web cookie out to the attacker, causing the attacker to unintentionally identify themselves later. Many organizations have expressed concern over web cookies [Junkbusters 2002] [EPIC 2002].

5. LoJack™-like tools. It may be possible to include, or insert, a program that can respond to later queries. This can be effective, for example, when trying to identify attackers using a stolen laptop; it may be possible to embed an identifying message that the laptop sends (possibly triggered by some other action) that would enable a victim to track down the location of the attacker. A widely-published example of this approach involves R.D. Bridges, who tracked down his sister's stolen iMac using the remote access software Timbuktu [Cohen 2002].

6. Watermarking. In this approach, an attacker receives data that, while not active, enables later identification that the attacker truly is the attacker because that data is unique in some way. For example, [Johnson 1999] provides an introduction to recovering watermarks from images.

An advantage of this technique is that when it works, it can directly reveal the attacker regardless of the number of layers and indirect systems the attacker uses to foil attribution. However, many of these techniques depend on highly technical and specialized approaches that are often easy to foil once an attacker knows about them.

2.10 Observe Honeypot/honeynet

Honeypots and honeynets are systems that appear to be normal systems, but in fact are never used by the defender for normal purposes. Thus, any use of the honeypot or

[6] Transclusion is the act of quoting another document on the network, without having to actually copy and paste the content. With transclusion, an author can quote original source material on the web, from the server on which it resides. The term originates from Ted Nelson's Xanadu project, one of the precursors to the World Wide Web. A brief discussion about transclusion is at http://www.theobvious.com/archive/1996/07/08.html.

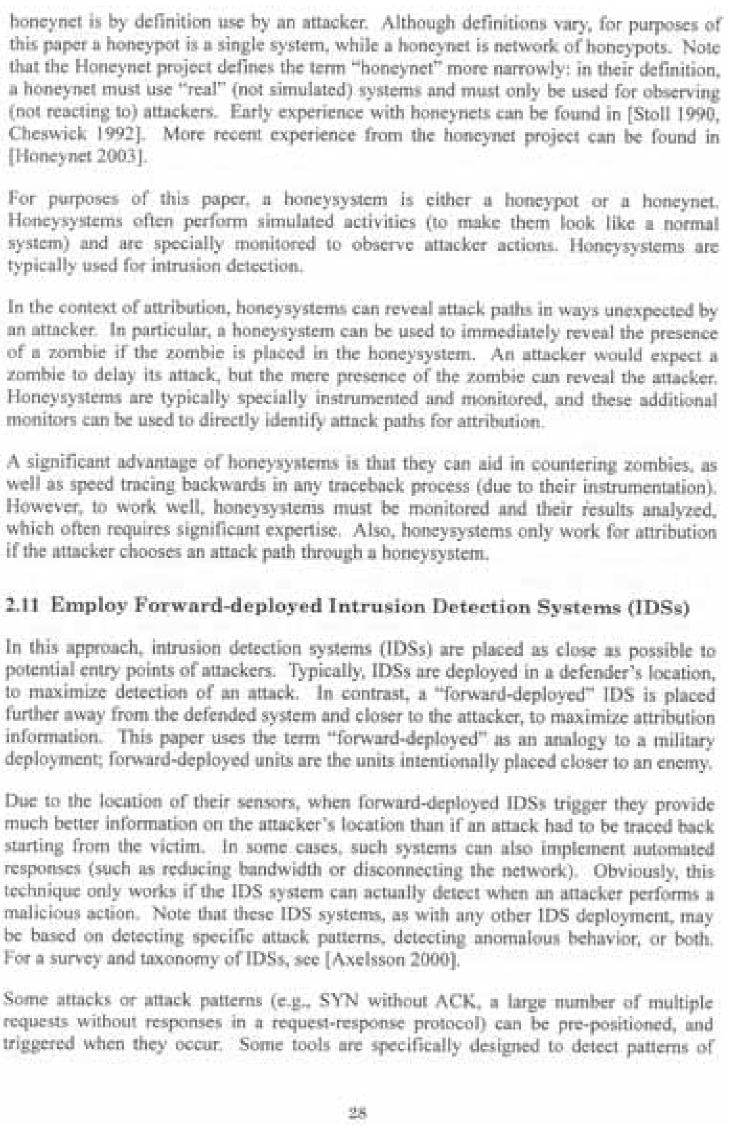

honeynet is by definition use by an attacker. Although definitions vary, for purposes of this paper a honeypot is a single system, while a honeynet is network of honeypots. Note that the Honeynet project defines the term "honeynet" more narrowly: in their definition, a honeynet must use "real" (not simulated) systems and must only be used for observing (not reacting to) attackers. Early experience with honeynets can be found in [Stoll 1990, Cheswick 1992]. More recent experience from the honeynet project can be found in [Honeynet 2003].

For purposes of this paper, a honeysystem is either a honeypot or a honeynet. Honeysystems often perform simulated activities (to make them look like a normal system) and are specially monitored to observe attacker actions. Honeysystems are typically used for intrusion detection.

In the context of attribution, honeysystems can reveal attack paths in ways unexpected by an attacker. In particular, a honeysystem can be used to immediately reveal the presence of a zombie if the zombie is placed in the honeysystem. An attacker would expect a zombie to delay its attack, but the mere presence of the zombie can reveal the attacker. Honeysystems are typically specially instrumented and monitored, and these additional monitors can be used to directly identify attack paths for attribution.

A significant advantage of honeysystems is that they can aid in countering zombies, as well as speed tracing backwards in any traceback process (due to their instrumentation). However, to work well, honeysystems must be monitored and their results analyzed, which often requires significant expertise. Also, honeysystems only work for attribution if the attacker chooses an attack path through a honeysystem.

2.11 Employ Forward-deployed Intrusion Detection Systems (IDSs)

In this approach, intrusion detection systems (IDSs) are placed as close as possible to potential entry points of attackers. Typically, IDSs are deployed in a defender's location, to maximize detection of an attack. In contrast, a "forward-deployed" IDS is placed further away from the defended system and closer to the attacker, to maximize attribution information. This paper uses the term "forward-deployed" as an analogy to a military deployment; forward-deployed units are the units intentionally placed closer to an enemy.

Due to the location of their sensors, when forward-deployed IDSs trigger they provide much better information on the attacker's location than if an attack had to be traced back starting from the victim. In some cases, such systems can also implement automated responses (such as reducing bandwidth or disconnecting the network). Obviously, this technique only works if the IDS system can actually detect when an attacker performs a malicious action. Note that these IDS systems, as with any other IDS deployment, may be based on detecting specific attack patterns, detecting anomalous behavior, or both. For a survey and taxonomy of IDSs, see [Axelsson 2000].

Some attacks or attack patterns (e.g., SYN without ACK, a large number of multiple requests without responses in a request-response protocol) can be pre-positioned, and triggered when they occur. Some tools are specifically designed to detect patterns of

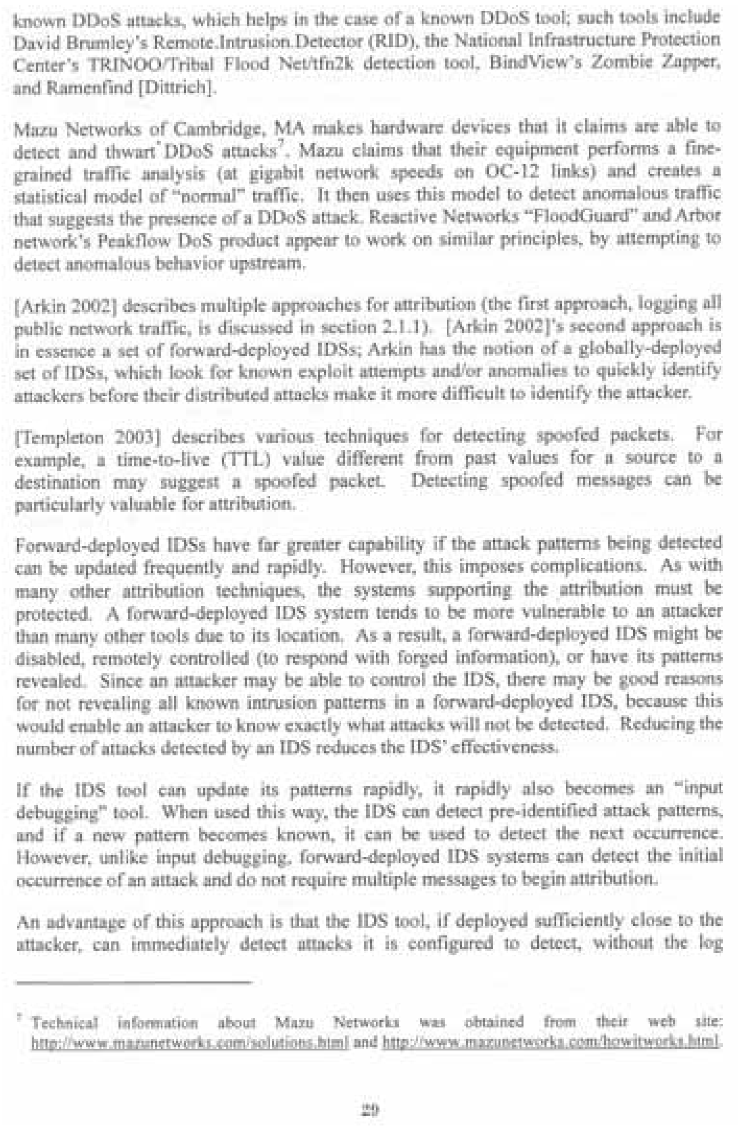

known DDoS attacks, which helps in the case of a known DDoS tool; such tools include David Brumley's Remote.Intrusion.Detector (RID), the National Infrastructure Protection Center's TRINOO/Tribal Flood Net/tfn2k detection tool, BindView's Zombie Zapper, and Ramenfind [Dittrich].

Mazu Networks of Cambridge, MA makes hardware devices that it claims are able to detect and thwart DDoS attacks[7]. Mazu claims that their equipment performs a fine-grained traffic analysis (at gigabit network speeds on OC-12 links) and creates a statistical model of "normal" traffic. It then uses this model to detect anomalous traffic that suggests the presence of a DDoS attack. Reactive Networks "FloodGuard" and Arbor network's Peakflow DoS product appear to work on similar principles, by attempting to detect anomalous behavior upstream.

[Arkin 2002] describes multiple approaches for attribution (the first approach, logging all public network traffic, is discussed in section 2.1.1). [Arkin 2002]'s second approach is in essence a set of forward-deployed IDSs; Arkin has the notion of a globally-deployed set of IDSs, which look for known exploit attempts and/or anomalies to quickly identify attackers before their distributed attacks make it more difficult to identify the attacker.

[Templeton 2003] describes various techniques for detecting spoofed packets. For example, a time-to-live (TTL) value different from past values for a source to a destination may suggest a spoofed packet. Detecting spoofed messages can be particularly valuable for attribution.

Forward-deployed IDSs have far greater capability if the attack patterns being detected can be updated frequently and rapidly. However, this imposes complications. As with many other attribution techniques, the systems supporting the attribution must be protected. A forward-deployed IDS system tends to be more vulnerable to an attacker than many other tools due to its location. As a result, a forward-deployed IDS might be disabled, remotely controlled (to respond with forged information), or have its patterns revealed. Since an attacker may be able to control the IDS, there may be good reasons for not revealing all known intrusion patterns in a forward-deployed IDS, because this would enable an attacker to know exactly what attacks will not be detected. Reducing the number of attacks detected by an IDS reduces the IDS' effectiveness.

If the IDS tool can update its patterns rapidly, it rapidly also becomes an "input debugging" tool. When used this way, the IDS can detect pre-identified attack patterns, and if a new pattern becomes known, it can be used to detect the next occurrence. However, unlike input debugging, forward-deployed IDS systems can detect the initial occurrence of an attack and do not require multiple messages to begin attribution.

An advantage of this approach is that the IDS tool, if deployed sufficiently close to the attacker, can immediately detect attacks it is configured to detect, without the log

[7] Technical information about Mazu Networks was obtained from their web site: http://www.mazunetworks.com/solutions.html and http://www.mazunetworks.com/howitworks.html.

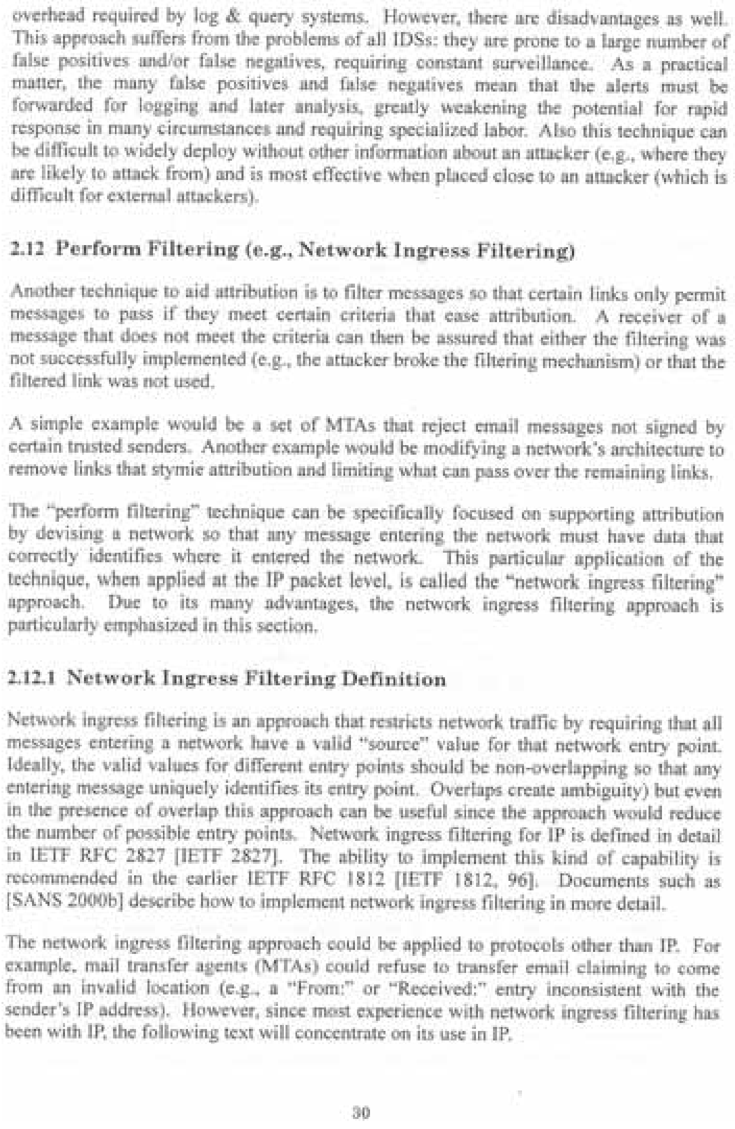

overhead required by log & query systems. However, there are disadvantages as well. This approach suffers from the problems of all IDSs: they are prone to a large number of false positives and/or false negatives, requiring constant surveillance. As a practical matter, the many false positives and false negatives mean that the alerts must be forwarded for logging and later analysis, greatly weakening the potential for rapid response in many circumstances and requiring specialized labor. Also this technique can be difficult to widely deploy without other information about an attacker (e.g., where they are likely to attack from) and is most effective when placed close to an attacker (which is difficult for external attackers).

2.12 Perform Filtering (e.g., Network Ingress Filtering)

Another technique to aid attribution is to filter messages so that certain links only permit messages to pass if they meet certain criteria that ease attribution. A receiver of a message that does not meet the criteria can then be assured that either the filtering was not successfully implemented (e.g., the attacker broke the filtering mechanism) or that the filtered link was not used.

A simple example would be a set of MTAs that reject email messages not signed by certain trusted senders. Another example would be modifying a network's architecture to remove links that stymie attribution and limiting what can pass over the remaining links.

The "perform filtering" technique can be specifically focused on supporting attribution by devising a network so that any message entering the network must have data that correctly identifies where it entered the network. This particular application of the technique, when applied at the IP packet level, is called the "network ingress filtering" approach. Due to its many advantages, the network ingress filtering approach is particularly emphasized in this section.

2.12.1 Network Ingress Filtering Definition

Network ingress filtering is an approach that restricts network traffic by requiring that all messages entering a network have a valid "source" value for that network entry point. Ideally, the valid values for different entry points should be non-overlapping so that any entering message uniquely identifies its entry point. Overlaps create ambiguity) but even in the presence of overlap this approach can be useful since the approach would reduce the number of possible entry points. Network ingress filtering for IP is defined in detail in IETF RFC 2827 [IETF 2827]. The ability to implement this kind of capability is recommended in the earlier IETF RFC 1812 [IETF 1812, 96]. Documents such as [SANS 2000b] describe how to implement network ingress filtering in more detail.

The network ingress filtering approach could be applied to protocols other than IP. For example, mail transfer agents (MTAs) could refuse to transfer email claiming to come from an invalid location (e.g., a "From:" or "Received:" entry inconsistent with the sender's IP address). However, since most experience with network ingress filtering has been with IP, the following text will concentrate on its use in IP.

2.12.2 Network Ingress Filtering Implementation

Implementing network ingress filtering of IP packets only requires simple packet filtering, a capability built into most of today's routers and a fundamental capability included in any firewall. Thus, implementation primarily involves reconfiguring the existing routers (or in rare cases, inserting firewalls) that connect other networks to the network being protected. In short, all connections to other networks must ensure that the source information is valid for the given connection.

Figure 4 shows a sample network ingress filtering configuration. In this example, the filtered network has connections to an external network through multiple gateways, as well as connections to many internal networks through major routers. All of those routers (E1, E2, E3, GW1, and GW2) are configured so that a packet's source address must be in a valid range to pass through those routers. Thus, in this figure, the attacker connected through router E11 who is attacking victim 1 must reveal that they are located in network E1, given certain caveats discussed below.

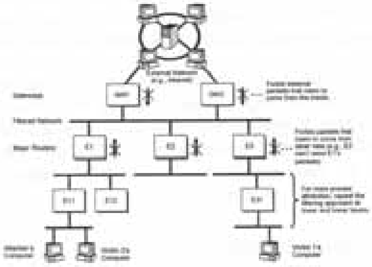

Figure 4. Network Ingress Filtering

For a typical wide area network (WAN) that connects to both the "outside" and to "internal" networks, two kinds of configurations are necessary to change an unfiltered WAN into a filtered WAN:

1. Routers connecting to the outside must forbid messages that try to enter the filtered network from the outside, yet claim to come from the inside. This is a generally accepted practice, as doing otherwise can permit many security problems, and is the fundamental rule implemented by even trivial firewalls.

2. Routers connecting to "internal" networks must forbid messages that try to enter the filtered network from the internal network, yet claim to come from somewhere other than that specific internal network. For example, if an internal network is allocated the IP address range 204.69.207.x (more

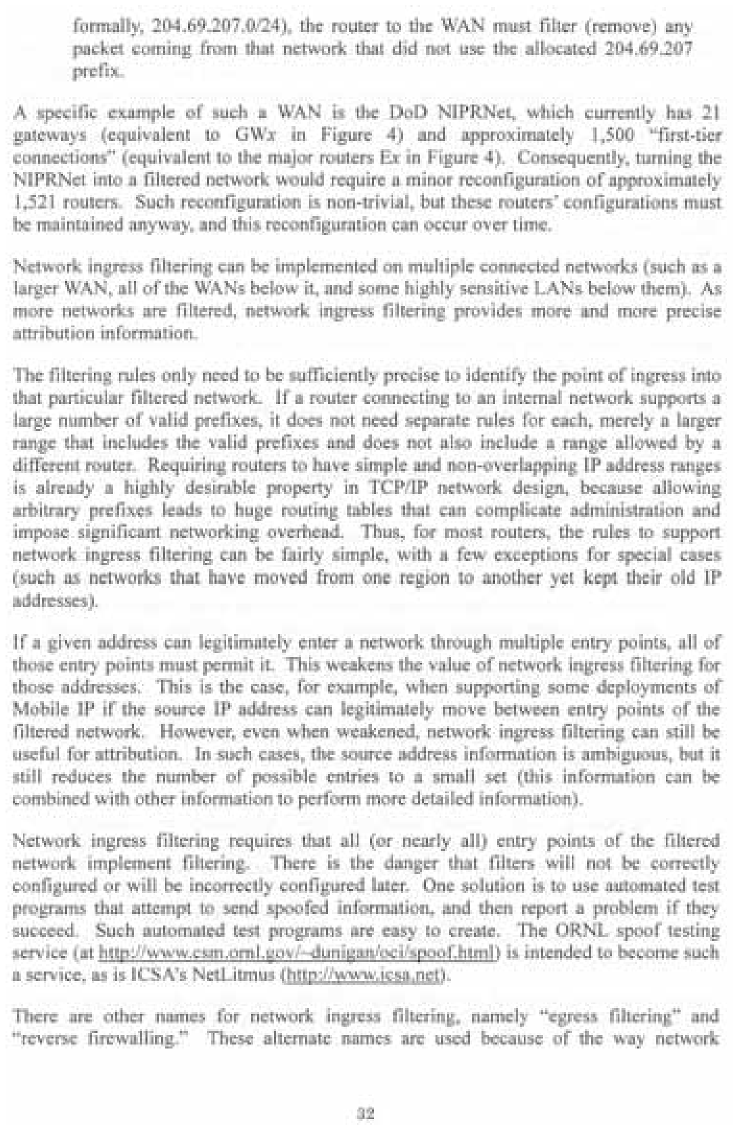

formally, 204.69.207.0/24), the router to the WAN must filter (remove) any packet coming from that network that did not use the allocated 204.69.207 prefix.

A specific example of such a WAN is the DoD NIPRNet, which currently has 21 gateways (equivalent to GW*x* in Figure 4) and approximately 1,500 "first-tier connections" (equivalent to the major routers E*x* in Figure 4). Consequently, turning the NIPRNet into a filtered network would require a minor reconfiguration of approximately 1,521 routers. Such reconfiguration is non-trivial, but these routers' configurations must be maintained anyway, and this reconfiguration can occur over time.

Network ingress filtering can be implemented on multiple connected networks (such as a larger WAN, all of the WANs below it, and some highly sensitive LANs below them). As more networks are filtered, network ingress filtering provides more and more precise attribution information.

The filtering rules only need to be sufficiently precise to identify the point of ingress into that particular filtered network. If a router connecting to an internal network supports a large number of valid prefixes, it does not need separate rules for each, merely a larger range that includes the valid prefixes and does not also include a range allowed by a different router. Requiring routers to have simple and non-overlapping IP address ranges is already a highly desirable property in TCP/IP network design, because allowing arbitrary prefixes leads to huge routing tables that can complicate administration and impose significant networking overhead. Thus, for most routers, the rules to support network ingress filtering can be fairly simple, with a few exceptions for special cases (such as networks that have moved from one region to another yet kept their old IP addresses).

If a given address can legitimately enter a network through multiple entry points, all of those entry points must permit it. This weakens the value of network ingress filtering for those addresses. This is the case, for example, when supporting some deployments of Mobile IP if the source IP address can legitimately move between entry points of the filtered network. However, even when weakened, network ingress filtering can still be useful for attribution. In such cases, the source address information is ambiguous, but it still reduces the number of possible entries to a small set (this information can be combined with other information to perform more detailed information).

Network ingress filtering requires that all (or nearly all) entry points of the filtered network implement filtering. There is the danger that filters will not be correctly configured or will be incorrectly configured later. One solution is to use automated test programs that attempt to send spoofed information, and then report a problem if they succeed. Such automated test programs are easy to create. The ORNL spoof testing service (at http://www.csm.ornl.gov/~dunigan/oci/spoof.html) is intended to become such a service, as is ICSA's NetLitmus (http://www.icsa.net).

There are other names for network ingress filtering, namely "egress filtering" and "reverse firewalling." These alternate names are used because of the way network

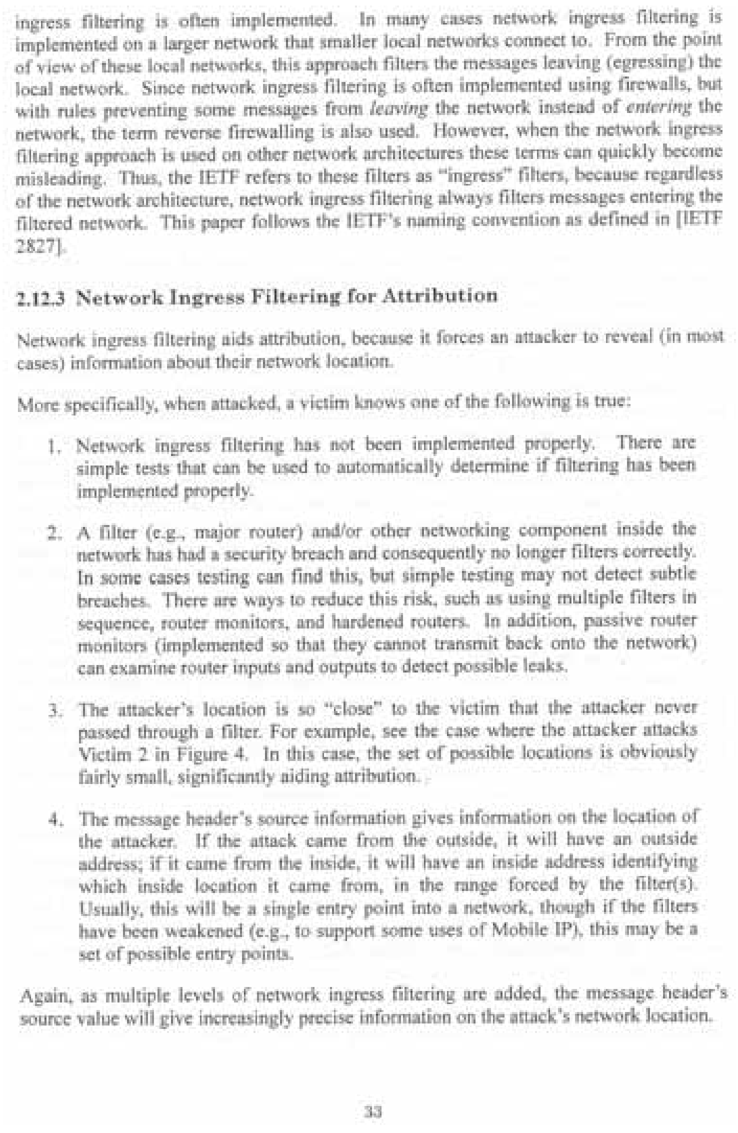

ingress filtering is often implemented. In many cases network ingress filtering is implemented on a larger network that smaller local networks connect to. From the point of view of these local networks, this approach filters the messages leaving (egressing) the local network. Since network ingress filtering is often implemented using firewalls, but with rules preventing some messages from *leaving* the network instead of *entering* the network, the term reverse firewalling is also used. However, when the network ingress filtering approach is used on other network architectures these terms can quickly become misleading. Thus, the IETF refers to these filters as "ingress" filters, because regardless of the network architecture, network ingress filtering always filters messages entering the filtered network. This paper follows the IETF's naming convention as defined in [IETF 2827].

2.12.3 Network Ingress Filtering for Attribution

Network ingress filtering aids attribution, because it forces an attacker to reveal (in most cases) information about their network location.

More specifically, when attacked, a victim knows one of the following is true:

1. Network ingress filtering has not been implemented properly. There are simple tests that can be used to automatically determine if filtering has been implemented properly.

2. A filter (e.g., major router) and/or other networking component inside the network has had a security breach and consequently no longer filters correctly. In some cases testing can find this, but simple testing may not detect subtle breaches. There are ways to reduce this risk, such as using multiple filters in sequence, router monitors, and hardened routers. In addition, passive router monitors (implemented so that they cannot transmit back onto the network) can examine router inputs and outputs to detect possible leaks.

3. The attacker's location is so "close" to the victim that the attacker never passed through a filter. For example, see the case where the attacker attacks Victim 2 in Figure 4. In this case, the set of possible locations is obviously fairly small, significantly aiding attribution.

4. The message header's source information gives information on the location of the attacker. If the attack came from the outside, it will have an outside address; if it came from the inside, it will have an inside address identifying which inside location it came from, in the range forced by the filter(s). Usually, this will be a single entry point into a network, though if the filters have been weakened (e.g., to support some uses of Mobile IP), this may be a set of possible entry points.

Again, as multiple levels of network ingress filtering are added, the message header's source value will give increasingly precise information on the attack's network location.

2.12.4 Network Ingress Filtering Advantages

There are a number of advantages to network ingress filtering for attribution:

1. It is relatively easily implemented today at low cost with existing infrastructure. Implementation requires policies, configuration of existing filters (routers and firewalls) by network administrators, upgrades of components in rare cases, and simple occasional tests to ensure that the filtering is in operation. There is little to buy, existing personnel and products can be used to implement it, and since it is a relatively simple configuration, training costs are not expected to be high. This does not mean that it is trivial; any policy change requires coordination, and reconfiguring all routers on any major network is not trivial. However, it is less expensive than many other approaches, since many other approaches require deploying a large number of new hardware and/or software components as well as training personnel to learn new skills.

2. It can be deployed incrementally, with incremental improvements in attribution. This can be viewed in two ways:

   a. The routers on a given filtered network can have their rules modified in stages, instead of trying to modify all routers at once. Although the benefits of ingress filtering are primarily obvious when all (or nearly all) routers of a network do the filtering, even implementing the approach on a subset of routers can aid attribution by reducing the number of possible paths that must be traced back. Adding the rules incrementally can also aid in identifying the cause of any unintentional problem.

   b. Network ingress filtering can be beneficially implemented on a single major backbone and provide a benefit, and later deployed on additional networks to provide increasingly more precise attribution. This property – of not requiring all networks to be changed simultaneously – also simplifies deployment.

3. Implementation of network ingress filtering can be made a requirement for connection to certain networks. Thus, once a network ingress filtering regime is set up, it should be easy to maintain.

4. Network ingress filtering supports attribution without requiring message logs of unrelated messages or additional network bandwidth. Logging systems have complications due to the difficulties of acquiring data at speed, storing logs, and retrieving the correct data later. Techniques that send new messages or extend messages take additional network bandwidth. The filtering technique (including specifically the network ingress filtering approach) requires neither.

5. Network ingress filtering is generally transparent to users. Since the filtering rules simply enforce what users "should" do, users are generally unaware of the filtering; tools simply work as usual.

6. There are no known legal impediments. Since the filtering rules simply prevent forging of source addresses, there are no known laws that forbid deployment. In certain cases some state laws even forbid forging "from" information (e.g., some state anti-spam laws).

2.12.5 Network Ingress Filtering Disadvantages

Of course, network ingress filtering is not perfect. There are a number of disadvantages to network ingress filtering when used for attribution:

1. Network ingress filtering must be implemented by *nearly every* entry point into a given filtered network to be effective. If some entry points do not implement the rules, then any attack message *may* have also been sent through those entry points. If there are more than a few entry points that do not implement the filtering rules, the value of network ingress filtering for attribution goes down rapidly (unless supported by other techniques).

2. Networks that must support multiple entry paths, such as some uses of Mobile IP and permanently enabled backup routes, supply weaker attribution information for those messages. In the worse case, it can somewhat interfere with fault tolerance, since the rules could forbid alternative routes that might be desirable. This is actually a variation of the first point; if the filtering rules must allow multiple entry paths, then a given message may have taken any of those entry paths, making the message harder to attribute.

3. Network ingress filtering is primarily useful for internal network attribution and to determine if an attack came from the "outside." Since generally only "internal" networks can be required to implement network ingress filtering rules, the approach is only useful for those networks. Thus, in the shorter term, this approach is probably more useful for organizations that are concerned about threats from within networks they control.

   In the longer term, this approach could be applied to countries or even internationally to give more attribution information. For example, U.S. law could be modified to require network ingress filtering on U.S. Internet Service Providers (ISPs). In that case, network ingress filtering could aid attribution of a (probably intermediate) system inside the U.S. and identify messages that originated from outside the U.S. as well. Note, however, that it would not reliably identify the network for messages originating from outside the U.S., because attackers could send messages through interconnected non-U.S. networks. Thus, attackers would quickly move to use at least some intermediaries outside U.S. jurisdiction.

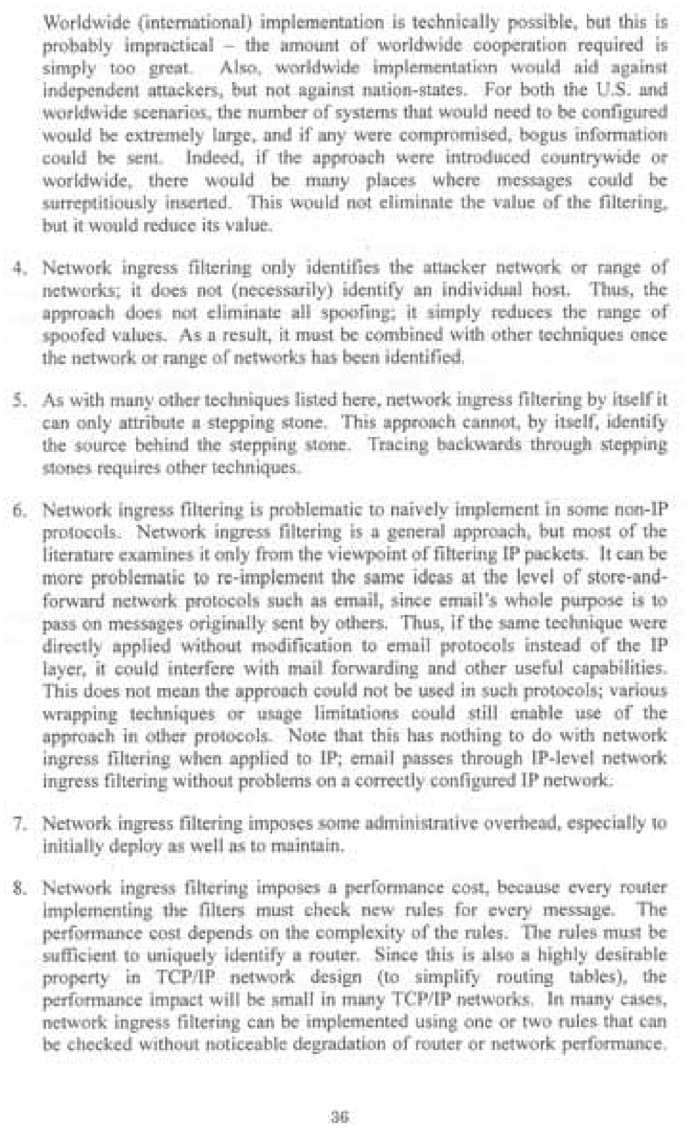

Worldwide (international) implementation is technically possible, but this is probably impractical – the amount of worldwide cooperation required is simply too great. Also, worldwide implementation would aid against independent attackers, but not against nation-states. For both the U.S. and worldwide scenarios, the number of systems that would need to be configured would be extremely large, and if any were compromised, bogus information could be sent. Indeed, if the approach were introduced countrywide or worldwide, there would be many places where messages could be surreptitiously inserted. This would not eliminate the value of the filtering, but it would reduce its value.

4. Network ingress filtering only identifies the attacker network or range of networks; it does not (necessarily) identify an individual host. Thus, the approach does not eliminate all spoofing; it simply reduces the range of spoofed values. As a result, it must be combined with other techniques once the network or range of networks has been identified.

5. As with many other techniques listed here, network ingress filtering by itself it can only attribute a stepping stone. This approach cannot, by itself, identify the source behind the stepping stone. Tracing backwards through stepping stones requires other techniques.

6. Network ingress filtering is problematic to naively implement in some non-IP protocols. Network ingress filtering is a general approach, but most of the literature examines it only from the viewpoint of filtering IP packets. It can be more problematic to re-implement the same ideas at the level of store-and-forward network protocols such as email, since email's whole purpose is to pass on messages originally sent by others. Thus, if the same technique were directly applied without modification to email protocols instead of the IP layer, it could interfere with mail forwarding and other useful capabilities. This does not mean the approach could not be used in such protocols; various wrapping techniques or usage limitations could still enable use of the approach in other protocols. Note that this has nothing to do with network ingress filtering when applied to IP; email passes through IP-level network ingress filtering without problems on a correctly configured IP network.

7. Network ingress filtering imposes some administrative overhead, especially to initially deploy as well as to maintain.

8. Network ingress filtering imposes a performance cost, because every router implementing the filters must check new rules for every message. The performance cost depends on the complexity of the rules. The rules must be sufficient to uniquely identify a router. Since this is also a highly desirable property in TCP/IP network design (to simplify routing tables), the performance impact will be small in many TCP/IP networks. In many cases, network ingress filtering can be implemented using one or two rules that can be checked without noticeable degradation of router or network performance.

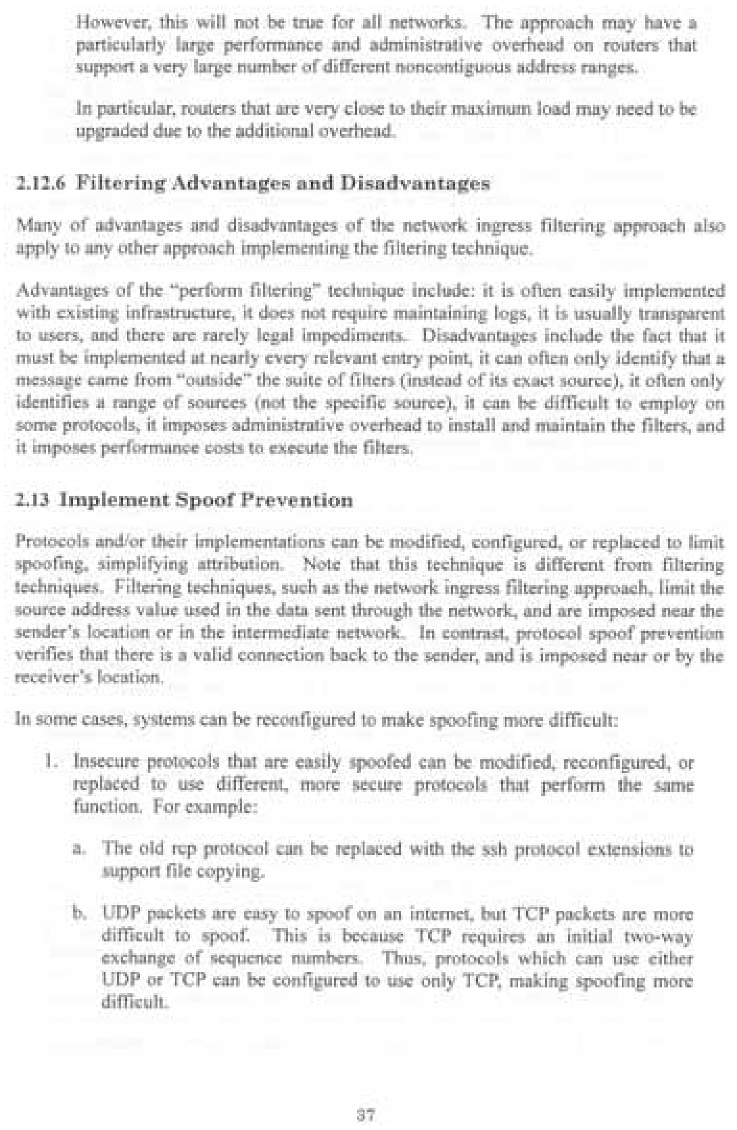

However, this will not be true for all networks. The approach may have a particularly large performance and administrative overhead on routers that support a very large number of different noncontiguous address ranges.

In particular, routers that are very close to their maximum load may need to be upgraded due to the additional overhead.

2.12.6 Filtering Advantages and Disadvantages

Many of advantages and disadvantages of the network ingress filtering approach also apply to any other approach implementing the filtering technique.

Advantages of the "perform filtering" technique include: it is often easily implemented with existing infrastructure, it does not require maintaining logs, it is usually transparent to users, and there are rarely legal impediments. Disadvantages include the fact that it must be implemented at nearly every relevant entry point, it can often only identify that a message came from "outside" the suite of filters (instead of its exact source), it often only identifies a range of sources (not the specific source), it can be difficult to employ on some protocols, it imposes administrative overhead to install and maintain the filters, and it imposes performance costs to execute the filters.

2.13 Implement Spoof Prevention

Protocols and/or their implementations can be modified, configured, or replaced to limit spoofing, simplifying attribution. Note that this technique is different from filtering techniques. Filtering techniques, such as the network ingress filtering approach, limit the source address value used in the data sent through the network, and are imposed near the sender's location or in the intermediate network. In contrast, protocol spoof prevention verifies that there is a valid connection back to the sender, and is imposed near or by the receiver's location.

In some cases, systems can be reconfigured to make spoofing more difficult:

1. Insecure protocols that are easily spoofed can be modified, reconfigured, or replaced to use different, more secure protocols that perform the same function. For example:

   a. The old rcp protocol can be replaced with the ssh protocol extensions to support file copying.

   b. UDP packets are easy to spoof on an internet, but TCP packets are more difficult to spoof. This is because TCP requires an initial two-way exchange of sequence numbers. Thus, protocols which can use either UDP or TCP can be configured to use only TCP, making spoofing more difficult.

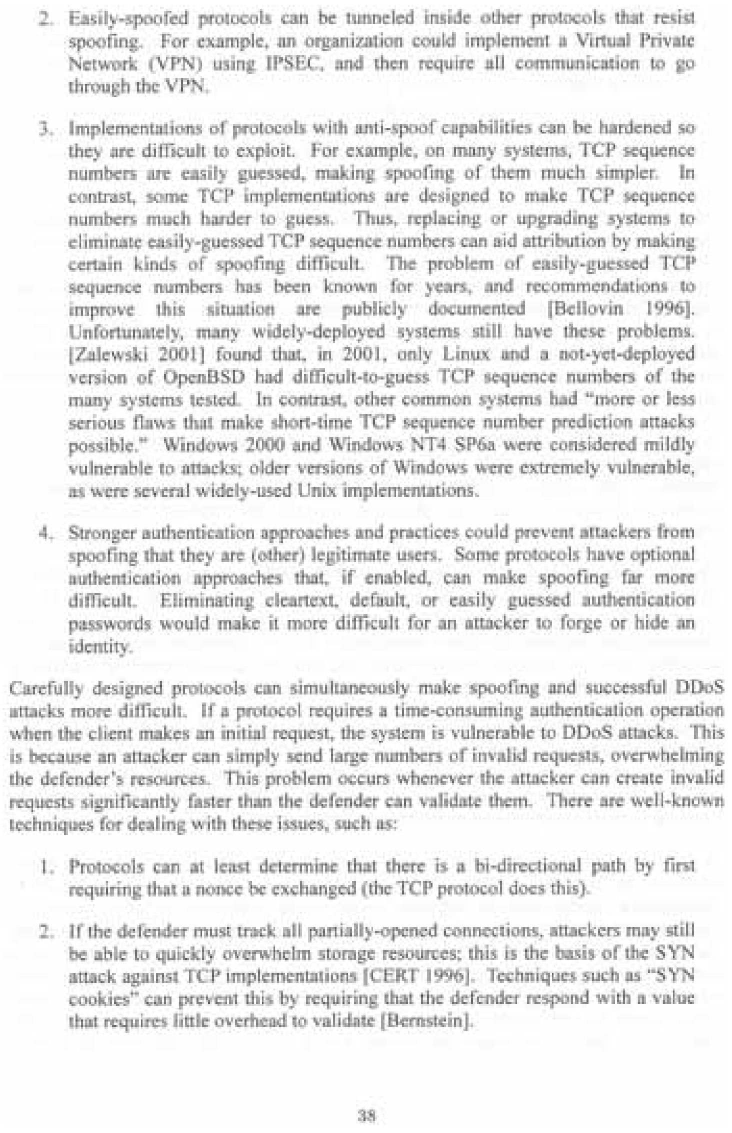

2. Easily-spoofed protocols can be tunneled inside other protocols that resist spoofing. For example, an organization could implement a Virtual Private Network (VPN) using IPSEC, and then require all communication to go through the VPN.

3. Implementations of protocols with anti-spoof capabilities can be hardened so they are difficult to exploit. For example, on many systems, TCP sequence numbers are easily guessed, making spoofing of them much simpler. In contrast, some TCP implementations are designed to make TCP sequence numbers much harder to guess. Thus, replacing or upgrading systems to eliminate easily-guessed TCP sequence numbers can aid attribution by making certain kinds of spoofing difficult. The problem of easily-guessed TCP sequence numbers has been known for years, and recommendations to improve this situation are publicly documented [Bellovin 1996]. Unfortunately, many widely-deployed systems still have these problems. [Zalewski 2001] found that, in 2001, only Linux and a not-yet-deployed version of OpenBSD had difficult-to-guess TCP sequence numbers of the many systems tested. In contrast, other common systems had "more or less serious flaws that make short-time TCP sequence number prediction attacks possible." Windows 2000 and Windows NT4 SP6a were considered mildly vulnerable to attacks; older versions of Windows were extremely vulnerable, as were several widely-used Unix implementations.

4. Stronger authentication approaches and practices could prevent attackers from spoofing that they are (other) legitimate users. Some protocols have optional authentication approaches that, if enabled, can make spoofing far more difficult. Eliminating cleartext, default, or easily guessed authentication passwords would make it more difficult for an attacker to forge or hide an identity.

Carefully designed protocols can simultaneously make spoofing and successful DDoS attacks more difficult. If a protocol requires a time-consuming authentication operation when the client makes an initial request, the system is vulnerable to DDoS attacks. This is because an attacker can simply send large numbers of invalid requests, overwhelming the defender's resources. This problem occurs whenever the attacker can create invalid requests significantly faster than the defender can validate them. There are well-known techniques for dealing with these issues, such as:

1. Protocols can at least determine that there is a bi-directional path by first requiring that a nonce be exchanged (the TCP protocol does this).

2. If the defender must track all partially-opened connections, attackers may still be able to quickly overwhelm storage resources; this is the basis of the SYN attack against TCP implementations [CERT 1996]. Techniques such as "SYN cookies" can prevent this by requiring that the defender respond with a value that requires little overhead to validate [Bernstein].

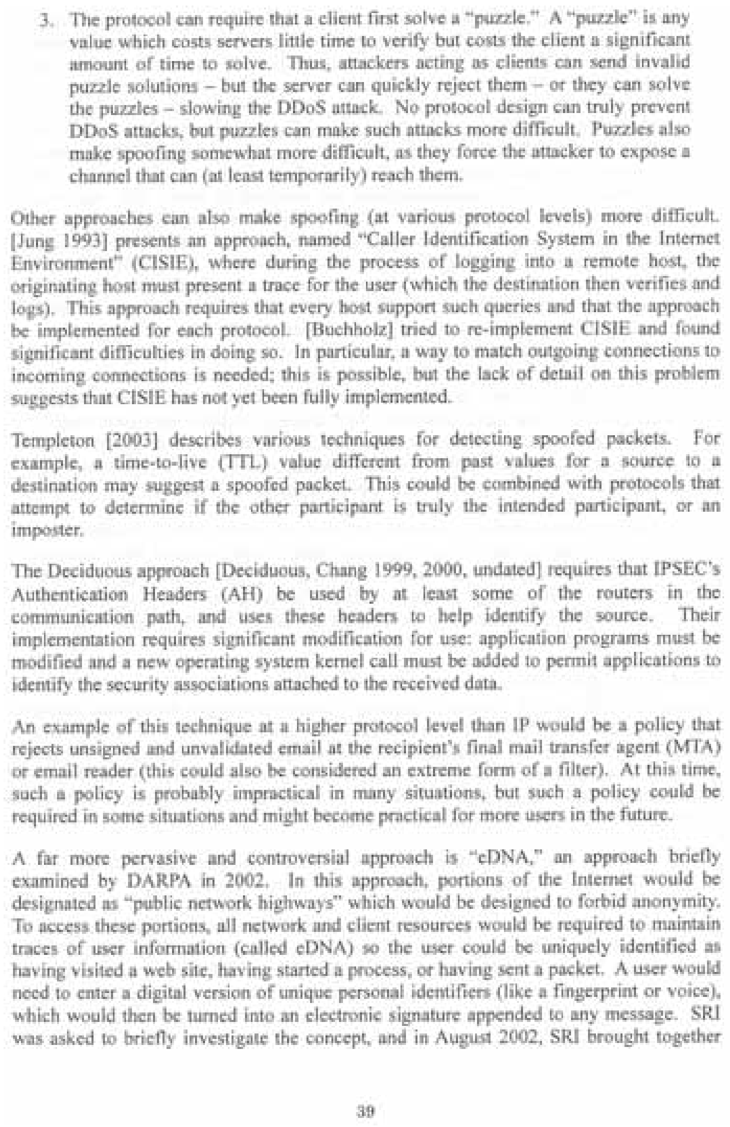

3. The protocol can require that a client first solve a "puzzle." A "puzzle" is any value which costs servers little time to verify but costs the client a significant amount of time to solve. Thus, attackers acting as clients can send invalid puzzle solutions – but the server can quickly reject them – or they can solve the puzzles – slowing the DDoS attack. No protocol design can truly prevent DDoS attacks, but puzzles can make such attacks more difficult. Puzzles also make spoofing somewhat more difficult, as they force the attacker to expose a channel that can (at least temporarily) reach them.

Other approaches can also make spoofing (at various protocol levels) more difficult. [Jung 1993] presents an approach, named "Caller Identification System in the Internet Environment" (CISIE), where during the process of logging into a remote host, the originating host must present a trace for the user (which the destination then verifies and logs). This approach requires that every host support such queries and that the approach be implemented for each protocol. [Buchholz] tried to re-implement CISIE and found significant difficulties in doing so. In particular, a way to match outgoing connections to incoming connections is needed; this is possible, but the lack of detail on this problem suggests that CISIE has not yet been fully implemented.

Templeton [2003] describes various techniques for detecting spoofed packets. For example, a time-to-live (TTL) value different from past values for a source to a destination may suggest a spoofed packet. This could be combined with protocols that attempt to determine if the other participant is truly the intended participant, or an imposter.

The Deciduous approach [Deciduous, Chang 1999, 2000, undated] requires that IPSEC's Authentication Headers (AH) be used by at least some of the routers in the communication path, and uses these headers to help identify the source. Their implementation requires significant modification for use: application programs must be modified and a new operating system kernel call must be added to permit applications to identify the security associations attached to the received data.

An example of this technique at a higher protocol level than IP would be a policy that rejects unsigned and unvalidated email at the recipient's final mail transfer agent (MTA) or email reader (this could also be considered an extreme form of a filter). At this time, such a policy is probably impractical in many situations, but such a policy could be required in some situations and might become practical for more users in the future.

A far more pervasive and controversial approach is "eDNA," an approach briefly examined by DARPA in 2002. In this approach, portions of the Internet would be designated as "public network highways" which would be designed to forbid anonymity. To access these portions, all network and client resources would be required to maintain traces of user information (called eDNA) so the user could be uniquely identified as having visited a web site, having started a process, or having sent a packet. A user would need to enter a digital version of unique personal identifiers (like a fingerprint or voice), which would then be turned into an electronic signature appended to any message. SRI was asked to briefly investigate the concept, and in August 2002, SRI brought together

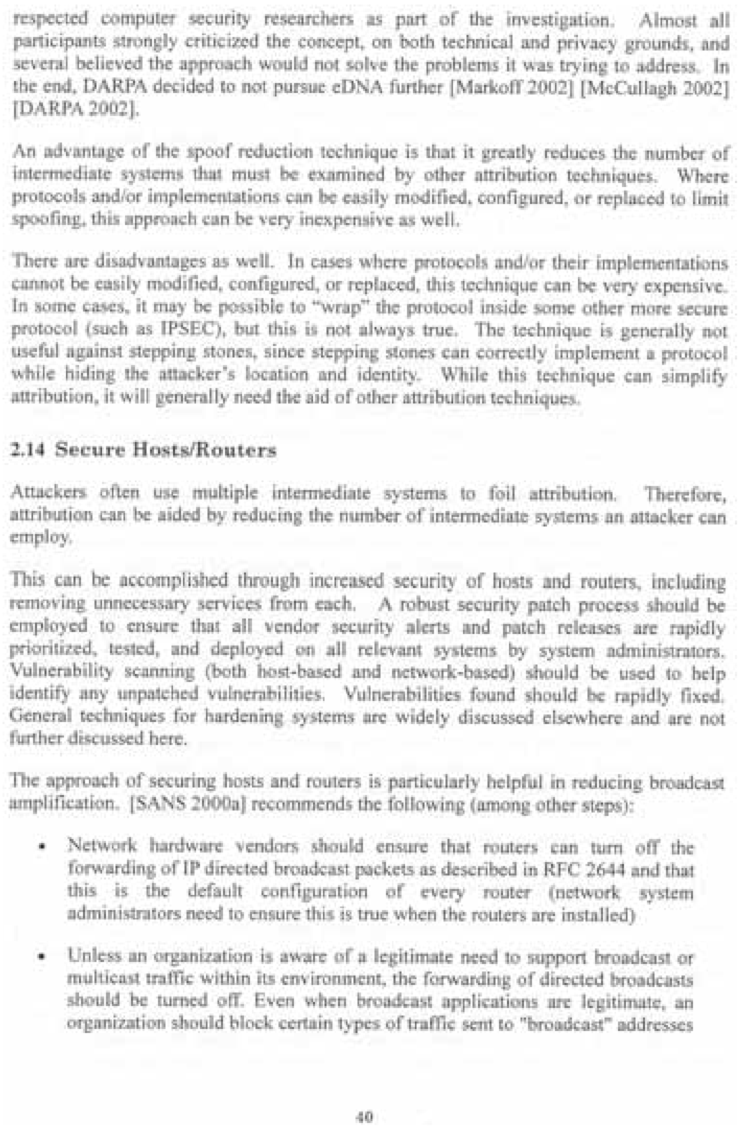

respected computer security researchers as part of the investigation. Almost all participants strongly criticized the concept, on both technical and privacy grounds, and several believed the approach would not solve the problems it was trying to address. In the end, DARPA decided to not pursue eDNA further [Markoff 2002] [McCullagh 2002] [DARPA 2002].

An advantage of the spoof reduction technique is that it greatly reduces the number of intermediate systems that must be examined by other attribution techniques. Where protocols and/or implementations can be easily modified, configured, or replaced to limit spoofing, this approach can be very inexpensive as well.

There are disadvantages as well. In cases where protocols and/or their implementations cannot be easily modified, configured, or replaced, this technique can be very expensive. In some cases, it may be possible to "wrap" the protocol inside some other more secure protocol (such as IPSEC), but this is not always true. The technique is generally not useful against stepping stones, since stepping stones can correctly implement a protocol while hiding the attacker's location and identity. While this technique can simplify attribution, it will generally need the aid of other attribution techniques.

2.14 Secure Hosts/Routers

Attackers often use multiple intermediate systems to foil attribution. Therefore, attribution can be aided by reducing the number of intermediate systems an attacker can employ.

This can be accomplished through increased security of hosts and routers, including removing unnecessary services from each. A robust security patch process should be employed to ensure that all vendor security alerts and patch releases are rapidly prioritized, tested, and deployed on all relevant systems by system administrators. Vulnerability scanning (both host-based and network-based) should be used to help identify any unpatched vulnerabilities. Vulnerabilities found should be rapidly fixed. General techniques for hardening systems are widely discussed elsewhere and are not further discussed here.

The approach of securing hosts and routers is particularly helpful in reducing broadcast amplification. [SANS 2000a] recommends the following (among other steps):

- Network hardware vendors should ensure that routers can turn off the forwarding of IP directed broadcast packets as described in RFC 2644 and that this is the default configuration of every router (network system administrators need to ensure this is true when the routers are installed)

- Unless an organization is aware of a legitimate need to support broadcast or multicast traffic within its environment, the forwarding of directed broadcasts should be turned off. Even when broadcast applications are legitimate, an organization should block certain types of traffic sent to "broadcast" addresses

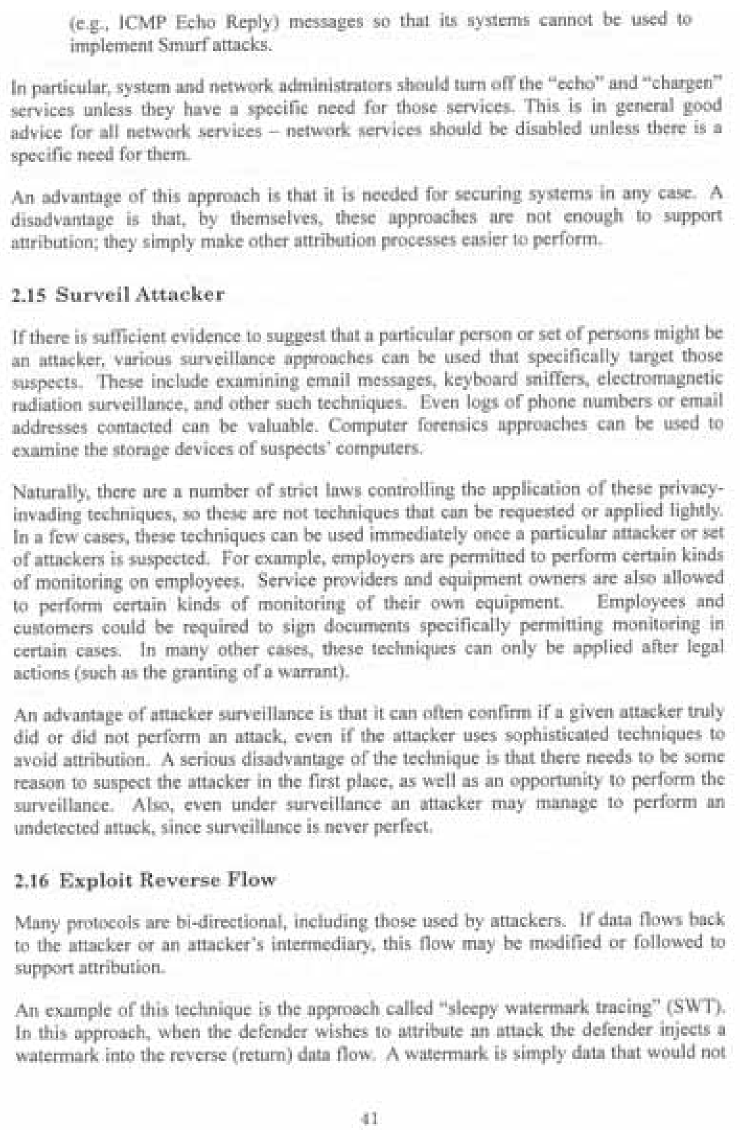

(e.g., ICMP Echo Reply) messages so that its systems cannot be used to implement Smurf attacks.

In particular, system and network administrators should turn off the "echo" and "chargen" services unless they have a specific need for those services. This is in general good advice for all network services – network services should be disabled unless there is a specific need for them.

An advantage of this approach is that it is needed for securing systems in any case. A disadvantage is that, by themselves, these approaches are not enough to support attribution; they simply make other attribution processes easier to perform.

2.15 Surveil Attacker

If there is sufficient evidence to suggest that a particular person or set of persons might be an attacker, various surveillance approaches can be used that specifically target those suspects. These include examining email messages, keyboard sniffers, electromagnetic radiation surveillance, and other such techniques. Even logs of phone numbers or email addresses contacted can be valuable. Computer forensics approaches can be used to examine the storage devices of suspects' computers.

Naturally, there are a number of strict laws controlling the application of these privacy-invading techniques, so these are not techniques that can be requested or applied lightly. In a few cases, these techniques can be used immediately once a particular attacker or set of attackers is suspected. For example, employers are permitted to perform certain kinds of monitoring on employees. Service providers and equipment owners are also allowed to perform certain kinds of monitoring of their own equipment. Employees and customers could be required to sign documents specifically permitting monitoring in certain cases. In many other cases, these techniques can only be applied after legal actions (such as the granting of a warrant).

An advantage of attacker surveillance is that it can often confirm if a given attacker truly did or did not perform an attack, even if the attacker uses sophisticated techniques to avoid attribution. A serious disadvantage of the technique is that there needs to be some reason to suspect the attacker in the first place, as well as an opportunity to perform the surveillance. Also, even under surveillance an attacker may manage to perform an undetected attack, since surveillance is never perfect.

2.16 Exploit Reverse Flow

Many protocols are bi-directional, including those used by attackers. If data flows back to the attacker or an attacker's intermediary, this flow may be modified or followed to support attribution.

An example of this technique is the approach called "sleepy watermark tracing" (SWT). In this approach, when the defender wishes to attribute an attack the defender injects a watermark into the reverse (return) data flow. A watermark is simply data that would not normally be detected by an attacker. For example, in the telnet and rlogin protocols, a defender returning the string:

> "**See me**abc\b\b\b \b"

would look the same to an attacker as a defender who returned the string:

> "**See me**"

In SWT, network server applications (such as telnetd and rlogind) on the defender's host system are modified to be "watermark-enabled," so that on command they can insert these watermarks. SWT guardian gateways are then used to detect and report the presence of these watermarks. SWT is described in [Wang 2001a]. SWT response options are further described in [Wang 2001b].

One advantage of the technique is that it can attribute immediately through a large number of stepping stones, if the data is not transformed through processes such as encryption. However, there are many disadvantages. Most such implementations (such as SWT) require significant changes to pre-existing implementations. Detectors of the data in the reverse flow must be placed in locations that can actually observe the data, and there is always the danger of false positives from the detectors. Also, hosts that transform the data (such as encrypting the data) may foil the technique.

2.17 Combine Techniques

Since every technique has its strengths and weaknesses, it is probably more effective to combine the various techniques to perform attribution. For example, network sensors could be forward deployed closer to where an attacker might attack from. If the network sensors include initial rules for known attacks, they would be considered forward-deployed IDSs in this grouping. If they also supported rapid run-time requests for new patterns, they could also support input debugging. Network ingress filtering might be useful to reduce the number of possible networks an attacker came from, and then other attribution techniques could be used to identify the attacker's location more precisely.

Also note that different protocol layers can provide different information that together provide better attribution information. Many implementers concentrate on the IP layer; since IP is a common layer, attribution approaches based on the IP layer are more general. However, combining information from various protocol layers (such as IP, authentication logs, MTA logs, higher-level protocols, and lower level protocols) can add information that examining only one layer will miss.

In theory, an advantage of combining techniques is that it can overcome many of the disadvantages of individual techniques. However, currently there is relatively little experience (or available automation) for combining techniques. Also, combining techniques cannot overcome the old phrase "garbage in, garbage out" – if the results of the individual techniques are worthless, combining them will not help.

3. Issues in Attribution Techniques

This section discusses some of the issues common to many attribution techniques.

3.1 Prepositioning of Tools and Trust is Critical

Many attribution techniques cannot be applied to an attack unless the attribution implementations and trust relationships have been prepositioned. This is particularly obvious with logging systems; it is impossible to query a log unless the logging system has already been deployed. However, this is true for many other techniques, such as network ingress filtering.

Even if the technology does not need to be prepositioned, trust relationships need to be prepositioned for attacks to be attributed in a timely manner. For example, input debugging for a single attack using a simple pattern does not take much time to implement technically as long as the routers are inside a single administrative domain. However, rapidly attributing attackers through paths going through external administrative domains often requires some sort of pre-existing trust relationship between the person performing the attribution and the administrators of those external domains. The external domain administrators need to know if the request is coming from a legitimate source (with a legitimate reason to know the answer), and the requestor needs to know if they are truly communicating with the correct external domain administrators. Thus, trust relationships (manual or automated) must be developed so that when requests are made they will be honored in a timely way.

3.2 Prepositioning Tools and Trust in External Networks is Difficult

Attacks often originate "outside" of the network being attacked. However, as noted above, to be effective many attribution techniques require some sort of cooperation by networks along the path from the attacker to the victim. Gaining such trust, unfortunately, can be very difficult. Even when trust is gained, convincing others to implement attribution tools can be a significant challenge.

3.3 Networks and Systems Can be Configured to Ease Attribution: Changing the Terrain

Networks and systems can be configured to simplify attribution in a variety of ways. Network routers and systems can be hardened against attack, spoofable protocols can be eliminated or limited, cleartext passwords can be eliminated, and broadcast amplification/reflection can be disabled. Attribution can be aided even more directly

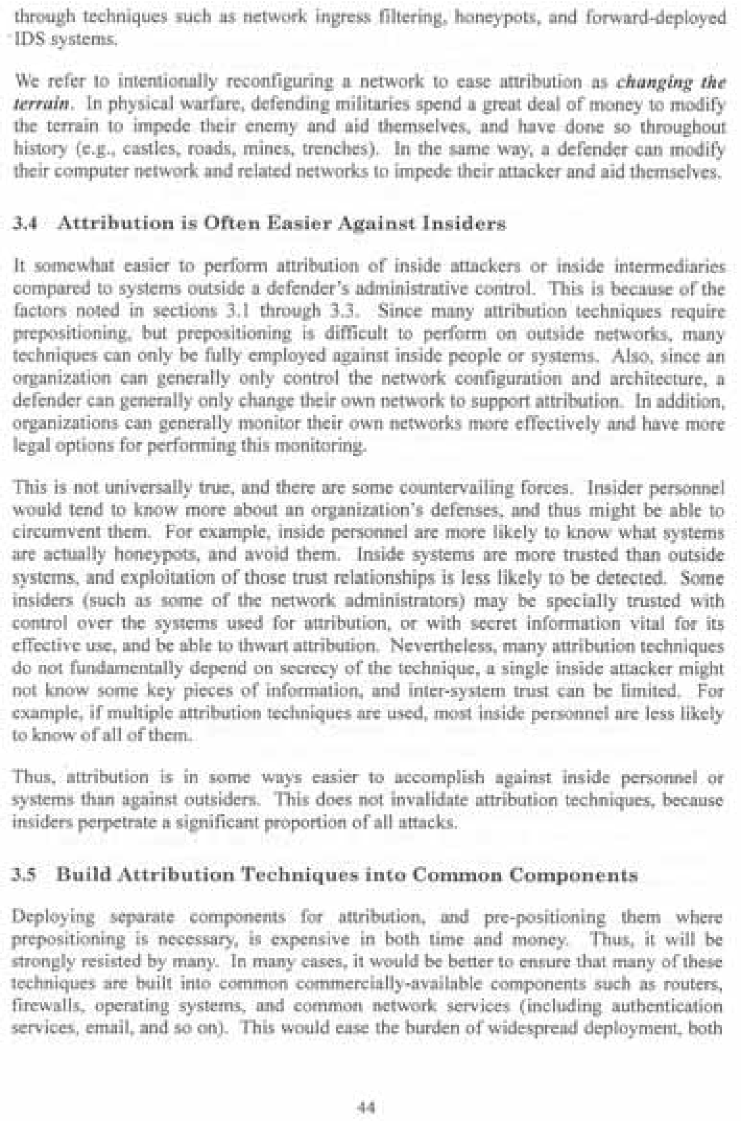

through techniques such as network ingress filtering, honeypots, and forward-deployed IDS systems.

We refer to intentionally reconfiguring a network to ease attribution as ***changing the terrain***. In physical warfare, defending militaries spend a great deal of money to modify the terrain to impede their enemy and aid themselves, and have done so throughout history (e.g., castles, roads, mines, trenches). In the same way, a defender can modify their computer network and related networks to impede their attacker and aid themselves.

3.4 Attribution is Often Easier Against Insiders

It somewhat easier to perform attribution of inside attackers or inside intermediaries compared to systems outside a defender's administrative control. This is because of the factors noted in sections 3.1 through 3.3. Since many attribution techniques require prepositioning, but prepositioning is difficult to perform on outside networks, many techniques can only be fully employed against inside people or systems. Also, since an organization can generally only control the network configuration and architecture, a defender can generally only change their own network to support attribution. In addition, organizations can generally monitor their own networks more effectively and have more legal options for performing this monitoring.

This is not universally true, and there are some countervailing forces. Insider personnel would tend to know more about an organization's defenses, and thus might be able to circumvent them. For example, inside personnel are more likely to know what systems are actually honeypots, and avoid them. Inside systems are more trusted than outside systems, and exploitation of those trust relationships is less likely to be detected. Some insiders (such as some of the network administrators) may be specially trusted with control over the systems used for attribution, or with secret information vital for its effective use, and be able to thwart attribution. Nevertheless, many attribution techniques do not fundamentally depend on secrecy of the technique, a single inside attacker might not know some key pieces of information, and inter-system trust can be limited. For example, if multiple attribution techniques are used, most inside personnel are less likely to know of all of them.

Thus, attribution is in some ways easier to accomplish against inside personnel or systems than against outsiders. This does not invalidate attribution techniques, because insiders perpetrate a significant proportion of all attacks.

3.5 Build Attribution Techniques into Common Components

Deploying separate components for attribution, and pre-positioning them where prepositioning is necessary, is expensive in both time and money. Thus, it will be strongly resisted by many. In many cases, it would be better to ensure that many of these techniques are built into common commercially-available components such as routers, firewalls, operating systems, and common network services (including authentication services, email, and so on). This would ease the burden of widespread deployment, both

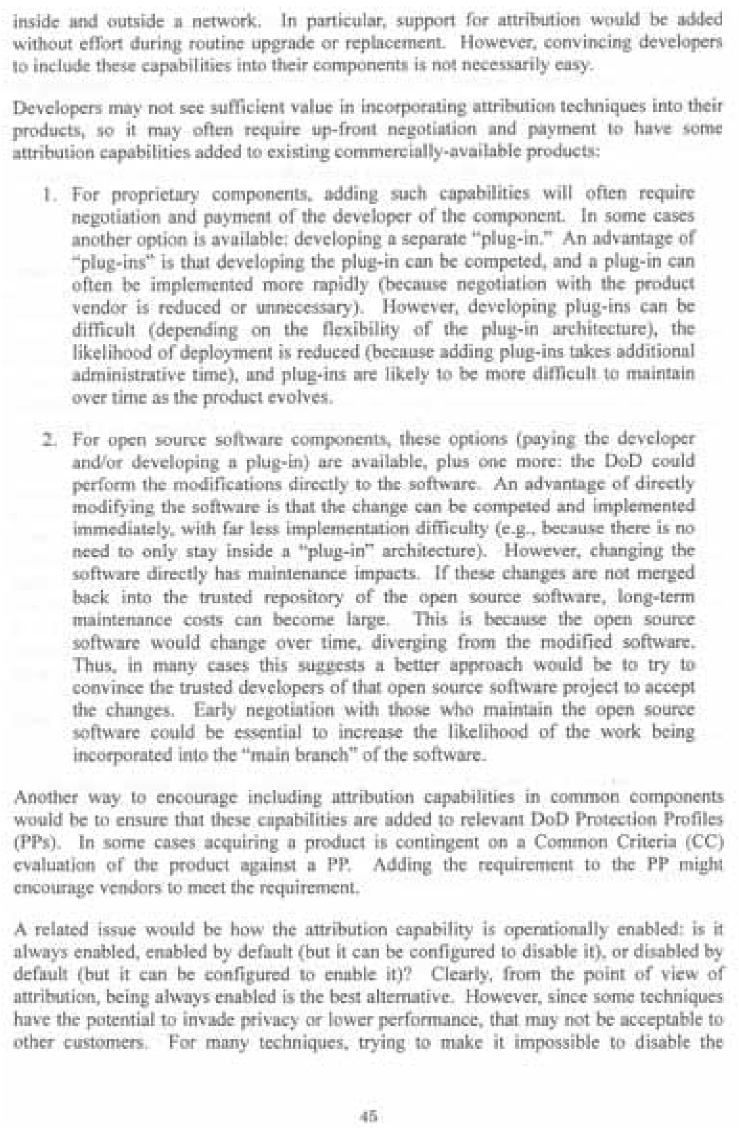

inside and outside a network. In particular, support for attribution would be added without effort during routine upgrade or replacement. However, convincing developers to include these capabilities into their components is not necessarily easy.

Developers may not see sufficient value in incorporating attribution techniques into their products, so it may often require up-front negotiation and payment to have some attribution capabilities added to existing commercially-available products:

1. For proprietary components, adding such capabilities will often require negotiation and payment of the developer of the component. In some cases another option is available: developing a separate "plug-in." An advantage of "plug-ins" is that developing the plug-in can be competed, and a plug-in can often be implemented more rapidly (because negotiation with the product vendor is reduced or unnecessary). However, developing plug-ins can be difficult (depending on the flexibility of the plug-in architecture), the likelihood of deployment is reduced (because adding plug-ins takes additional administrative time), and plug-ins are likely to be more difficult to maintain over time as the product evolves.

2. For open source software components, these options (paying the developer and/or developing a plug-in) are available, plus one more: the DoD could perform the modifications directly to the software. An advantage of directly modifying the software is that the change can be competed and implemented immediately, with far less implementation difficulty (e.g., because there is no need to only stay inside a "plug-in" architecture). However, changing the software directly has maintenance impacts. If these changes are not merged back into the trusted repository of the open source software, long-term maintenance costs can become large. This is because the open source software would change over time, diverging from the modified software. Thus, in many cases this suggests a better approach would be to try to convince the trusted developers of that open source software project to accept the changes. Early negotiation with those who maintain the open source software could be essential to increase the likelihood of the work being incorporated into the "main branch" of the software.

Another way to encourage including attribution capabilities in common components would be to ensure that these capabilities are added to relevant DoD Protection Profiles (PPs). In some cases acquiring a product is contingent on a Common Criteria (CC) evaluation of the product against a PP. Adding the requirement to the PP might encourage vendors to meet the requirement.

A related issue would be how the attribution capability is operationally enabled: is it always enabled, enabled by default (but it can be configured to disable it), or disabled by default (but it can be configured to enable it)? Clearly, from the point of view of attribution, being always enabled is the best alternative. However, since some techniques have the potential to invade privacy or lower performance, that may not be acceptable to other customers. For many techniques, trying to make it impossible to disable the capability is a waste of time. Administrators who truly want to disable the capability can often combine the component with other components to thwart attribution, or use a different component. However, it may be fruitful to try to have some capabilities enabled by default. Any discussion with a developer of a commercially-available product (proprietary or open source) should include discussions on how the capability will be enabled.

3.6 Attribution Requires Funding

Clearly it will cost money to build or buy attribution capabilities (either as separate components or as additions to other components), and there are administrative costs for installing, maintaining, and using components supporting attribution. How will these capabilities be paid for?

There is little evidence that the commercial sector is willing to primarily shoulder the costs of these capabilities. Commercial companies *are* concerned about DDoS attacks, for example, but they are often only interested in reducing their effect, not in actually identifying the attacker. Even if the cost of attribution were reduced to zero (an unlikely scenario), there seems little benefit to identifying attackers in many cases. Bringing a lawsuit against an attacker is quite likely to be very expensive, and it is unlikely that these costs would be recovered. A company bringing a lawsuit risks failing to convict, and it is unclear if a conviction would actually reduce attacks. Some companies are very concerned about unwanted publicity any such lawsuit would entail. Indeed, many companies are unlikely to see attributing attackers as their job. Most commercial companies appear to view identifying attackers as a law enforcement or military task, not a commercial one.

Laws could be enacted requiring certain attribution capabilities be embedded in products for sale, or requiring providers of services to implement certain capabilities. This is unlikely to be effective for most techniques. Most techniques' costs are sufficiently large (especially if research is also required) that any such effort would be strongly resisted. One exception to this may be network ingress filtering; it might be possible to impose a requirement on Internet Service Providers to require that any data entering their backbones go through such a filter, at least for non-peers.

Another approach would be for the U.S. government or DoD to require that certain products must have certain attribution capabilities before they will be acquired. In theory, vendors would add those capabilities and then pass those costs back to users through higher prices. This alternative approach only works with proprietary vendors, who receive funding through usage licenses. This would require convincing the vendor that the cost can be recovered with a profit - an argument that they may not accept. It should be noted that in many markets the U.S. government and/or DoD has a very small portion of the market. A vendor can be disadvantaged by spending money to create a specialized feature only wanted by a small portion of the market, because competing vendors will spend money on capabilities desired by more customers instead. As a result, the DoD may not have any viable application with attribution capabilities, or may only have inferior applications from vendors who decide to add the attribution capability in the hope that the U.S. government or DoD will have to buy it instead of a competing superior product.

If the government (including the DoD) wants the ability to attribute attacks, in many cases the government may need to pay for it directly. One approach is to fund development and deployment of these abilities for widely-used applications. More than one product of each category should be funded, so that the government is not locked into a single product. If the government cannot switch to another product, the vendor will probably raise prices substantially and is less likely to provide good service.

3.7 Standards are Needed

Standards are critically necessary for attribution for the following reasons:

- Interoperability. Many techniques require automated interaction (for speed) between many different organizations and echelons, and it is improbable that exactly the same vendor would be used for all of them. Different organizations will have preferences for different vendors for a variety of reasons (pre-existing relationships, low cost, enhanced functionality, working well with existing infrastructure, and so on). Thus, for attribution to be effective on a wide scale, attribution standards to support inter-vendor interoperability will be necessary.

- Lower cost. By avoiding a proprietary solution from a single vendor, users may be able to select between a variety of offerers. Competition between vendors usually results in a lower price for consumers.

- Lower risk. If a vendor goes out of business or stops supporting a product, another can be used.

- Increased flexibility. If a product doesn't provide what is desired, another product can be used or the extensions can be developed (the latter is particularly easy if it's an open source product).

In theory, a standard could be held secret inside the DoD, but a component that is so widely deployed would be difficult to keep secret. In addition, to attribute external attackers the information would have to be released anyway, so it is almost certainly better to start by developing publicly available standards from the beginning. Note that some will have significant privacy concerns, so standards development should include efforts to address those concerns.

Standards should be open, in particular:

- Standards should be publicly defined and held. This way, no single vendor controls others, permitting competition. Organizations that support development of publicly defined and held standards include the Internet Engineering Task Force (IETF), the World Wide Web Consortium (W3C), the

Institute of Electrical and Electronics Engineers, Inc. (IEEE), the American National Standards Institute (ANSI), and the International Organization for Standardization (ISO). The IETF and W3C are more commonly used for internet-related standards; they are also faster to respond and redistribute standards freely (increasing the number of potential competitors). Thus, for many attribution standards, these organizations might be preferred.

- Standards should not be patent-encumbered. A standard that cannot be implemented without a patent license gives a special advantage to the patent holder(s). Such patents constrain or prevent competition, and thus undermine the advantages of standards listed above. Both the W3C and IETF strongly discourage patent-encumbered standards for these reasons.

Some specifications that could form the basis of standards for attribution include the following:

- IDIP (of CITRA). This has two layers, the application layer (that uses CISL) and the message layer.

- Common Intrusion Specification Language (CISL)

- Ident extensions (for SPIE)

- ICMP Traceback Messages (iTrace) – this is an IETF draft

- Results of the Intrusion Detection Exchange Format Working Group (IETF idwg), including the Intrusion Detection Message Exchange Format (IDMEF) and intrusion detection exchange protocol (IDXP). Information on this working group is available at http://www.ietf.org/html.charters/idwg-charter.html

- RID-DoS, a simple draft protocol for inter-network provider communication. This protocol defines messages for trace request, trace authorization, and source found. More information is in [Moriarty 2003].

- Network Ingress Filtering, IETF RFC 2827 (note that this is already a standard)

3.8 Attribution Techniques Must Be Secured

Clearly, attribution techniques themselves need to be secured. They should be resistant to subversion by an attacker, in particular, attackers should not be able to corrupt the data used for attribution or prevent attribution by directly attacking the attribution components. An attribution technique should not create a new avenue of exploitation (e.g., by creating a new technique for performing a denial of service attack against the system). In many cases, this will require authentication and checking for authorization, and intrusion detection systems should note unauthorized requests.

Many of these techniques require trust between multiple different organizations, making securing these components more difficult.

3.9 Attribution Should Usually Be Hidden from the Attacker

In many cases, an attribution technique should not reveal to an attacker that an attribution process exists or that one is being executed. This is especially difficult if the administrators of domains along the path are colluding with the attacker. "Random" queries can help (where occasionally messages are randomly selected for attribution). However, employing a large number of random queries is only practical if the technique is highly automated and does not interfere with normal operation.

However, in some cases hiding attribution capabilities may not be desirable. Some attackers may decide to not attack at all if they knew that they risked attribution, or may break off an attack if they believe an attribution attack is ongoing. Thus, a known attribution capability may sometimes serve as a useful deterrence. The attribution capability or technique may not even be real, or it may appear to use one technique when in fact another is being used.

Organizations will need to decide if they wish to hide or reveal the existence of an attribution process, as well as the details of that process. Organizations will also need to determine how they intend to implement the hiding or revealing.

3.10 Sensor Placement Is Important

Many attribution techniques are based on the principle of establishing sensors of some kind, then analyzing and using the information from that sensor. Clearly, the information gained depends on the placement of those sensors. Two sensor placement issues are particularly relevant: sensor location, and whether or not the sensor is "in-line":

1. *Location.* Clearly sensors can only be useful if they are placed where they can acquire useful data. This suggests that for attribution purposes, sensors should be placed not only near a defender, but also as near to the attacker as possible so the attacker can be more accurately attributed. To support traceback, sensors must be located at relevant intermediate points as well, to enable a defender to quickly locate the attack path.

2. *In-line or Non-in-line.* Sensors can be placed as in-line sensors or as non-in-line (monitoring) sensors. In-line sensors require that all sensing operations be complete (e.g., initial logging) before additional normal processing occurs. Non-in-line (monitoring) sensors passively observe operations instead, but if they cannot keep up they lose data. The disadvantage of in-line sensors is that they may slow down overall processing. The disadvantage of monitoring sensors is that, if a network or system becomes overwhelmed, such sensors may lose critical data, and this is exactly the time where such data may be needed. Fundamentally, this is a trade-off between the quality of attribution information (in-line) and performance (non-in-line).

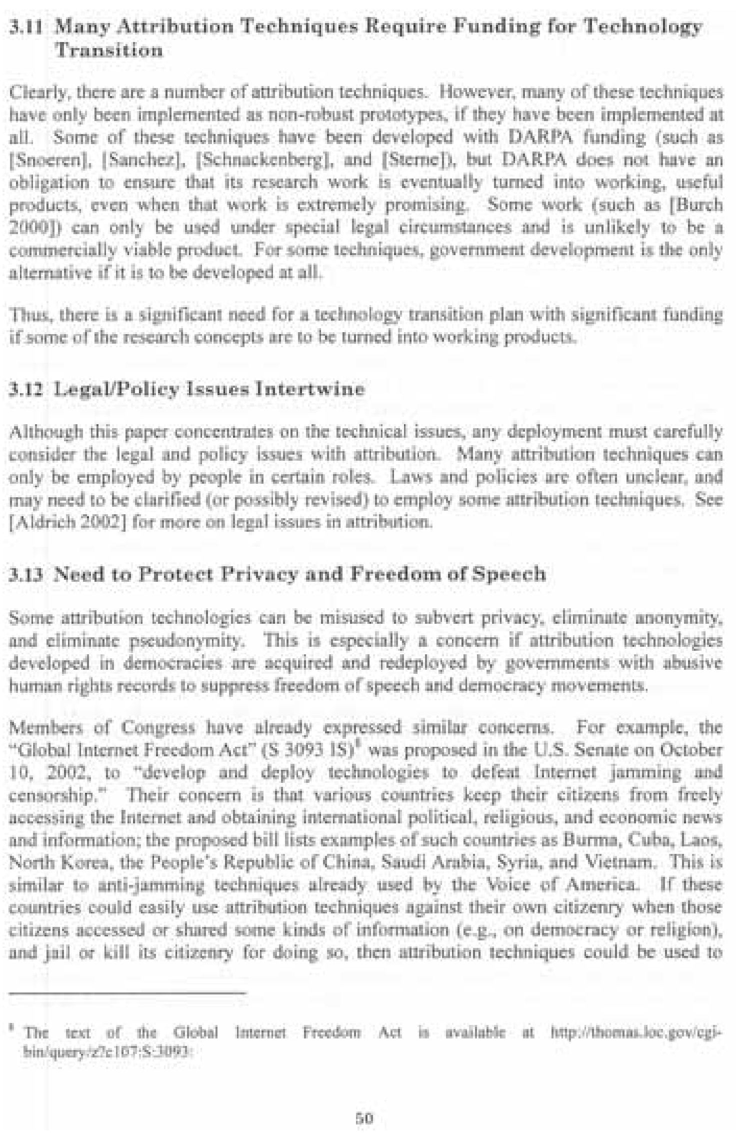

3.11 Many Attribution Techniques Require Funding for Technology Transition

Clearly, there are a number of attribution techniques. However, many of these techniques have only been implemented as non-robust prototypes, if they have been implemented at all. Some of these techniques have been developed with DARPA funding (such as [Snoeren], [Sanchez], [Schnackenberg], and [Sterne]), but DARPA does not have an obligation to ensure that its research work is eventually turned into working, useful products, even when that work is extremely promising. Some work (such as [Burch 2000]) can only be used under special legal circumstances and is unlikely to be a commercially viable product. For some techniques, government development is the only alternative if it is to be developed at all.

Thus, there is a significant need for a technology transition plan with significant funding if some of the research concepts are to be turned into working products.

3.12 Legal/Policy Issues Intertwine

Although this paper concentrates on the technical issues, any deployment must carefully consider the legal and policy issues with attribution. Many attribution techniques can only be employed by people in certain roles. Laws and policies are often unclear, and may need to be clarified (or possibly revised) to employ some attribution techniques. See [Aldrich 2002] for more on legal issues in attribution.

3.13 Need to Protect Privacy and Freedom of Speech

Some attribution technologies can be misused to subvert privacy, eliminate anonymity, and eliminate pseudonymity. This is especially a concern if attribution technologies developed in democracies are acquired and redeployed by governments with abusive human rights records to suppress freedom of speech and democracy movements.

Members of Congress have already expressed similar concerns. For example, the "Global Internet Freedom Act" (S 3093 IS)[8] was proposed in the U.S. Senate on October 10, 2002, to "develop and deploy technologies to defeat Internet jamming and censorship." Their concern is that various countries keep their citizens from freely accessing the Internet and obtaining international political, religious, and economic news and information; the proposed bill lists examples of such countries as Burma, Cuba, Laos, North Korea, the People's Republic of China, Saudi Arabia, Syria, and Vietnam. This is similar to anti-jamming techniques already used by the Voice of America. If these countries could easily use attribution techniques against their own citizenry when those citizens accessed or shared some kinds of information (e.g., on democracy or religion), and jail or kill its citizenry for doing so, then attribution techniques could be used to

[8] The text of the Global Internet Freedom Act is available at http://thomas.loc.gov/cgi-bin/query/z?c107:S:3093:

suppress independent thought in other countries. This result is not in the best interests of the U.S.; indeed, it's not in the best interest of humanity.

Even without the concern of abuse by foreign governments, U.S. citizens will certainly want their privacy against what they may perceive as unwarranted government intrusion. Indeed, the fourth amendment to the U.S. Constitution guarantees that people must be secure "against unreasonable searches and seizures."

Clearly, attribution techniques that pose less danger to privacy should be the ones most encouraged.

3.14 Required Attribution Times Will Continue to Shrink

Some attacks will be slow, over a period of possibly months. But other attacks will be rapid, on the order of milliseconds. [Staniford 2002] discusses techniques for attacking large numbers of systems in very short times. For rapid attacks, attribution techniques will need to rapidly attribute the attacker before the attacker can "get away" or any useful data is hidden by a mass of distracting data. This suggests that automated attribution will be increasingly necessary, and that manual techniques will become increasingly worthless against certain kinds of attacks.

3.15 Attribution is Inherently Limited

All technical means for attribution are inherently limited. These limitations include attribution delay, failed attribution, and misattribution.

3.15.1 Attribution Delays

If an attacker uses a zombie to perform a significantly delayed automated attack, it becomes extremely difficult to attribute the attack path preceding that zombie. Even when the attacker does not intentionally include a delay, there is usually a delay in the defender's response that an attacker can exploit. This delay in the defender's response has many sources: the defender must determine that the message is an attack (or at least that it is worth attributing), perform the attribution, decide on a response, and implement the response. An attacker may be gone before attribution has identified or located the attacker, and/or the attribution may have been made but too late to perform an effective response.

These weaknesses can be partially countered by considering certain kinds of pre-attack activity to be a form of attack and performing attribution. Examples of such pre-attack activity include footprinting, scanning, and enumeration of systems on a network. However, some of these activities are legitimate and/or not really attacks, and they occur constantly on the open Internet. Thus, attribution activities that have a high cost should not normally be used simply to attribute actions that may not be precursors to an attack or are reoccurring.

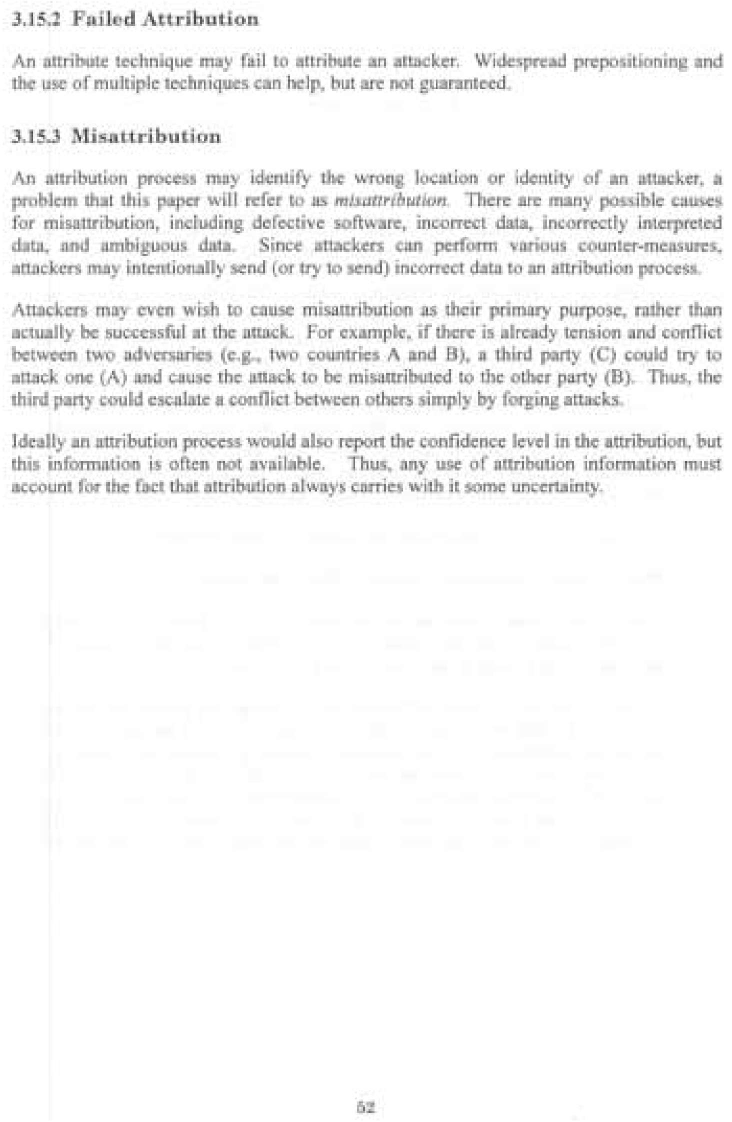

3.15.2 Failed Attribution

An attribute technique may fail to attribute an attacker. Widespread prepositioning and the use of multiple techniques can help, but are not guaranteed.

3.15.3 Misattribution

An attribution process may identify the wrong location or identity of an attacker, a problem that this paper will refer to as *misattribution*. There are many possible causes for misattribution, including defective software, incorrect data, incorrectly interpreted data, and ambiguous data. Since attackers can perform various counter-measures, attackers may intentionally send (or try to send) incorrect data to an attribution process.

Attackers may even wish to cause misattribution as their primary purpose, rather than actually be successful at the attack. For example, if there is already tension and conflict between two adversaries (e.g., two countries A and B), a third party (C) could try to attack one (A) and cause the attack to be misattributed to the other party (B). Thus, the third party could escalate a conflict between others simply by forging attacks.

Ideally an attribution process would also report the confidence level in the attribution, but this information is often not available. Thus, any use of attribution information must account for the fact that attribution always carries with it some uncertainty.

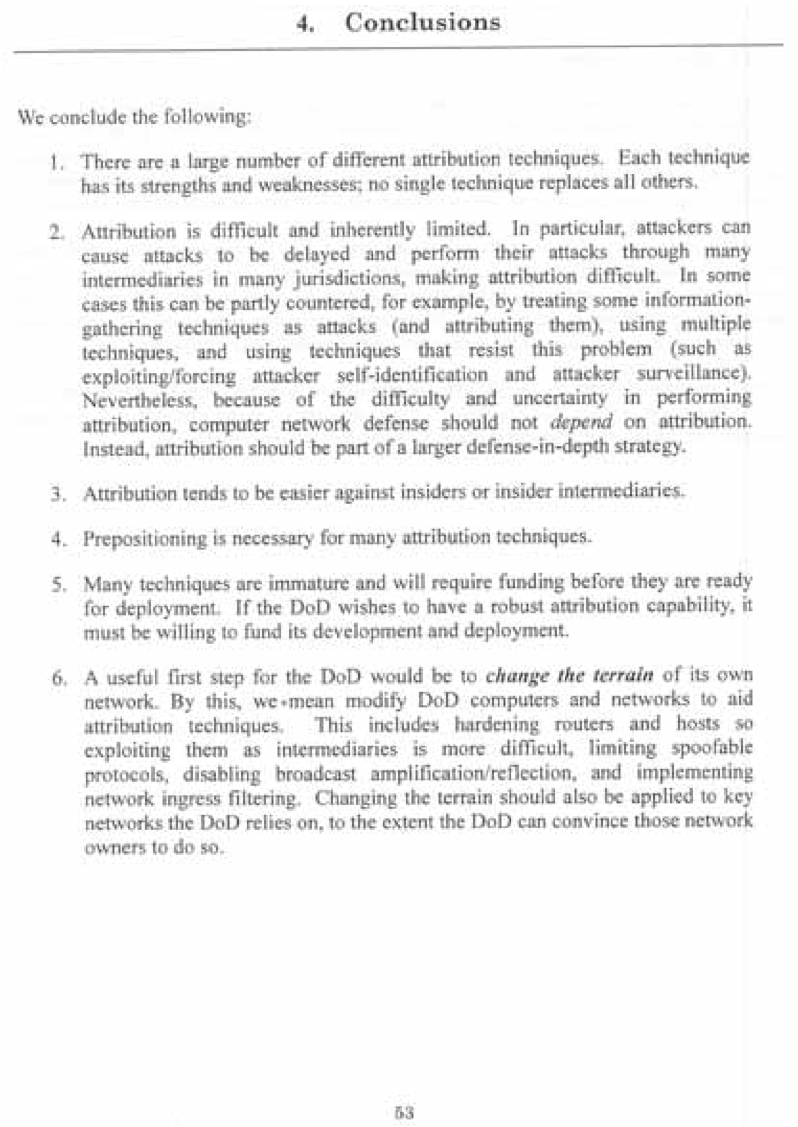

4. Conclusions

We conclude the following:

1. There are a large number of different attribution techniques. Each technique has its strengths and weaknesses; no single technique replaces all others.

2. Attribution is difficult and inherently limited. In particular, attackers can cause attacks to be delayed and perform their attacks through many intermediaries in many jurisdictions, making attribution difficult. In some cases this can be partly countered, for example, by treating some information-gathering techniques as attacks (and attributing them), using multiple techniques, and using techniques that resist this problem (such as exploiting/forcing attacker self-identification and attacker surveillance). Nevertheless, because of the difficulty and uncertainty in performing attribution, computer network defense should not *depend* on attribution. Instead, attribution should be part of a larger defense-in-depth strategy.

3. Attribution tends to be easier against insiders or insider intermediaries.

4. Prepositioning is necessary for many attribution techniques.

5. Many techniques are immature and will require funding before they are ready for deployment. If the DoD wishes to have a robust attribution capability, it must be willing to fund its development and deployment.

6. A useful first step for the DoD would be to ***change the terrain*** of its own network. By this, we mean modify DoD computers and networks to aid attribution techniques. This includes hardening routers and hosts so exploiting them as intermediaries is more difficult, limiting spoofable protocols, disabling broadcast amplification/reflection, and implementing network ingress filtering. Changing the terrain should also be applied to key networks the DoD relies on, to the extent the DoD can convince those network owners to do so.

Appendix. Attribution Technique Taxonomy

There are existing taxonomies of attribution techniques. For example, [Wang 2001a] divides "tracing approaches" into two categories: host-based and network-based, each of which can be classified as being active or passive. However, these taxonomies do not appear to suggest the many possible techniques described in this paper.

Figure A.1 presents a possible taxonomy of the attribution techniques as defined in this paper. The figure shows that the task of attributing attackers can be divided into techniques that actually perform attribution, as well as techniques that modify the environment to simplify attribution. Performing attribution can be further subdivided into techniques that trace backwards from some given point, techniques that send data forward from a given point, and techniques that view network/host from an external view and extract attribution information using that viewpoint.

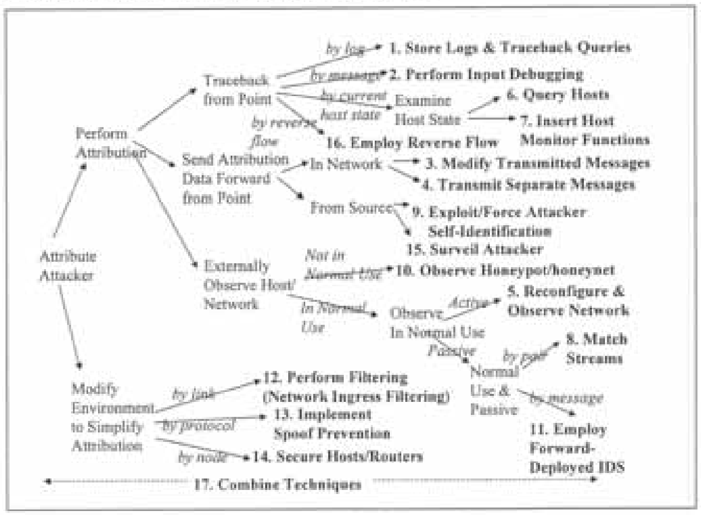
Attribute Attacker
Perform Attribution
Traceback from Point
by log
1. Store Logs & Traceback Queries
by message
2. Perform Input Debugging
by current host state
Examine Host State
6. Query Hosts
7. Insert Host Monitor Functions
by reverse flow
16. Employ Reverse Flow
Send Attribution Data Forward from Point
In Network
3. Modify Transmitted Messages
4. Transmit Separate Messages
From Source
9. Exploit/Force Attacker Self-Identification
15. Surveil Attacker
Externally Observe Host/ Network
Not in Normal Use
10. Observe Honeypot/honeynet
In Normal Use
Observe In Normal Use
Active
5. Reconfigure & Observe Network
Passive
Normal Use & Passive
by pair
8. Match Streams
by message
11. Employ Forward-Deployed IDS
Modify Environment to Simplify Attribution
by link
12. Perform Filtering (Network Ingress Filtering)
by protocol
13. Implement Spoof Prevention
by node
14. Secure Hosts/Routers
17. Combine Techniques

Figure A.1. Attribution Technique Taxonomy

Tracing backwards from some point (typically the defender) back toward the attacker can use different kinds of information. One approach is to require intermediate systems to store logs (historical information of some kind) that can be later queried. Another approach is to request intermediate systems to report the "next time" a message of a

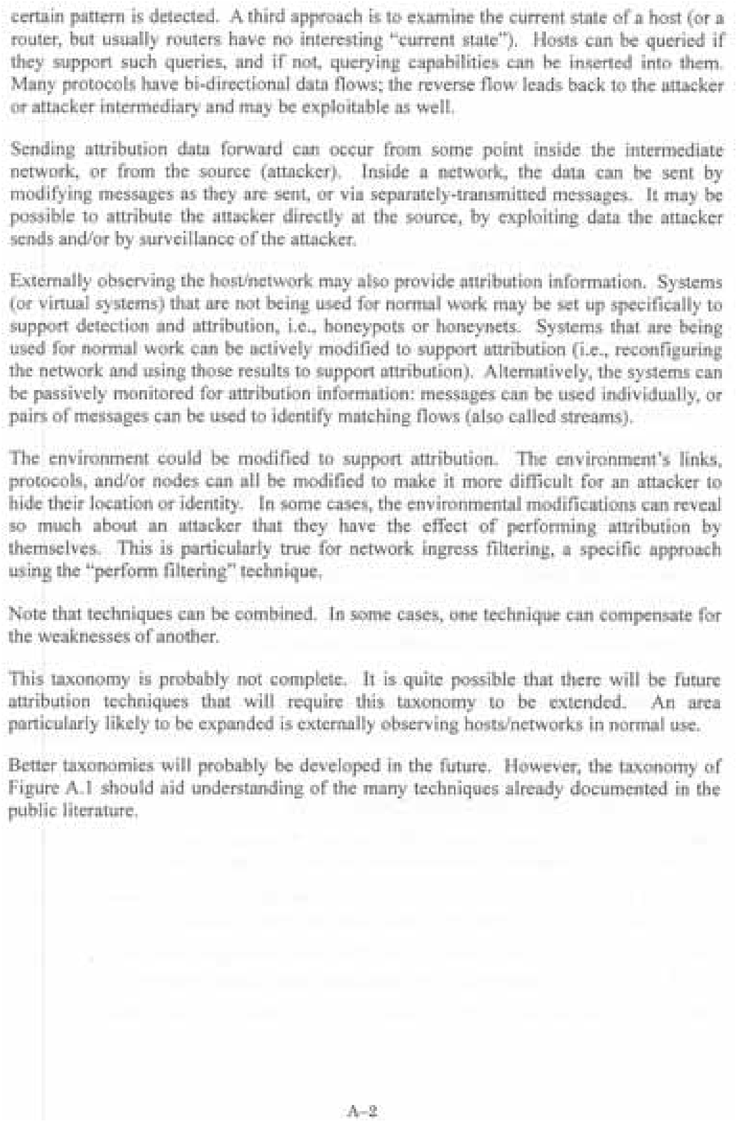

certain pattern is detected. A third approach is to examine the current state of a host (or a router, but usually routers have no interesting "current state"). Hosts can be queried if they support such queries, and if not, querying capabilities can be inserted into them. Many protocols have bi-directional data flows; the reverse flow leads back to the attacker or attacker intermediary and may be exploitable as well.

Sending attribution data forward can occur from some point inside the intermediate network, or from the source (attacker). Inside a network, the data can be sent by modifying messages as they are sent, or via separately-transmitted messages. It may be possible to attribute the attacker directly at the source, by exploiting data the attacker sends and/or by surveillance of the attacker.

Externally observing the host/network may also provide attribution information. Systems (or virtual systems) that are not being used for normal work may be set up specifically to support detection and attribution, i.e., honeypots or honeynets. Systems that are being used for normal work can be actively modified to support attribution (i.e., reconfiguring the network and using those results to support attribution). Alternatively, the systems can be passively monitored for attribution information: messages can be used individually, or pairs of messages can be used to identify matching flows (also called streams).

The environment could be modified to support attribution. The environment's links, protocols, and/or nodes can all be modified to make it more difficult for an attacker to hide their location or identity. In some cases, the environmental modifications can reveal so much about an attacker that they have the effect of performing attribution by themselves. This is particularly true for network ingress filtering, a specific approach using the "perform filtering" technique.

Note that techniques can be combined. In some cases, one technique can compensate for the weaknesses of another.

This taxonomy is probably not complete. It is quite possible that there will be future attribution techniques that will require this taxonomy to be extended. An area particularly likely to be expanded is externally observing hosts/networks in normal use.

Better taxonomies will probably be developed in the future. However, the taxonomy of Figure A.1 should aid understanding of the many techniques already documented in the public literature.

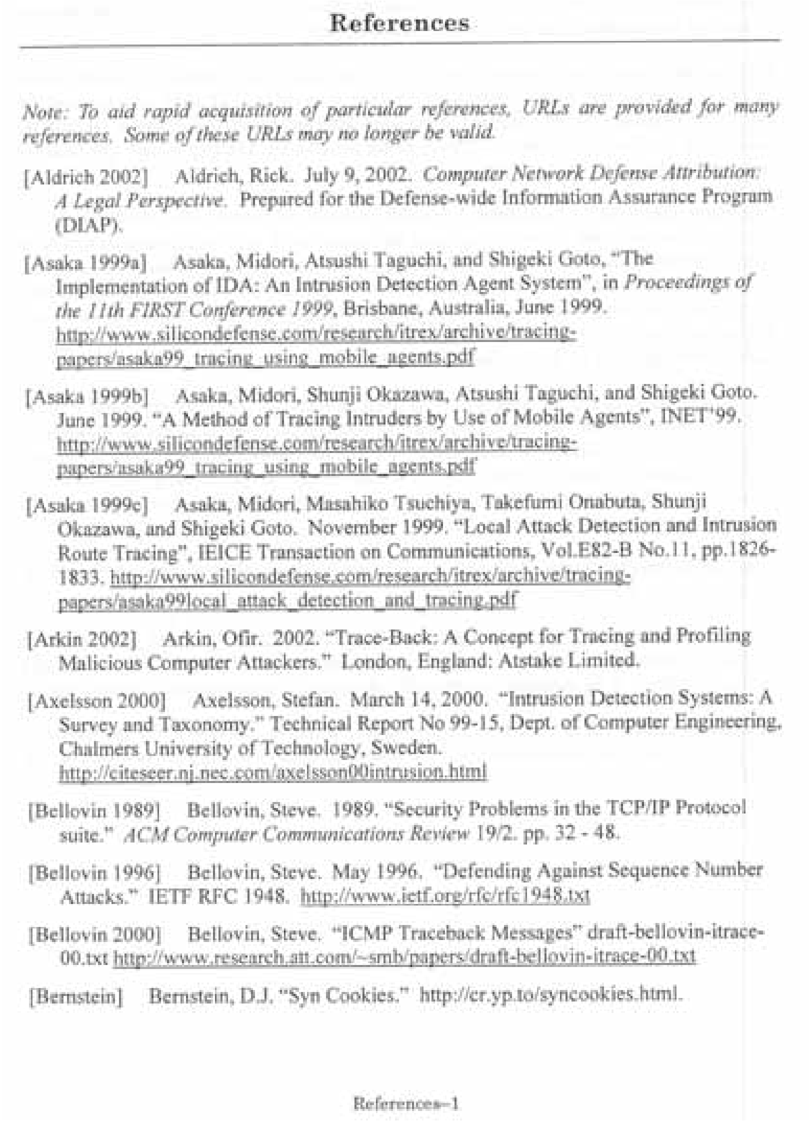

References

Note: To aid rapid acquisition of particular references, URLs are provided for many references. Some of these URLs may no longer be valid.

[Aldrich 2002] Aldrich, Rick. July 9, 2002. *Computer Network Defense Attribution: A Legal Perspective.* Prepared for the Defense-wide Information Assurance Program (DIAP).

[Asaka 1999a] Asaka, Midori, Atsushi Taguchi, and Shigeki Goto, "The Implementation of IDA: An Intrusion Detection Agent System", in *Proceedings of the 11th FIRST Conference 1999*, Brisbane, Australia, June 1999. http://www.silicondefense.com/research/itrex/archive/tracing-papers/asaka99_tracing_using_mobile_agents.pdf

[Asaka 1999b] Asaka, Midori, Shunji Okazawa, Atsushi Taguchi, and Shigeki Goto. June 1999. "A Method of Tracing Intruders by Use of Mobile Agents", INET'99. http://www.silicondefense.com/research/itrex/archive/tracing-papers/asaka99_tracing_using_mobile_agents.pdf

[Asaka 1999c] Asaka, Midori, Masahiko Tsuchiya, Takefumi Onabuta, Shunji Okazawa, and Shigeki Goto. November 1999. "Local Attack Detection and Intrusion Route Tracing", IEICE Transaction on Communications, Vol.E82-B No.11, pp.1826-1833. http://www.silicondefense.com/research/itrex/archive/tracing-papers/asaka99local_attack_detection_and_tracing.pdf

[Arkin 2002] Arkin, Ofir. 2002. "Trace-Back: A Concept for Tracing and Profiling Malicious Computer Attackers." London, England: Atstake Limited.

[Axelsson 2000] Axelsson, Stefan. March 14, 2000. "Intrusion Detection Systems: A Survey and Taxonomy." Technical Report No 99-15, Dept. of Computer Engineering, Chalmers University of Technology, Sweden. http://citeseer.nj.nec.com/axelsson00intrusion.html

[Bellovin 1989] Bellovin, Steve. 1989. "Security Problems in the TCP/IP Protocol suite." *ACM Computer Communications Review* 19/2. pp. 32 - 48.

[Bellovin 1996] Bellovin, Steve. May 1996. "Defending Against Sequence Number Attacks." IETF RFC 1948. http://www.ietf.org/rfc/rfc1948.txt

[Bellovin 2000] Bellovin, Steve. "ICMP Traceback Messages" draft-bellovin-itrace-00.txt http://www.research.att.com/~smb/papers/draft-bellovin-itrace-00.txt

[Bernstein] Bernstein, D.J. "Syn Cookies." http://cr.yp.to/syncookies.html.

[Buchholz] Buchholz, Florian, Thomas E. Daniels, Benjamin Kuperman, Clay Shields. "Packet Tracker: Final Report." CERIAS. http://www.cerias.purdue.edu/events/traceback/packettrackerreport.pdf

[Burch 2000] Burch, H., and B. Cheswick. "Tracing anonymous packets to their approximate source." Proc. Usenix LISA '00, December 2000.

[Carrier 2002] Carrier, Brian, and Clay Shields. 2002. A Recursive Session Token Protocol For Use in Computer Forensics and TCP Traceback. http://citeseer.nj.nec.com/510558.html

[CERT 1996] CERT. September 19, 1996. "TCP SYN Flooding and IP Spoofing Attacks." CERT Advisory CA-1996-21. http://www.cert.org/advisories/CA-1996-21.html

[Chang 1999] Chang, H.Y., R. Narayan, S.F. Wu, B.M. Vetter, M. Brown, J.J. Yuill, X. Wang, C. Sargor, F. Jou, F. Gong. May 1999. "Deciduous: Decentralized Source Identification for Network-based Intrusions," 6th IFIP/IEEE International Symposium on Integrated Network Management, IEEE Communications Society Press. http://www.silicondefense.com/research/itrex/archive/tracing-papers/chang99deciduos.pdf

[Chang 2000] Chang, H.Y., P. Chen, A. Hayatnagarkar, R. Narayan, P. Sheth, N. Vo, C. L. Wu, S.F. Wu, L. Zhang, X. Zhang, F. Gong, F. Jou, C. Sargor, and X. Wu. January 2000. "Design and Implementation of A Real-Time Decentralized Source Identification System for Untrusted IP Packets." *Proceedings of the DARPA Information Survivability Conference & Exposition (DISCEX 2000). IEEE Computer Society Press.* http://www.silicondefense.com/research/itrex/archive/tracing-papers/chang00design_and_implementation_of_realtime.pdf or http://shang.csc.ncsu.edu/papers/desimpdeciduous.pdf

[Chang undated] Chang, Ho-Yen, S.Felix Wu, C. Sargor, X. Wu. Undated. "Towards Tracing Hidden Attackers on Untrusted IP Networks." http://www.silicondefense.com/research/itrex/archive/tracing-papers/chang00towards_tracing_hidden_attackers.pdf

[Cheswick 1992] Cheswick, Bill. January 1992. "An Evening with Berferd: In Which a Cracker is Lured, Endured, and Studied." *Proceedings of the Usenix Winter 92 Conference.*

[Cisco 1] Cisco Systems. Characterizing and Tracing Packet Floods Using Cisco Routers. http://www.cisco.com/warp/public/707/22.html

[Cisco 2] Cisco Systems. IP Source Tracking on Cisco 12000 Series Internet Routers. http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s21/ipst.htm

[Cohen 2002] Cohen, Peter. January 25, 2002. "Timbuktu used to recover stolen iMac." *MacCentral.* Mac Publishing LLC. http://maccentral.macworld.com/news/0201/25.timbuktu.php

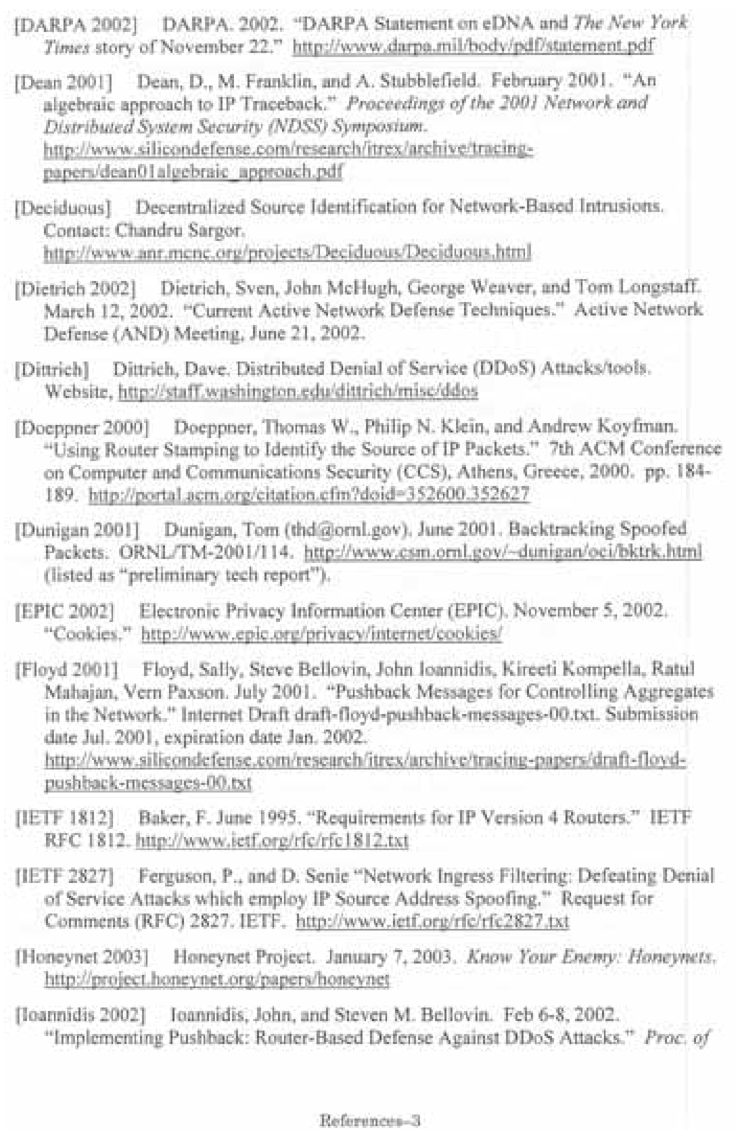

[DARPA 2002] DARPA. 2002. "DARPA Statement on eDNA and *The New York Times* story of November 22." http://www.darpa.mil/body/pdf/statement.pdf

[Dean 2001] Dean, D., M. Franklin, and A. Stubblefield. February 2001. "An algebraic approach to IP Traceback." *Proceedings of the 2001 Network and Distributed System Security (NDSS) Symposium.* http://www.silicondefense.com/research/itrex/archive/tracing-papers/dean01algebraic_approach.pdf

[Deciduous] Decentralized Source Identification for Network-Based Intrusions. Contact: Chandru Sargor. http://www.anr.mcnc.org/projects/Deciduous/Deciduous.html

[Dietrich 2002] Dietrich, Sven, John McHugh, George Weaver, and Tom Longstaff. March 12, 2002. "Current Active Network Defense Techniques." Active Network Defense (AND) Meeting, June 21, 2002.

[Dittrich] Dittrich, Dave. Distributed Denial of Service (DDoS) Attacks/tools. Website, http://staff.washington.edu/dittrich/misc/ddos

[Doeppner 2000] Doeppner, Thomas W., Philip N. Klein, and Andrew Koyfman. "Using Router Stamping to Identify the Source of IP Packets." 7th ACM Conference on Computer and Communications Security (CCS), Athens, Greece, 2000. pp. 184-189. http://portal.acm.org/citation.cfm?doid=352600.352627

[Dunigan 2001] Dunigan, Tom (thd@ornl.gov). June 2001. Backtracking Spoofed Packets. ORNL/TM-2001/114. http://www.csm.ornl.gov/~dunigan/oci/bktrk.html (listed as "preliminary tech report").

[EPIC 2002] Electronic Privacy Information Center (EPIC). November 5, 2002. "Cookies." http://www.epic.org/privacy/internet/cookies/

[Floyd 2001] Floyd, Sally, Steve Bellovin, John Ioannidis, Kireeti Kompella, Ratul Mahajan, Vern Paxson. July 2001. "Pushback Messages for Controlling Aggregates in the Network." Internet Draft draft-floyd-pushback-messages-00.txt. Submission date Jul. 2001, expiration date Jan. 2002. http://www.silicondefense.com/research/itrex/archive/tracing-papers/draft-floyd-pushback-messages-00.txt

[IETF 1812] Baker, F. June 1995. "Requirements for IP Version 4 Routers." IETF RFC 1812. http://www.ietf.org/rfc/rfc1812.txt

[IETF 2827] Ferguson, P., and D. Senie "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing." Request for Comments (RFC) 2827. IETF. http://www.ietf.org/rfc/rfc2827.txt

[Honeynet 2003] Honeynet Project. January 7, 2003. *Know Your Enemy: Honeynets.* http://project.honeynet.org/papers/honeynet

[Ioannidis 2002] Ioannidis, John, and Steven M. Bellovin. Feb 6-8, 2002. "Implementing Pushback: Router-Based Defense Against DDoS Attacks." *Proc. of*

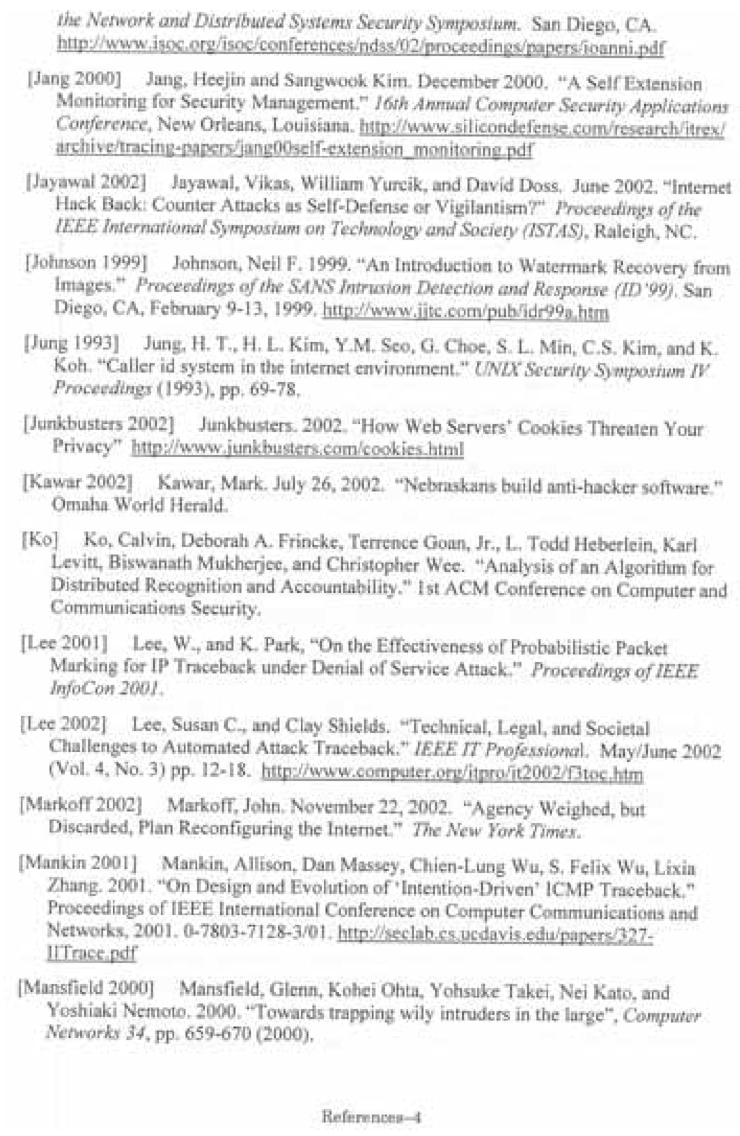

the Network and Distributed Systems Security Symposium. San Diego, CA. http://www.isoc.org/isoc/conferences/ndss/02/proceedings/papers/ioanni.pdf

[Jang 2000] Jang, Heejin and Sangwook Kim. December 2000. "A Self Extension Monitoring for Security Management." *16th Annual Computer Security Applications Conference*, New Orleans, Louisiana. http://www.silicondefense.com/research/itrex/archive/tracing-papers/jang00self-extension_monitoring.pdf

[Jayawal 2002] Jayawal, Vikas, William Yurcik, and David Doss. June 2002. "Internet Hack Back: Counter Attacks as Self-Defense or Vigilantism?" *Proceedings of the IEEE International Symposium on Technology and Society (ISTAS)*, Raleigh, NC.

[Johnson 1999] Johnson, Neil F. 1999. "An Introduction to Watermark Recovery from Images." *Proceedings of the SANS Intrusion Detection and Response (ID'99)*. San Diego, CA, February 9-13, 1999. http://www.jjtc.com/pub/idr99a.htm

[Jung 1993] Jung, H. T., H. L. Kim, Y.M. Seo, G. Choe, S. L. Min, C.S. Kim, and K. Koh. "Caller id system in the internet environment." *UNIX Security Symposium IV Proceedings* (1993), pp. 69-78.

[Junkbusters 2002] Junkbusters. 2002. "How Web Servers' Cookies Threaten Your Privacy" http://www.junkbusters.com/cookies.html

[Kawar 2002] Kawar, Mark. July 26, 2002. "Nebraskans build anti-hacker software." Omaha World Herald.

[Ko] Ko, Calvin, Deborah A. Frincke, Terrence Goan, Jr., L. Todd Heberlein, Karl Levitt, Biswanath Mukherjee, and Christopher Wee. "Analysis of an Algoritlım for Distributed Recognition and Accountability." 1st ACM Conference on Computer and Communications Security.

[Lee 2001] Lee, W., and K. Park, "On the Effectiveness of Probabilistic Packet Marking for IP Traceback under Denial of Service Attack." *Proceedings of IEEE InfoCon 2001*.

[Lee 2002] Lee, Susan C., and Clay Shields. "Technical, Legal, and Societal Challenges to Automated Attack Traceback." *IEEE IT Professiona*l. May/June 2002 (Vol. 4, No. 3) pp. 12-18. http://www.computer.org/itpro/it2002/f3toc.htm

[Markoff 2002] Markoff, John. November 22, 2002. "Agency Weighed, but Discarded, Plan Reconfiguring the Internet." *The New York Times*.

[Mankin 2001] Mankin, Allison, Dan Massey, Chien-Lung Wu, S. Felix Wu, Lixia Zhang. 2001. "On Design and Evolution of 'Intention-Driven' ICMP Traceback." Proceedings of IEEE International Conference on Computer Communications and Networks, 2001. 0-7803-7128-3/01. http://seclab.cs.ucdavis.edu/papers/327-IITrace.pdf

[Mansfield 2000] Mansfield, Glenn, Kohei Ohta, Yohsuke Takei, Nei Kato, and Yoshiaki Nemoto. 2000. "Towards trapping wily intruders in the large", *Computer Networks 34*, pp. 659-670 (2000).

http://www.silicondefense.com/research/itrex/archive/tracing-papers/mansfield00wily_hacker.pdf

[McCullagh 2002] McCullagh, Declan. "Pentagon backs off on Net ID tags." http://zdnet.com.com/2100-1105-966894.html

[Mirkovic] Mirkovic, Jelena, Janice Martin and Peter Reiher. "A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms." Computer Science Department, University of California, Los Angeles. Technical report #020018. http://www.lasr.cs.ucla.edu/ddos/ucla_tech_report_020018.pdf

[Merriam-Webster 1983] Merriam-Webster. 1983. Webster's Ninth New Collegiate Dictionary.Springfield, MA: Merriam-Webster Inc. ISBN 0-87779-508-8.

[Miles 1999] Miles, Stephanie. February 24, 1999. "Intel downplays chip hack report." C/Net. http://news.com.com/2100-1001-222182.html?legacy=cnet

[Moriarty 2003] Moriarty, Kathleen M. February 10, 2003. "Distributed Denial of Service Incident Handling: Real-Time Inter-Network Defense." Work in Progress, Internet-Draft. draft-moriarty-ddos-rid-03.txt. Expires August 10, 2003. ftp://ftp.isi.edu/internet-drafts/draft-moriarty-ddos-rid-03.txt

[NAI] Advanced Intrusion Tracing and Response. http://download.nai.com/products/media/nai/pdf/NAI-Labs-AITR-1-5-01.pdf. See also http://www.pgp.com/research/nailabs/adaptive-networks.asp

[Ohta 2000] Ohta, Kohei, Glenn Mansfield, Yohsuke Takei, Nei Kato. July 2000. "Detection, Defense, and Tracking of Internet-Wide Illegal Access in a Distributed Manner." Proceedings of the 10th Annual Internet Society Conference (INET 2000). http://www.isoc.org/isoc/conferences/inet/00/cdproceedings/1f/1f_2.htm

[ORNL] Oak Ridge National Lab. "Backtracking Spoofed Packets" (web site) http://www.csm.ornl.gov/~dunigan/oci/bktrk.html

[Park 2000] Park, Kihong, and Heejo Lee. June 2000. "On the Effectiveness of Probabilistic Packet Marking for IP Traceback under Denial of Service Attack" Technical Report CSD-00-013, Department of Computer Sciences, Purdue University. http://www.silicondefense.com/research/itrex/archive/tracing-papers/park00effectiveness_technical_paper.pdf. An extended abstract of the report was published with the same title in *Proceedings of the IEEE INFOCOM '01*, pp. 338-347, 2001.

[Partridge 2001] Partridge, C., C. Jones, D. Waitzman, A. Snoeren. November 2001. "New Protocols to Support Internet Traceback." http://www.ir.bbn.com/documents/internet_drafts/draft-partridge-ippt-discuss-00.txt

[Provos 2002] Provos, Niels, and Peter Honeyman. "Detecting Steganographic Content on the Internet." http://isoc.org/isoc/conferences/ndss/02/proceedings

[Purdue 2000] Purdue. Results of the "Attack Traceback Summit Proceedings" of September 6-8, 2000. http://www.cerias.purdue.edu/events/traceback/

[Ricciuti 1999] Ricciuti, Mike. March 7, 1999. "Microsoft admits privacy problem, plans fix." C|Net. http://news.com.com/2100-1040-222673.html?legacy=cnet

[Sager 1998] Sager, Glenn. November 20, 1998. "Security Fun with OCxmon and cflowd." Internet-2 Measurement Working Group. http://www.caida.org/projects/ngi/content/security/1198

[Sanchez 2001] Sanchez, Luis A., Walter C. Milliken, Alex C. Snoeren, Fabrice Tchakountio, Christine E. Jones, Stephen T. Kent, Craig Partrige, and W. Timothy Strayer. "Hardware Support for a Hash-Based IP Traceback." Proceedings of the 2nd DARPA Information Survivability Conference and Exposition (DISCEX II), pp. 146-152, Anaheim, CA, June 2001. http://www.ir.bbn.com/projects/SPIE/pubs/discex01.html

[SANS 2000a] SANS. February 23, 2000. Consensus Roadmap for Defeating Distributed Denial of Service Attacks: A Project of the Partnership for Critical Infrastructure Security. Version 1.10. http://www.sans.org/ddos_roadmap.htm

[SANS 2000b] SANS. Egress Filtering v 0.2. February 29, 2000. http://www.sans.org/y2k/egress.htm

[Savage 2000] Savage, Stefan, David Wetherall, Anna Karlin and Tom Anderson, "Practical Network Support for IP Traceback", *Proceedings of the 2000 ACM SIGCOMM Conference*, pp. 295-306, Stockholm, Sweden, August 2000. See [Savage 2001]. http://www.cs.washington.edu/homes/savage/traceback.html

[Savage 2001] Stefan Savage , David Wetherall , Anna Karlin , and Tom Anderson. June 2001. Network Support for IP Traceback. *IEEE/ACM Transactions on Networking (TON)*. Volume 9, Issue 3.

[Schaeffer 2000] Schaeffer, Richard C., Jr. "TESTIMONY of Richard C. Schaeffer, Jr., Director, Infrastructure and Information Assurance, Office of the Assistant Secretary of Defense (Command, Control, Communication, and Intelligence) before a hearing of the Subcommittee on Government Management, Information, and Technology, July 26, 2000, Computer Security: Cyber Attacks - War without Borders."

[Schnackenberg 2000] Schnackenberg, D., K. Djahandari, and D. Sterne. "Infrastructure for intrusion detection and response." *Proc. First DARPA Information Survivability Conference and Exposition*, Jan. 2000. http://www.silicondefense.com/research/itrex/archive/tracing-papers/schnackenberg00DISCEX_intrusion_detection_and_response.pdf

[Sharp] Sharp, Walter Gary. (then at MITRE) "Key Legal Implications of Computer Network Defense." http://www.blackhat.com/html/bh-multi-media-archives.html

[Silicon Defense] Silicon Defense. *Traceback and Related Papers Archive.* http://www.silicondefense.com/research/itrex/archive/tracing-papers

[Smith] Smith, Richard M. "FAQ: Web Bugs." Verified October 24, 2002. http://www.privacyfoundation.org/resources/webbug.asp

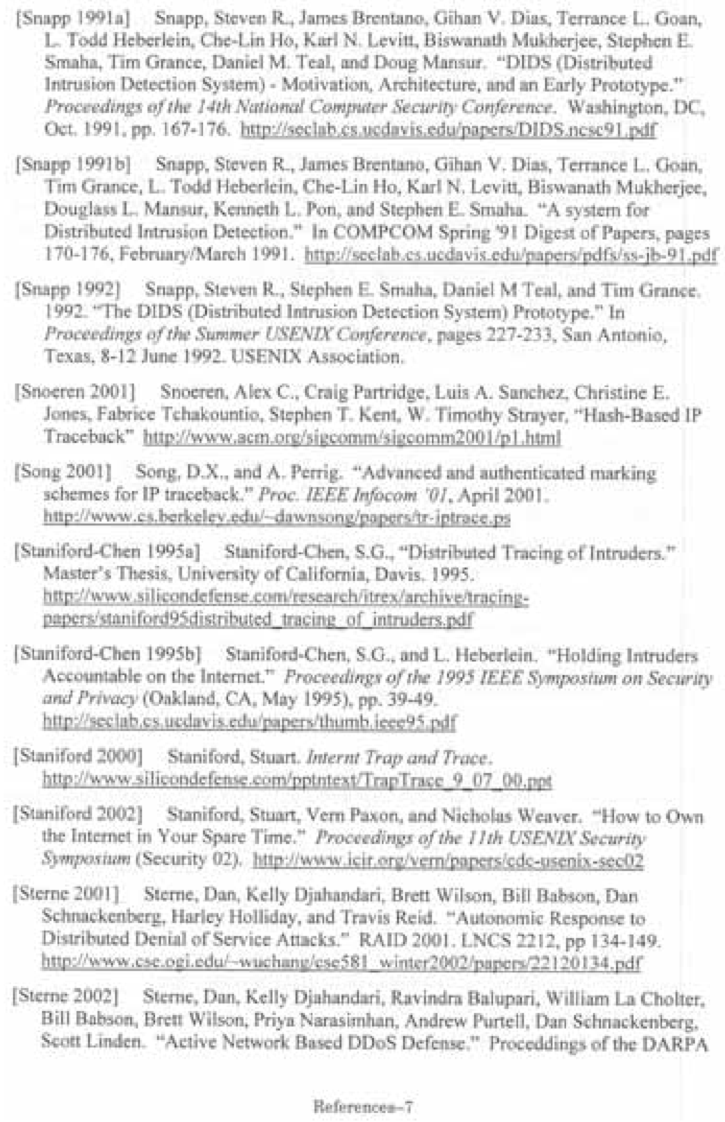

[Snapp 1991a] Snapp, Steven R., James Brentano, Gihan V. Dias, Terrance L. Goan, L. Todd Heberlein, Che-Lin Ho, Karl N. Levitt, Biswanath Mukherjee, Stephen E. Smaha, Tim Grance, Daniel M. Teal, and Doug Mansur. "DIDS (Distributed Intrusion Detection System) - Motivation, Architecture, and an Early Prototype." *Proceedings of the 14th National Computer Security Conference*. Washington, DC, Oct. 1991, pp. 167-176. http://seclab.cs.ucdavis.edu/papers/DIDS.ncsc91.pdf

[Snapp 1991b] Snapp, Steven R., James Brentano, Gihan V. Dias, Terrance L. Goan, Tim Grance, L. Todd Heberlein, Che-Lin Ho, Karl N. Levitt, Biswanath Mukherjee, Douglass L. Mansur, Kenneth L. Pon, and Stephen E. Smaha. "A system for Distributed Intrusion Detection." In COMPCOM Spring '91 Digest of Papers, pages 170-176, February/March 1991. http://seclab.cs.ucdavis.edu/papers/pdfs/ss-jb-91.pdf

[Snapp 1992] Snapp, Steven R., Stephen E. Smaha, Daniel M Teal, and Tim Grance. 1992. "The DIDS (Distributed Intrusion Detection System) Prototype." In *Proceedings of the Summer USENIX Conference*, pages 227-233, San Antonio, Texas, 8-12 June 1992. USENIX Association.

[Snoeren 2001] Snoeren, Alex C., Craig Partridge, Luis A. Sanchez, Christine E. Jones, Fabrice Tchakountio, Stephen T. Kent, W. Timothy Strayer, "Hash-Based IP Traceback" http://www.acm.org/sigcomm/sigcomm2001/p1.html

[Song 2001] Song, D.X., and A. Perrig. "Advanced and authenticated marking schemes for IP traceback." *Proc. IEEE Infocom '01*, April 2001. http://www.cs.berkeley.edu/~dawnsong/papers/tr-iptrace.ps

[Staniford-Chen 1995a] Staniford-Chen, S.G., "Distributed Tracing of Intruders." Master's Thesis, University of California, Davis. 1995. http://www.silicondefense.com/research/itrex/archive/tracing-papers/staniford95distributed_tracing_of_intruders.pdf

[Staniford-Chen 1995b] Staniford-Chen, S.G., and L. Heberlein. "Holding Intruders Accountable on the Internet." *Proceedings of the 1995 IEEE Symposium on Security and Privacy* (Oakland, CA, May 1995), pp. 39-49. http://seclab.cs.ucdavis.edu/papers/thumb.ieee95.pdf

[Staniford 2000] Staniford, Stuart. *Internt Trap and Trace*. http://www.silicondefense.com/pptntext/TrapTrace_9_07_00.ppt

[Staniford 2002] Staniford, Stuart, Vern Paxon, and Nicholas Weaver. "How to Own the Internet in Your Spare Time." *Proceedings of the 11th USENIX Security Symposium* (Security 02). http://www.icir.org/vern/papers/cdc-usenix-sec02

[Sterne 2001] Sterne, Dan, Kelly Djahandari, Brett Wilson, Bill Babson, Dan Schnackenberg, Harley Holliday, and Travis Reid. "Autonomic Response to Distributed Denial of Service Attacks." RAID 2001. LNCS 2212, pp 134-149. http://www.cse.ogi.edu/~wuchang/cse581_winter2002/papers/22120134.pdf

[Sterne 2002] Sterne, Dan, Kelly Djahandari, Ravindra Balupari, William La Cholter, Bill Babson, Brett Wilson, Priya Narasimhan, Andrew Purtell, Dan Schnackenberg, Scott Linden. "Active Network Based DDoS Defense." Proceddings of the DARPA

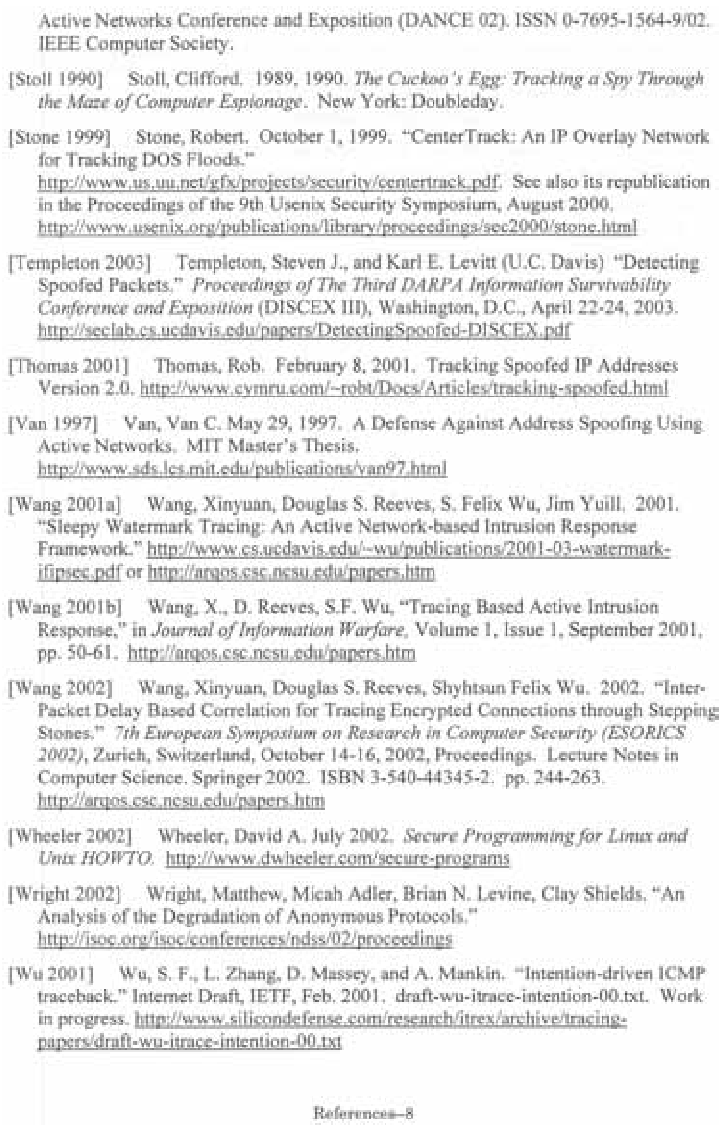

Active Networks Conference and Exposition (DANCE 02). ISSN 0-7695-1564-9/02. IEEE Computer Society.

[Stoll 1990] Stoll, Clifford. 1989, 1990. *The Cuckoo's Egg: Tracking a Spy Through the Maze of Computer Espionage*. New York: Doubleday.

[Stone 1999] Stone, Robert. October 1, 1999. "CenterTrack: An IP Overlay Network for Tracking DOS Floods." http://www.us.uu.net/gfx/projects/security/centertrack.pdf. See also its republication in the Proceedings of the 9th Usenix Security Symposium, August 2000. http://www.usenix.org/publications/library/proceedings/sec2000/stone.html

[Templeton 2003] Templeton, Steven J., and Karl E. Levitt (U.C. Davis) "Detecting Spoofed Packets." *Proceedings of The Third DARPA Information Survivability Conference and Exposition* (DISCEX III), Washington, D.C., April 22-24, 2003. http://seclab.cs.ucdavis.edu/papers/DetectingSpoofed-DISCEX.pdf

[Thomas 2001] Thomas, Rob. February 8, 2001. Tracking Spoofed IP Addresses Version 2.0. http://www.cymru.com/~robt/Docs/Articles/tracking-spoofed.html

[Van 1997] Van, Van C. May 29, 1997. A Defense Against Address Spoofing Using Active Networks. MIT Master's Thesis. http://www.sds.lcs.mit.edu/publications/van97.html

[Wang 2001a] Wang, Xinyuan, Douglas S. Reeves, S. Felix Wu, Jim Yuill. 2001. "Sleepy Watermark Tracing: An Active Network-based Intrusion Response Framework." http://www.cs.ucdavis.edu/~wu/publications/2001-03-watermark-ifipsec.pdf or http://arqos.csc.ncsu.edu/papers.htm

[Wang 2001b] Wang, X., D. Reeves, S.F. Wu, "Tracing Based Active Intrusion Response," in *Journal of Information Warfare,* Volume 1, Issue 1, September 2001, pp. 50-61. http://arqos.csc.ncsu.edu/papers.htm

[Wang 2002] Wang, Xinyuan, Douglas S. Reeves, Shyhtsun Felix Wu. 2002. "Inter-Packet Delay Based Correlation for Tracing Encrypted Connections through Stepping Stones." *7th European Symposium on Research in Computer Security (ESORICS 2002)*, Zurich, Switzerland, October 14-16, 2002, Proceedings. Lecture Notes in Computer Science. Springer 2002. ISBN 3-540-44345-2. pp. 244-263. http://arqos.csc.ncsu.edu/papers.htm

[Wheeler 2002] Wheeler, David A. July 2002. *Secure Programming for Linux and Unix HOWTO.* http://www.dwheeler.com/secure-programs

[Wright 2002] Wright, Matthew, Micah Adler, Brian N. Levine, Clay Shields. "An Analysis of the Degradation of Anonymous Protocols." http://isoc.org/isoc/conferences/ndss/02/proceedings

[Wu 2001] Wu, S. F., L. Zhang, D. Massey, and A. Mankin. "Intention-driven ICMP traceback." Internet Draft, IETF, Feb. 2001. draft-wu-itrace-intention-00.txt. Work in progress. http://www.silicondefense.com/research/itrex/archive/tracing-papers/draft-wu-itrace-intention-00.txt

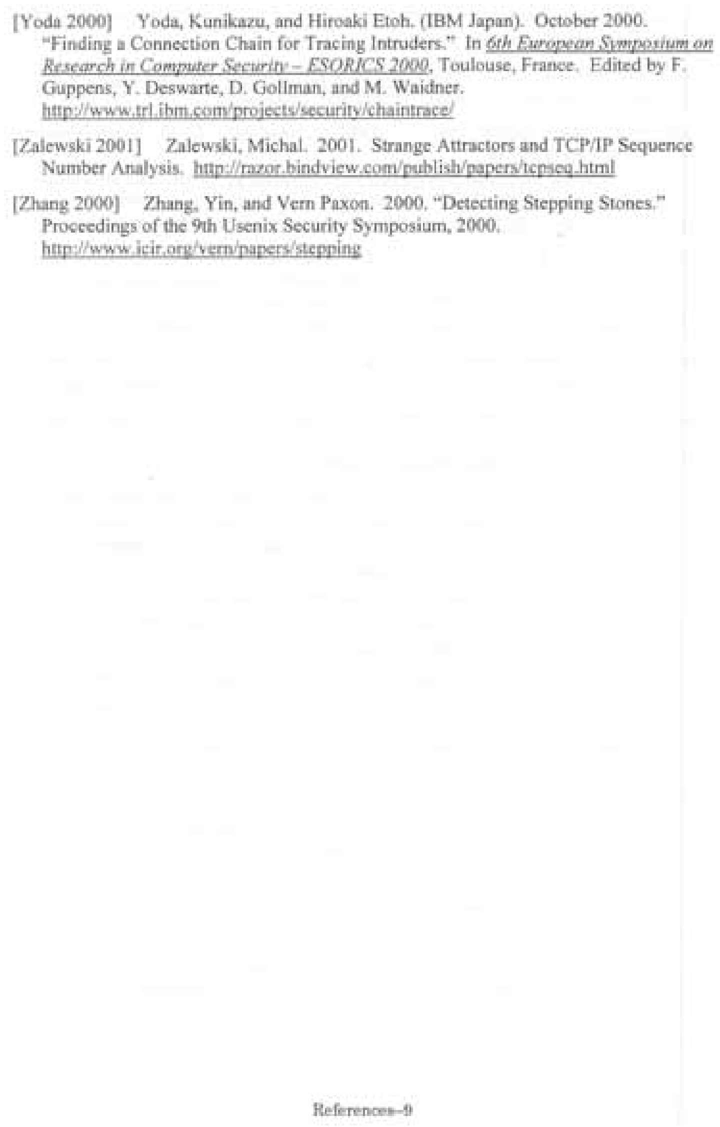

[Yoda 2000] Yoda, Kunikazu, and Hiroaki Etoh. (IBM Japan). October 2000. "Finding a Connection Chain for Tracing Intruders." In *6th European Symposium on Research in Computer Security – ESORICS 2000*, Toulouse, France. Edited by F. Guppens, Y. Deswarte, D. Gollman, and M. Waidner. http://www.trl.ibm.com/projects/security/chaintrace/

[Zalewski 2001] Zalewski, Michal. 2001. Strange Attractors and TCP/IP Sequence Number Analysis. http://razor.bindview.com/publish/papers/tcpseq.html

[Zhang 2000] Zhang, Yin, and Vern Paxon. 2000. "Detecting Stepping Stones." Proceedings of the 9th Usenix Security Symposium, 2000. http://www.icir.org/vern/papers/stepping

Acronyms and Abbreviations

ACC	Aggregate-based Congestion Control
ACK	Acknowledge(d)
AH	Authentication Header
ANSI	American National Standards Institute
BGP	Border Gateway Protocol
CC	Common Criteria
CEF	Cisco Express Forwarding
CERT	formerly Computer Emergency Response Team; now just CERT
CISIE	Caller Identification System in the Internet Environment
CISL	Common Intrusion Specification Language
CITRA	Cooperative Intrusion Traceback and Response Architecture
CNA	Computer Network Attack
CND	Computer Network Defense
CPU	Central Processing Unit
DARPA	Defense Advanced Research Projects Agency
dCEF	Distributed CEF
DDoS	Distributed Denial of Service (a type of attack)
DHCP	Dynamic Host Configuration Protocol
DIAP	Defense-wide Information Assurance Program
DoD	Department of Defense
DOS	Denial of Service
ftp	File Transfer Protocol
HTML	HyperText Markup Language
HTTP	HyperText Transport Protocol
ICMP	Internet Control Message Protocol
ID	Identification
IDA	Institute for Defense Analyses
IDIP	Intruder Detection and Isolation Protocol
IDMEF	Intrusion Detection Message Exchange Format
IDS	Intrusion Detection System
IEEE	Institute of Electrical and Electronics Engineers, Inc.
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPSEC	IP Security
IPv4	Internet Protocol, version 4
IPv6	Internet Protocol, version 6
ISO	International Organization for Standardization
ISP	Internet Service Provider

ISP	Internet Service Provider
iTrace	ICMP Traceback
LAN	Local Area Network
MLSI	Mark Left by Suspected Intruder
MTA	Mail Transfer Agent
NAI Labs	Network Associates Laboratories
ORNL	Oak Ridge National Laboratory
OSI	Open Systems Interconnection
P2P	Peer to Peer
PP	Protection Profile
PPM	Probabilistic Packet Marking
PR/TT	Pen Register/Trap and Trace
RFC	Request for Comments
RID	Remote. Intrusion. Detector
SPIE	Source Path Isolation Engine
STOP	Session Token Protocol
SWT	Sleepy Watermark Tracing
Syn	Synchronize
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
U.S.	United States
WAN	Wide Area Network
W3C	World Wide Web Consortium

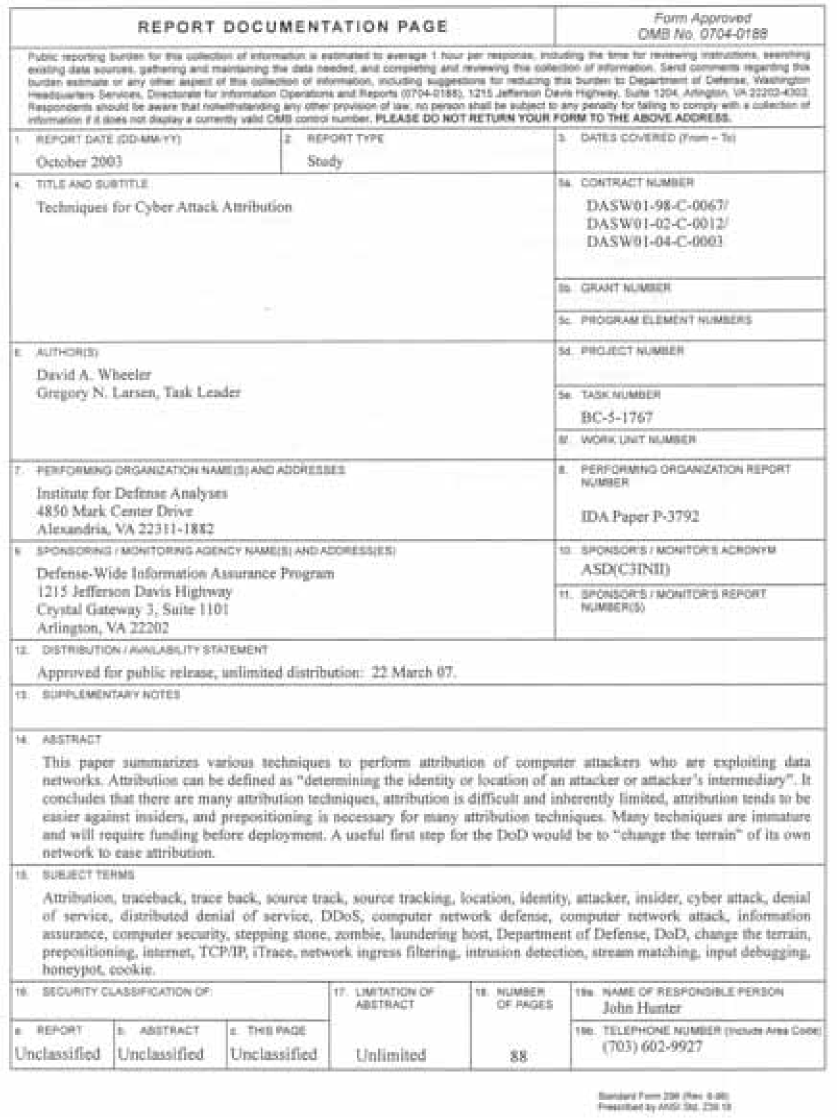

REPORT DOCUMENTATION PAGE

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 this 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 Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YY)
October 2003

2. REPORT TYPE
Study

3. DATES COVERED (From – To)

4. TITLE AND SUBTITLE
Techniques for Cyber Attack Attribution

5a. CONTRACT NUMBER
DASW01-98-C-0067/
DASW01-02-C-0012/
DASW01-04-C-0003

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBERS

6. AUTHOR(S)
David A. Wheeler
Gregory N. Larsen, Task Leader

5d. PROJECT NUMBER

5e. TASK NUMBER
BC-5-1767

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESSES
Institute for Defense Analyses
4850 Mark Center Drive
Alexandria, VA 22311-1882

8. PERFORMING ORGANIZATION REPORT NUMBER
IDA Paper P-3792

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Defense-Wide Information Assurance Program
1215 Jefferson Davis Highway
Crystal Gateway 3, Suite 1101
Arlington, VA 22202

10. SPONSOR'S / MONITOR'S ACRONYM
ASD(C3INII)

11. SPONSOR'S / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release, unlimited distribution: 22 March 07.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This paper summarizes various techniques to perform attribution of computer attackers who are exploiting data networks. Attribution can be defined as "determining the identity or location of an attacker or attacker's intermediary". It concludes that there are many attribution techniques, attribution is difficult and inherently limited, attribution tends to be easier against insiders, and prepositioning is necessary for many attribution techniques. Many techniques are immature and will require funding before deployment. A useful first step for the DoD would be to "change the terrain" of its own network to ease attribution.

15. SUBJECT TERMS
Attribution, traceback, trace back, source track, source tracking, location, identity, attacker, insider, cyber attack, denial of service, distributed denial of service, DDoS, computer network defense, computer network attack, information assurance, computer security, stepping stone, zombie, laundering host, Department of Defense, DoD, change the terrain, prepositioning, internet, TCP/IP, iTrace, network ingress filtering, intrusion detection, stream matching, input debugging, honeypot, cookie.

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON John Hunter
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include Area Code)
Unclassified	Unclassified	Unclassified	Unlimited	88	(703) 602-9927

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Biography for David A. Wheeler

Dr. David A. Wheeler has been in the computing field since 1980, and is an expert on computer security, open source software, open standards, and software development approaches. He has worked at the Institute for Defense Analyses (IDA) since 1987.

As part of his work in computer security, Dr. Wheeler led the development of "Key Practices" guidance to perform supply chain risk management in the U.S. Department of Defense. He is co-author of the DoD/NDIA document "Engineering for System Assurance." He has written a book ("Secure Programming for Linux and Unix

HOWTO"), written various articles (including the "Secure Programmer" series), and given many presentations on how to develop secure software. His Ph.D. dissertation, "Fully Countering Trusting Trust Through Diverse Double-Compiling," proves and demonstrates that the "Diverse Double-Compiling" (DDC) process (a process he named) counters the "trusting trust" attack. The trusting trust attack is a computer attack that previously had no effective countermeasure. He is also the author of an IDA report surveying how to attribute cyber attackers, "Techniques for Cyber Attack Attribution."

Dr. Wheeler lectures worldwide as an invited expert on open source software and/or security, including in Belgium, Brazil, Saudi Arabia, and numerous times in the U.S. As part of his work in open source software, he helped develop the official DoD memo "Clarifying Guidance Regarding Open Source Software (OSS)" and was the primary author of the supporting document "DoD Open Source Software (OSS) FAQ."

Dr. Wheeler has been involved in many efforts related to open standards. He represented the Missile Defense Agency (MDA) in the development of the DoD Information Technology Standards Registry (DISR), formerly named the Joint Technical Architecture (JTA). He also initiated and led development of OpenFormula, an open standard for the interchange of spreadsheet formulas which is planned to be part of the OpenDocument standard (ISO/IEC 26300).

Dr. Wheeler has long been involved in efforts to improve software development approaches and technology. For example, he led the evaluation of software development processes and software development environments across missile defense programs. He is the lead editor and co-author of the IEEE Computer Society Press book "Software Inspection: An Industry Best Practice" and is the sole author of Springer-Verlag's book "Ada 95: The Lovelace Tutorial." His more recent work has focused on how to change software development practices to improve the security and assurance of the resulting software.

Chairman WU. Thank you very much, Dr. Wheeler.

Mr. Knake, please proceed.

STATEMENT OF ROBERT KNAKE, INTERNATIONAL AFFAIRS FELLOW, COUNCIL ON FOREIGN RELATIONS

Mr. KNAKE. Thank you, Chairman Wu and distinguished Members of the House Subcommittee on Technology and Innovation for the opportunity to discuss the role of attack attribution in preventing cyber attacks. My name is Rob Knake. I am an international affairs fellow at the Council on Foreign Relations where I have spent the last year studying state conflict in cyberspace, so I will focus my comments on the attribution problem at that level first.

It is my view that the problem of attribution has been largely overstated. For the high-end threats that my work is focused on, attribution will almost certainly be possible due to the limited number of actors that possess the capability to present a national security challenge in cyberspace. While we have all heard tales of teenagers with laptops sending viruses across the Internet, these sorts of threats do not amount to a national security concern and cannot cause the type of havoc that many envision a cyber attack can. Estimates vary, but analysts who have studied the capabilities of both foreign governments and private groups have concluded that no more than 100 groups and possibly as few as four foreign militaries possess the capability to cause real-world harm through cyber attacks. Moreover, such an attack would take significant investments of both time and money and teams of highly skilled specialists. While technical attribution may only provide limited evidence of who was behind the attack, traditional intelligence and law enforcement investigation can make up the difference. I have no doubt that in the event of a so-called cyber Pearl Harbor, cyber 9/11 or cyber Katrina, that we will be able to amass enough evidence for the President to take action.

For lower-level threats, everything from nuisance behavior like spam to cyber criminal activity, many in the cybersecurity community have viewed the development of ironclad attribution in real time as the Holy Grail. In one widely discussed scenario, all packets could be labeled with a unique identifier that would tie it to an individual, a so-called license plate for the Internet. It is my view that such a concept would be far more useful for authoritarian regimes to monitor and control Internet use by their citizens than it would be in combating cyber warfare, crime and nuisance behavior. Criminals would find ways around this tracking mechanism while average users would experience a near-total loss of privacy. Moreover, such attribution would in no way force noncooperative regimes to cooperate in investigating cyber crimes.

As the title of my written testimony suggests, instead of focusing on attribution, we need to move to accountability in cyberspace. Noncooperation in investigating international cyber attacks should be taken as a sign of culpability. States must be held responsible for securing their national cyberspace and should have an obligation to assist when their citizens or systems within their county are involved in a cyber attack.

Chinese government officials will often protest and lay the blame their country receives in the western press for cyber espionage against both government and corporate attacks by suggesting that the systems the attacks are traced to are simply compromised proxies that have been used to mask the identity of the real attackers. They will also suggest that systems in their country are used just disproportionately in these attacks because of the poor state of cybersecurity due to the widespread use of pirated software and low installation rates for even the most basic software security. This scenario may very well be plausible but even if true, I would argue that it is no longer an acceptable excuse. We need to move to a situation in which countries not only assist in investigating but also have mechanisms in place to shut down systems that are controlling attacks or participating in botnets. Failure to assist should be treated as complicity.

Let me conclude with a comment on the issue of deterrence. Much ink has been spilled trying to make the Cold War construct of deterrence applicable in cyberspace but I believe the results of these efforts are unpersuasive. Deterrence during the Cold War was predicated on mutual assured destruction. While better attribution can let us know who is attacking us, most potential adversaries do not have as heavy reliance on network technologies in their industries, government or militaries. Thus, in order to retaliate in any significant way, we would be forced to escalate out of the cyber domain and conduct kinetic attacks. That is not a situation we want to be in, and the threat to do so may be perceived as incredible, this limiting its deterrent factor. Instead, we need to focus on improving our defenses and making investments to secure our portion of cyberspace.

Thank you very much.

[The prepared statement of Mr. Knake follows:]

PREPARED STATEMENT OF ROBERT K. KNAKE

Untangling Attribution: Moving to Accountability in Cyberspace

Chairman Wu, Ranking Member Smith, and distinguished members of the House Subcommittee on Technology and Innovation, thank you for the opportunity to discuss the role of attack attribution in preventing cyber attacks and how attribution technologies can affect the anonymity and the privacy of Internet users. In your letter of invitation, you asked me to address the following series of questions:

1. As has been stated by many experts, deterrence is a productive way to prevent physical attacks. How can attack attribution play a role in deterring cyber attacks?
2. What are the proper roles of both the government and private industry in developing and improving attack attribution capabilities? What R&D is needed to address capability gaps in attack attribution and who should be responsible for completing that R&D?
3. What are the distinguishing factors between anonymity and privacy? How should we account for both in the development and use of attribution technologies?
4. Is there a need for standards in the development and implementation of attack attribution technologies? Is there a specific need for privacy standards and if so, what should be the government's role in the development of these standards?

Attributions Role in Deterring Cyber Attacks

Let me begin by stating my view that the utility of deterrence in cyber security may be limited and that the problem of attribution has been over-stated for the high end threats that represent a challenge to our national security. In its classic usage, deterrence is the idea of using fear of reprisal in order to dissuade an adversary from launching an attack. For deterrence to work, it is critically important that we know who has carried out the attack and thus attribution is a central component of deterrence strategy. I believe it may be too broad to view deterrence as a productive way to prevent all kinetic attacks. Deterrence was the central concept in preventing a nuclear exchange between the United States and the Soviet Union during the Cold War. It is not, however, a central part of U.S. strategy to prevent terrorist attacks and its importance in preventing conventional military attacks is more limited than in the nuclear case. During the Cold War, deterrence of the use of nuclear weapons was created through the establishment of "Mutually Assured Destruction" or MAD, in which both the United States and the Soviets understood that any use of nuclear weapons would be responded to in kind. The threat of total annihilation kept both sides at bay. Radar and other warning systems provided the mechanism for attributing any nuclear attack and possession of a second strike capability that could provide a nuclear response even after a successful Soviet launch kept the threat of retaliation credible. Equally important, however, was symmetry.

The Soviets as rational actors did not want to see the loss of their cities, industry, and regime in a retaliatory nuclear strike. As long as we had the ability to hold these assets under threat, a Soviet strike against us would not be to their advantage. Such parity does not exist in cyberspace. Attribution may be a secondary problem to the lack of symmetry. Many countries that possess sophisticated offensive capabilities do not have extensive societal reliance on the Internet or networked systems. If attribution could be achieved, deterrence might not follow because a state conducting an attack in cyberspace, may have little to lose through retaliation. The logical solution to this problem is to threaten retaliation through diplomatic or kinetic means outside of cyberspace, responses that could range from the imposition of sanctions to airstrikes. Thus far, despite the onslaught of attacks in cyberspace, no country has chosen to escalate their response outside of cyberspace. Moreover, it may be difficult to achieve proportionality in response to a cyber attack through other means. Deterrence may simply not be a useful concept to address our current state of cyber insecurity.

If deterrence is to be a central part of our cyber security strategy, I believe it is essential that we can answer three questions: First, what degree of certainty in attribution is necessary to take action? Second, what would that action look like? Third, how will we make potential adversaries understand the answers to these questions prior to an incident so that they will be deterred? To begin, I think it is important to breakdown the attribution problem in cyberspace. There are three broad categories of attack that have their own distinct attribution problem. The first attribution problem, the one on which most attention is focused is the attribution problem for attacks carried over the Internet. These attacks are difficult to deter because of the underlying architecture of the Internet, the lack of security on many hosts, and because the individuals or teams carrying out these attacks can do so remotely, from the safe confines of a non-cooperative country. The second attribution problem is for cyber attacks that are not carried over the Internet. Potentially, many of the most dangerous forms of cyber attacks will be carried out against systems that are not connected to the internet through other delivery mechanisms including attacks using microwave or other radio transmissions, thumb drives, and other portable media like CDs and DVDs. For these attacks against well-defended military and industrial systems, the attribution problem is similar to the attribution problem for kinetic attacks and can be addressed through real world forensics, investigation, and intelligence. Finally, there is the problem of attribution for the introduction of malicious code in the supply chain for hardware and software. The threat to the supply chain may be the area of most concern today, yet the attribution problem for the insertion of malicious content into software and hardware is no different from a traditional investigative challenge to identify the opportunity and the motive for inserting malicious content (see Figure 1 for a visual representation of these challenges).

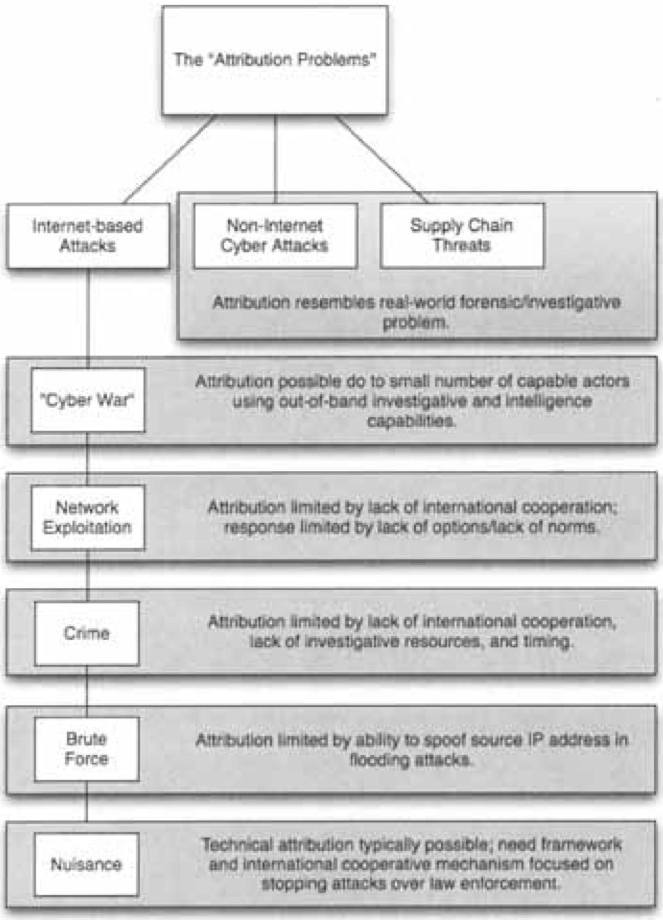
The "Attribution Problems"
Internet-based Attacks
Non-Internet Cyber Attacks
Supply Chain Threats
Attribution resembles real-world forensic/investigative problem.
"Cyber War"
Attribution possible do to small number of capable actors using out-of-band investigative and intelligence capabilities.
Network Exploitation
Attribution limited by lack of international cooperation; response limited by lack of options/lack of norms.
Crime
Attribution limited by lack of international cooperation, lack of investigative resources, and timing.
Brute Force
Attribution limited by ability to spoof source IP address in flooding attacks.
Nuisance
Technical attribution typically possible; need framework and international cooperative mechanism focused on stopping attacks over law enforcement.

Figure 1: The Attribution Problems

With the exception of flooding attacks, all other forms of Internet-based cyber attack require two way communication between the attacking computer and the victim computer. Sophisticated adversaries will take steps to obfuscate their true location and identity through the use of proxy systems, whether they are compromised computers or anonymization services or both. Despite these precautions, trace back techniques and digital forensics can provide the technical means to allow the attackers to be discovered. The barriers to the use of these techniques are more legal than technical, due to international boundaries and non-cooperative countries. If we breakdown the various threats carried over the Internet, the scope of the attribution problem can be brought into focus and different solutions for managing each threat begin to emerge.

Attacks can be divided into the following categories ordered by the threat they pose: cyber warfare, cyber espionage, brute force attacks, crime, and nuisance. For each of these, both the attribution problem and the issue of response are different. For the highest level threat, that of cyber warfare, the attribution problem is largely overstated. As with other Internet based attacks, technical attribution may be difficult and the forensics work will take time, but at present there are a limited number of actors that are capable of carrying out such attacks. Moreover, the resources, planning, and timeline for such attacks would provide many opportunities to identify and disrupt such attacks. Estimates vary, but on the low end, many experts believe that only four countries possess the capability to carry out a catastrophic attack in cyberspace, the so-called Cyber Pearl Harbor, Cyber 9/11, or Cyber Katrina. On the high end, up to 100 state actors and private groups closely affiliated with state actors may have the capability. No matter which estimate is accurate, this is a fairly small list of suspects that can be narrowed down through technical means, as well as out of band methods that include intelligence, analysis of capabilities and analysis of intent. If not already a priority, U.S. intelligence agencies should be focused on identifying actors with high-level capabilities and understanding their intentions. While it has become a truism that hacking tools can be downloaded off the Internet and used by an individual with little or no technical skills, these tools do not pose the kind of threat that could cause widespread destruction. If the operators of critical systems cannot defend against such attacks, they are not taking the threat seriously. As the relevant technologies continue to evolve, it is important that the difficulty in carrying out significant attacks increases. Our critical industries, military and government agencies must continue to raise their defense levels in order to keep the ability to cause destruction in the hands of a limited number of state actors.

In the event of a catastrophic cyber attack, attribution to at least some level will almost always be possible. The question becomes to what level of certainty must attribution be demonstrated in order for the President to take action? At the lowest level, attribution that traces an attack back one hop can provide the foundation for further investigations. If that first hop is in a non-cooperative country that is unwilling to assist in the investigation, that may be enough evidence to hold that country accountable. As with the 9/11 attacks when the Taliban refused to turn over Osama Bin Laden, it may be appropriate under such circumstances to hold a non-cooperative country accountable, a concept I will return to later in this testimony.

On the issue of espionage, the capability necessary for network exploitation is generally lower than that required for destructive attacks, particularly in the realm of economic espionage where private sector companies are targeted. What we lack is not so much an ability to attribute attacks, but international norms that keep espionage limited. Espionage is generally recognized to be permissible under certain circumstances and many scholars will argue that it has a stabilizing effect on the international system by reducing paranoia. As has been recently demonstrated by the discovery of a Russian spy ring in the United States, engaging in espionage is not necessarily considered a hostile act and can be resolved without further escalation. The challenge with cyber espionage is that we lack norms that limit the extent to which states engage in it. This problem is exacerbated by the fact that cyber espionage is not constrained by the costs, consequences and limitations of traditional espionage.

By way of example, consider the case of Robert Hanssen, a former FBI agent who spied for the Soviets and then the Russian Federation for over two decades. Over that period, Hanssen smuggled several hundred pages of classified material to the Russians, who paid him several hundred thousand dollars and maintained a network of handlers in order run this operation. Hanssen paid a heavy price for his betrayal. Having been sentenced to life in prison, he spends 23 hours a day in solitary confinement at a Supermax Facility and is addressed by the guards only in the third person ("the prisoner will exit the cell.") The American spies he betrayed inside Russia were not so lucky. Most were executed. During the Cold War, spying had consequences. Now, according to public media reports, foreign intelligence agencies have exfiltrated several terabytes of information from U.S. government systems.

Whatever country or countries are behind this espionage campaign, the people who are carrying it out are working safely from within the borders of their own country at little risk of being discovered or imprisoned. The low cost and low risk of cyber espionage is the problem, not the difficulty in attributing the source of the activity. If ironclad proof emerged of who was behind an incident of cyber espionage, what would the U.S. response be, particularly given the likely intelligence advantages that the United States gains from cyber espionage? It may be time that we recognize cyber espionage to be a different phenomenon from traditional espionage, one that requires a different set of norms and responses. I doubt, however, that we lack sufficient certainty of who is behind these campaigns that we are limited in our response simply because we do not know who is carrying them out.

Brute force attacks, so called distributed denial of service attacks or DDOS attacks, do present a specific technical attribution challenge. During these attacks, compromised systems formed into a botnet flood targets with large numbers of packets that do not require the targeted system to respond. The malware behind these attacks will provide false information on the source of the packets, so that the machines sending the packets cannot be identified. This particular problem is due to the trusting nature of the internet protocol which does not provide any security mechanism to keep this information from being falsified. To deter DDOS attacks, it may be necessary to strengthen the Internet Protocol so that attacks can be traced to the computers that are part of the attacking botnet, and from their to the command and control servers and potentially to the botnet master himself. It may be equally productive to simply locate compromised computers participating in the attack and shut these down.

For crime, the goal of attribution is to aid in investigation and result in criminal prosecution. Attribution is therefore necessary in the first instance to direct where an investigation should be targeted and for this first step, attribution needs to rise to the level sufficient for 'probable cause' to initiate the investigation. This first level of attribution may only need to lead to a system, not to an individual and an IP address is often times all that is sufficient. In turn, the investigation will need to establish attribution to an individual or group of individuals for the purpose of prosecution. For prosecution to be successful, attribution will need to rise to the level of guilt beyond a reasonable doubt. In between, there is the potential to pursue criminals through civil litigation, in which case the standard for attribution would be lower, and guilt would be assigned based upon a preponderance of the evidence. The problem is that currently, many countries lack both the legal framework and resources to pursue cybercrimes committed by their citizens or that use systems within their territory that target victims in another country. Even crimes committed by individuals in the United States against individuals in the United States will make use of intermediary systems in other countries, particularly those that are not likely or able to cooperate with an investigation. What is needed to deal with the problem of crime is not better attribution but stronger legal mechanisms for working across international borders, the ability to shutdown attacks as they are taking place, and more investigative resources. Ultimately, there must be penalties for states that do not cooperate in investigations and do not take steps to secure their portion of cyberspace.

For nuisance attacks, attribution is rarely a problem. The problem is that few if any investigative resources are assigned to cyber criminal activity that does not have a high monetary value associated with it. This is a situation in which the impact of the crimes committed is fairly low but the resources necessary to address them are high given the volume of the problem. As an example, look at the problem of SPAM. The 2003 CAN–SPAM Act requires spammers to provide accurate header information and to provide an opt-out method for recipients so they can choose not to receive future methods. Yet nearly a decade later, SPAM is flourishing as 9 out of 10 emails are SPAM. For most of these messages, the organization that sent the message is identifiable because they are selling a product. What we lack is an enforcement method that fits this problem, one that is focused on stopping the nuisance behavior rather than prosecuting those who are behind it. Similarly, nuisance level network attacks, the type that can be initiated through downloads off the Internet, are rarely investigated and prosecuted yet they distract system administrators and computer response teams from higher level threats. Investigating and prosecuting more of this behavior could deter many of the people who engage in it.

For most of these threats, the challenges are not so much related to attribution as they are to resources and international cooperation. Focusing on deterrence may simply be the wrong way to think about how to handle these problems. The threats are materializing every day, making the abstract theorizing that laid the foundation for deterrence in a nuclear confrontation unnecessary. They are also, in every respect, a lower level concern that in no way threatens the existence of the United States. Instead we should focus in two areas. We need to reduce the scale of the problem by stopping threats as they unfold and by reducing the vulnerabilities that the threat actors make use of in their attacks. An investigative and enforcement approach to all problems is simply not tenable. Instead of trying to trace every incident back to a human user, we need to develop a legal framework for stopping attacking systems. We must move beyond treating intermediary systems as victims, and start viewing them as accomplices. In the United States, such a framework could require ISPs to monitor their network for compromised systems that have become parts of botnets and quarantine those systems until the problem is resolved. Similarly, we need mechanisms that allow companies or individuals that are under attack and have traced the attack to a system or systems to request for those systems to be shutdown. This process needs to take place quickly and mechanisms must be developed to authenticate such requests across international borders. Such a framework, if developed in the United States, could be promoted as a global model.

For higher end threats, there are lessons we can learn from the last decade of dealing with terrorist threats. The key is to move beyond the search for perfect attribution and instead hold states that do not cooperate accountable. Currently, the situation can be summed up like this. When an attack is traced to another country that is not cooperative, the investigation dead ends. If that country is Russia, Russian authorities will typically say that the incident was carried out either by patriotic hackers or cyber criminal groups that the Russian government cannot control. If that country is China, Chinese officials will point out that China is often the victim of cybercrime and that do to the poor security on many Chinese systems, they are often compromised in an effort to cast blame on China. In both cases, national sovereignty will be raised to explain why cooperation cannot be more forthcoming.

To move beyond this stalemate, the United States should make public a position that treats failure to cooperate in investigating a cyber attack as culpability for the attack. Countries should know that they can choose to have the incident treated as a law enforcement matter by cooperating in the investigation or choose not to cooperate and have the incident treated as a hostile attack for which their country will be held accountable. Over the last decade the concept of state sovereignty has evolved so that sovereignty not only comes with rights in the international system but also responsibilities. The evolution of this concept is due to events in one of the least wired parts of the world: the Hindu Kush.

In 1999, Michael Sheehan, the U.S. Ambassador at Large for Counterterrorism delivered a demarche over the phone to the Taliban's foreign secretary. The message was clear: as long as the Taliban continued to harbor and support al Qaeda and its leaders, the United States would hold the Taliban responsible for any al Qaeda attacks against the United States or other countries. To drive home the point, Sheehan used an analogy. He told the Taliban's representative: "If you have an arsonist in your basement; and every night he goes out and burns down a neighbor's house, and you know this is going on, then you can't claim you aren't responsible." The United States made good on Ambassador Sheehan's word after 9/11, and as the international community attempts to address failed states that cannot control their borders or police their internal territory, this new concept of sovereign responsibility is taking hold.

Applying this new concept of sovereignty to cyberspace has its merits. As with al Qaeda in Afghanistan, failure of a state to prevent its territory from being used to stage an international cyber attack should not, in and of itself, constitute a violation of state responsibility. Indeed, a world in which states monitor and constrain citizen activities to prevent crimes before they take place would be a very frightening world. What is crucial, however, is how states respond when confronted with the use of systems within their territory for cyber attack. If the Taliban had responded to requests to turn over bin Laden, the invasion of Afghanistan might never have occurred. Based on this new paradigm of sovereignty, states should be expected to pass laws making international cybercrime illegal and enforce them. They should have mechanisms in place to respond to international requests for assistance and they should have some ability to oversee the hygiene of their national networks. Better attribution through post-incident forensic techniques will be a crucial part of this new paradigm, but the development of ironclad attribution, will not necessarily lead to better security in cyberspace.

The Role of Government and Private Industry in Improving Attack Attribution

In order to improve attack attribution, there are many things that can be done with current technology. The most crucial is for both government and private industry to do a better job detecting significant threats, mitigating them quickly, and capturing evidence that can be used by law enforcement for investigative purposes. Forensic techniques are getting better, but there are genuine civil liberties concerns with them getting too good.

The vision of perfect attribution can best be summed up as the idea of giving packets license plates. Under such a system, compromised systems or other proxies could not be used to hide the identity of attackers because each packet would be labeled with a unique identifier, possibly an IPv6 address that has been assigned to an individual after having that individual's identity authenticated in some verifiable way. Access to the network would require authentication, and each packet produced by the user would be traceable back to that user. The privacy implications of such a system would be obvious, turning the Internet into the ultimate tool of state surveillance. The security benefits for pursuing criminals and state actors, however, would be minimal. Without cooperation from all foreign states, criminal activity will simply gravitate to states that do not authenticate identity before issuing identification numbers or choose not to participate in the system at all. Many states benefit tremendously from cybercrime, both directly through the cash it brings into economies, and indirectly through the bolstering of technology development through the theft of intellectual capital. Moreover, for less capable states, cybercrime provides the necessary cover of darkness for espionage to take place. By cracking down on cybercriminal groups, the activities of state actors would stand out starkly. Ultimately, such a system would restrict the freedom and privacy of most users, while doing little to curb criminal elements or state actors who would find ways around the system.

As a baseline, of what we should expect from digital forensics, it may be instructive to look at the role forensics plays in the real world. Many people have become familiar with modern forensics techniques through the popular series CSI and its spinoffs, television shows about real-world crime scene investigators. Each episode begins with a body. The crime scene investigators come in and walk the scene collecting forensic evidence and then take it back to the lab and process it for clues. This activity takes us to the first commercial break in an hour-long drama. The forensics have yielded clues about who the victim was, how he or she was killed, and possible attributes of the killer. Then the detective work begins. The detectives try and establish a motive. They delve into the past of the victim. They ask themselves who would have wanted the victim dead? They ask a lot of questions of a lot of people. On television, this process is packed into an hour. In the real world it can take days to weeks, months and years.

Cyberspace isn't so different from the real world. We have digital forensic tools and trace-back techniques that in the latest incident with Google, allowed the company to conclude that the attacks emanated from China. We can't know more than that without some good old-fashioned investigative work but we can ascertain motive based on what systems were infiltrated and what data was stolen. We can narrow down the list of possible suspects by geography. We can further narrow down the set by capability. Only so many people in the world have the ability to put together the kind of code used in the hack. We also know whoever built the exploits wasn't working alone. That's enough leads to get an investigation going in the real world, and it is also enough in cyberspace.

While the Google case illustrates the attribution "problem", it also illustrates the need for Internet Freedom, something the Chinese government is trying to erode. Our law enforcement community might want ironclad attribution on the Internet to combat cyber crime, but the Chinese government and other authoritarian states want it to combat speech. We may want to know who carried out the hacking of Google but we also want to protect the identity of anonymous posters in online forums about Chinese human rights.

Creating the perfect surveillance state online is within our technical means. In real-world equivalents, we could label each packet with its digital DNA, tying it to a single real-world person, and recordings of everything that goes on so we can play back the tape. But cyberspace isn't so different from the real world, especially since more and more of what we used to do by walking we now do online. If we don't want to live in a surveillance society out here, we also do not want to live in one in cyberspace. The tools for digital forensics are getting better. We don't want them to get too good. What the Google incident really demonstrates, isn't a technical problem; it's a legal and diplomatic one. We lack norms for acceptable behavior by states in conducting espionage online and we lack agreements between states to partner in pursuing cross-border cyber criminal activity. Better surveillance wouldn't solve that problem.

In two narrow areas, government and private sector technology companies should collaborate to improve two of the basic protocols that govern internet transactions. First, government and industry must work together to develop a secure version of the basic internet protocol that authenticates the "from" information contained in packet headers. In distributed denial of service or DDOS attacks that do not require the return of information, the ability to supply false sender information makes it difficult to trace and block such attacks. Similarly, the underlying protocols for sending email allow an individual to spoof the identity of a sender so that someone with malicious intent can send email appearing to be from a bank, a friend, or a work colleague. This weakness is typically exploited in social engineering attacks in order to get the recipient to click on a link that will download malware or send back sensitive information. These problems are well known and well documented. After more than two decades, I believe it is safe to conclude that the informal, consensus-based processes used by the Internet Engineering Task Force to develop and adopt new protocols will not solve these problems. The Federal Government must step in, lay out the challenge, and lead the development and adoption of protocols that solve these problems. An "X-prize" strategy might prove useful in this context.

Privacy and Anonymity in Resolving Attack Attribution

In the early days of the Internet, anonymity was how privacy was obtained when online. As a general trend, anonymity on the web is eroding for most users due to the interactive nature of current web content but new ways of protecting privacy have not developed, at least not for the average user. In terms of protecting privacy, anonymity is only useful in a “web 1.0” context. In the web 1.0 era, users were passive recipients of information posted to the web. Anonymity on the web is still useful for accessing information that you do not want others to know you have accessed, whether it be pornographic material or information on democracy if you live under an authoritarian regime. Increasingly, however, access to information is not what the Internet is being used for. Managing health records and finances and communicating online cannot be done anonymously. What is needed is privacy, something that does not currently exist on the web that must be created through both technical and legal mechanisms.

Most of the so-called “free” web is funded through advertising, and advertising is increasingly targeted to individuals based on information collected about them from their IP address and from various types of cookies placed on their computers when they access sites. By the time my homepage at the nytimes.com has loaded, a total of 12 cookies have been loaded onto my computer, including “flash cookies” that cannot be deleted through standard browser settings. While some of these cookies are used to authenticate my username and password on the site, the vast majority are for advertising, meant to track my use of the internet in order to target advertising at me. Companies sell geo-location services that use IP information to determine where you live so that advertising can be targeted at you for local services. By default, my browser, my computer, and the websites I visit are set to allow all this to happen without me knowing it. Advanced users may have the skill set and the motivation to set their browser settings and take other steps to avoid privacy loss but most users do not.

At present, only the technically sophisticated, be they law-abiding citizens concerned with their civil liberties or criminal actors, can obtain anonymity, while the average Internet user experiences a total loss of privacy. As the technology develops to improve attribution, we need to ensure that our laws develop to protect their use, both by government and by the private sector. These points to the need for government intervention to require companies that collect information online and track users to be explicit about what they are doing. Surrendering your privacy online in exchange for “free” access to information should not be something that happens behind the scenes, but an explicit decision that users make. The equivalent of the Surgeon General’s warning, something short, explicit, prominent and standard should be displayed on sites that use privacy compromising methods to generate advertising revenue.

In order to protect private communication online, we need to implement both technical solutions and stronger legal protections for the content of communication. While law enforcement and intelligence agencies are restricted from accessing private information without due process, private sector entities and criminals have far fewer barriers. The average home users email messages are not secured end-to-end through encryption, and the laws that protect the intercept of these messages are far weaker than those that protect regular mail.

Taken together, these steps would replace the loss of anonymity that was the foundation of privacy on the early web, with privacy for all activities carried out over the Internet, including transactions and two-way communication.

Standards Development for Attack Attribution and Privacy

As stated previously, I believe it is necessary for the U.S. government to work with the Internet engineering community to address known problems in the current suite of protocols. In my view, these problems are both limited and correctable but both funding for development and incentives for adoption post-development are necessary. The goal should not be to create ironclad attribution that would turn the Internet into the ultimate tool of the surveillance state. Rather, the end state should be protocols that prevent the spoofing of IP addresses and email.

On privacy standards, I believe that it is government’s role to protect the privacy of individual users. Government must stop assuming that consumers have all the information they need to make informed decisions about privacy. The goal of government intervention in this area should be to make the decision to surrender privacy in exchange for access to information and services a transparent decision. Websites should be required to notify users if access requires the installation of cookies that will track users for the purpose of targeting advertising. Many if not most users may make the decision to surrender their privacy for access to so-called “free content”. Others may choose a pay option. Still others may seek out content that neither costs privacy or dollars.

These two issues overlap for Internet Service Providers. The activity of ISPs is largely unregulated in the United States. For ISPs, attribution on their networks is not a problem: they can see malicious activity and trace it back to a customer. When evidence of the next jump on a host has been deleted, ISPs are often able to trace the next hop of packets. Standards are necessary for what ISPs should and should not be required to track, for how long they should store such information, and how this information can be shared with law enforcement or private parties.

Finally, we need standards for the operation of anonymity services. Services like Hotspot Shield, Tor, and others provide a valuable service to many Internet users, particularly those living under authoritarian regimes where accessing certain websites may not be possible or may be tracked in order to identify dissidents. Yet these same systems can be used for criminal purposes. Standards are necessary for regulating these services and they must be promoted internationally. These services provide anonymity, which, as previously discussed, is only useful for accessing information sources and anonymous posting activity. These services should therefore restrict their users to web-based activity. They should also make it easy for companies and government agencies to block the outbound IP addresses to prevent users that have gained anonymity from attempting to access secure systems. If you are trying to access your own bank account online, there is no legitimate reason to use an anonymization service. Finally, these services should retain auditable logs for law enforcement purposes. Users should understand that this information will be kept private, and only released if the service has been used for criminal purposes. Ultimately, as with states, anonymization services should be held accountable for their users' behavior if they do not cooperate with law enforcement.

Conclusion

As I have expressed throughout this testimony, it is my view that the problem of attribution has been largely overstated. Ironclad or perfect attribution would not address the problems of cyber warfare, espionage, crime or other threats in cyberspace. Such a capability would, however, be injurious to freedom of expression and access to information for many people around the world. Stronger mechanisms for international law enforcement cooperation are necessary, as is the ability to stop attacks in progress, and improvements to the general hygiene of the Internet ecosystem. More than anything else, we need to develop better and stronger options for responding to threats in cyberspace and introduce consequences for states that do not cooperate in stopping attacks or in investigating them. Finally, we need to move beyond anonymity as the guarantor of privacy on the Internet and instead work to create privacy through both technical means and legal requirements. Thank you for the opportunity to testify on these important issues. I would be happy to answer any questions at this time.

BIOGRAPHY FOR ROBERT K. KNAKE

Robert K. Knake is an international affairs fellow in residence at the Council on Foreign Relations studying cyber war. He is currently working on a Council Special Report on internet governance and security. Prior to his fellowship, he was a principal at Good Harbor Consulting, a security strategy consulting firm with offices in Washington, DC; Boston, MA; and Abu Dhabi, UAE, where he served domestic and foreign clients on cyber security and homeland security projects. Rob joined Good Harbor after earning his MA from Harvard University's Kennedy School of Government. He has written extensively on cyber security, counterterrorism and homeland security issues. He is co-author (with Richard Clarke) of Cyber War: The Next Threat to National Security and What To Do About It (HarperCollins, April 2010).

Chairman WU. Mr. Giorgio.

STATEMENT OF ED GIORGIO, PRESIDENT AND CO-FOUNDER, PONTE TECHNOLOGIES

Mr. GIORGIO. Good morning. My name is Ed Giorgio and I am the President of Ponte Technologies. Let me begin by commending Chairman Wu and Committee Members for looking into this important matter. Having personally spent a career in science and technology and having witnessed numerous R&D innovations that improve the quality of our lives, economic livelihoods, security and privacy, I am confident that this Committee will undertake the proper initiatives to solve long-term and extremely difficult problems such as the one we face with cyber attack attribution.

Post-attack attribution today is not effective and the protocols we have today are insufficient to provide it. The recent attacks on Google are neither new or surprising. What is new is the extensive publicity they generated, but despite all this publicity, and a convincing that they were perpetrated by a state-sponsored actor in China, the rate of such cyber attacks coming from China has not decreased. Current attribution capabilities are clearly no deterrent.

We envision transitioning to a multi-protocol Internet infrastructure where service is offered over DoD network segments and sensitive commercial and financial networks would require transmission using new protocols that have accountability and attribution built into their design. On such networks, attack attribution would meet the requirements for legal evidence without giving away sensitive sources and methods. Other less-sensitive services might be offered over network segments such as Radio Free America, which allow or indeed welcome interaction with anonymous entities. This is another case where the current protocols are lacking. They have little support for anonymity or for real flexibility in how much personal information is revealed in a transaction. Each citizen should have access to a certificate or other token that uniquely identifies the holder along with others that provide less or even no identity information. It should be possible to acquire as many such identity certificates as are needed to support multiple online roles. Some organizations already provide physical analogs in the form of prepaid credit cards or anonymous pay-as-you-go cell phones.

As Americans, we fiercely defend our right to privacy and security and subsequently create a vision where we achieve both simultaneously. But transparency is also important. Indeed, one might argue that the history of human social development and even evolution was driven by transparency of action, but we have witnessed three transformations brought about by technology that are having profound impact on human behavior, from attributable to anonymous, from discoverable to forever hidden, and from understandable to magical. Wherever we lost transparency, whether into governments, corporations or individuals, bad actors eventually emerged and violated our trust and our laws.

The threat comes from all these actors, many of whom are beyond the reach of our American courts, whether it is the Chinese stealing our American innovations to produce less-expensive versions, the Russians engaging in financial crimes, the Israelis stealing our political intentions, the French dealing our competition sensitive materials, the Nigerians conning our elderly and so on. Closer to home, we face the same threats from within our borders. In the past, gross violations of domestic civil liberties were justified by reference to foreign threat. These are very dangerous constitutional grounds we tread and the gravity of the legal and constitutional dimensions cannot be trivialized.

So in conclusion, my comments are not focused on promoting what the ideal balance between privacy and security should be but rather a challenge to those embracing the utopian view that both may be simultaneously within our grasp. While we continue to insist that private information remains just that and that anonymous persona will be supported, the existence of a trusted third party may be the only way to ensure that. In my opinion, government has not yet earned the necessary trust to perform this role and we will require a lot more transparency and oversight before giving that trust.

Thank you very much, and I would be happy to answer any questions.

[The prepared statement of Mr. Giorgio follows:]

PREPARED STATEMENT OF EDWARD J. GIORGIO

1. Answers to Committee Questions

1.1 Is Attack Attribution a Deterrent?

Question 1: As has been stated by many experts, deterrence is a productive way to prevent physical attacks. How can attack attribution play a role in deterring cyber attacks?

Attack attribution is much easier in physical space, but also possible in cyber space. One of our goals is to discover who is attacking us, not whose computer systems they are using to launch their attack, or where geographically those systems are located. However, even this is not enough for a diplomatic or public opinion deterrent. Consider for instance the recent attacks on Google. There is little doubt that these were perpetrated by a state-sponsored actor in China, but has the attendant publicity done anything to reduce the number of cyber attacks coming from China?

Attack attribution is an essential part of our overall situational awareness and emergency response measures. For example, we can use attribution to shut down or otherwise protect ourselves from attacks in progress. We can even stop a DDoS attack without attribution as to the initiator of the attack. We just need to stop where it is coming from. However if attribution is to have any value as a deterrent then it needs to be both irrefutable and able to be revealed to the world without compromising privileged information or intelligence assets. In some cases you can show China was a transit point for an attack and didn't stop it; this has value too.

Current technologies allow us some level of attribution, most of which is plausibly deniable. Attribution can sometimes be made irrefutable by combining what is publicly known with the resources available to an intelligence agency such as NSA or the FBI, but this is rarely releasable beyond government circles—much less to the attacker—and thus has little if any value as a deterrent. There is also the option of turning it into a U.S. State Department demarche to the offending country, but even this has pitfalls (like revealing very sensitive sources and methods).

As with any other form of attack, there are numerous types of organizations or individual involved, and some of these may well be deterred from pursuing a cyber attack for fear of attribution and the legal or economic consequences thereof.

Entities whose systems are used as the launching point for somebody else's attack may also be motivated by attack attribution to secure their systems and either stop an attack in progress or prevent such abuse in the future. It is often possible to identify the reputable private institution who owns the offending computer—if this is made public, it can have an adverse impact on the brand of that institution, revealing ineffective controls and poor information security practices. Corporate executives could be held personally responsible for such failures and personally liable if there is damage to shareholder value.

The same could be true of the ISPs whose networks are used to propagate cyber attacks. Where strong competition is present in the market, attribution can play a valuable role in motivating ISPs to address user education, network monitoring, and endpoint security.

With attacks from nation states, or state-sponsored actors, the potential impact of attribution technologies really depends on the nation, and so our response needs to be carefully tailored to that nation to have maximum effect. Some nations will act cautiously, fearful of the consequences that could come from being exposed as a cyber attacker, such as economic damage, sanctions or even war. Other countries do not seem to care. For those nations that do care but also have a strong offensive cyber presence, masquerading as an organized crime entity, or as a country that is well known to be the source of cyber attacks, is an easy way to reduce such risks.

Terrorist groups will not be deterred by attack attribution—they may even welcome it. However, if attribution can be used as a means of geo-locating members of a terrorist group during an attack, this is something that can be used to disrupt their operational tempo.

For organized crime, attribution may serve as a deterrent if that attribution could be used to help build a criminal case against them that will stand up in court. Unfortunately, their chosen targets may not have the situational awareness to know that they are being attacked, or the resources to provide that deterrent. Organized crime groups will often target either bank customers or small companies with vulnerable credit card databases. When they target the government, they will often target individuals rather than organizations—for example to discredit police officers by planting incriminating evidence on their home computers, or to bribe or blackmail insiders to monitor or affect the course of criminal investigations.

When forensic analysis or other collateral information also permits us to identify the actual human offender, criminal charges, prosecution, and conviction will serve as strong deterrents. This will be somewhat expensive to do here in the U.S., very complicated with even close allies, and nearly impossible with the bad foreign actors mentioned above. Consider for example the case of Gary McKinnon, who after eight years is still awaiting extradition from the UK—a very close ally. The legal costs arising from the investigation and long extradition process, along with any future trial, could easily exceed the actual damage of which he is accused. Once a suspect is convicted, their subsequent imprisonment is also expensive. Is this actually a good use of taxpayers' money? We simply do not have the resources to pursue every hacker out there, or even a significant subset of them, much less extradite them to the U.S. and imprison them here.

The last significant group of attackers is the "script kiddies"—typically the easiest attackers to identify, as well as the easiest to protect against. While we should take measures to protect our systems against such attackers, and take measures to identify and deter them where possible, we should keep in mind that many of them really are children. Notwithstanding the damage they cause, our goal should be to guide them towards a more enlightened path in which they become useful and productive members of society, rather than criminalizing them at an early age, which could leave them with no job, no vote, and no stake in the common good.

1.2 Roles of Government & Industry in Technology Development

Question 2: What are the proper roles of both the government and private industry in developing and improving attack attribution capabilities? What R&D is needed to address capability gaps in attack attribution and who should be responsible for completing that R&D?

While company-to-company and nation-to-nation political dialog may well do with less stringent, but plausible, attribution, if attribution is to be used in court then it must be irrefutable and presentable as evidence in its own right. To achieve this, we will have to move to new protocols in the infrastructure which change the very foundation of our networks, building in attribution and accountability from the ground level. Governments and private enterprises are facing similar threats, and trying to solve much the same problems, and so partnerships with industry will help to develop the protocols of the future.

Having built the necessary protocols in collaboration with industry, we can begin to require that entities with a legitimate presence in DoD networks, or in some civil government or critical national infrastructure networks, implement the new protocols as a pre-condition to network access. Some corporate enterprises (particularly in the financial space) will be motivated to do the same for their own business reasons. In this way we can add to the security posture of those networks at the same time as we demonstrate the viability of the enhancements.

This is not something that any one government can push through for broad use in the Internet as a whole. Evidence of this is in the recent claims over the "militarization" of the internet which is not embraced by business, academia, and civil libertarians alike, and even debated within government circles. This is somewhat recognizant of the crypto wars fought two decades ago which ultimately resulted in government conceding the issue. The fact that we may have to make concessions on this issue, should not prevent us from pursuing R&D which will be necessary if/when some politically viable path emerges.

In spite of this resistance to militarization, there are strong economic drivers in global electronic commerce that are pushing towards solving security problems in the infrastructure rather than in the application space. Applications can't sit around waiting to do a time critical task while depending on an unreliable infrastructure. The infrastructure will ultimately enforce stronger authentication for users and terminals, stronger integrity, and non-repudiation assurances for the transactions. These properties, once built into the infrastructure, will serve to decrease gaps in attack attribution capabilities. Infrastructure will always move more slowly than applications, and we should not ignore how quickly application changes can deliver either (and sometimes both) improved privacy and improved attack attribution.

Many credible experts claim the goal, even if deemed reasonable, is not technically feasible. That may be the case to a purist, but the fact that we can't find perfect security solutions anywhere has not deterred us from raising the bar very substantially through many hard fought for improvements.

While government cannot by itself mandate changes in underlying infrastructure technologies (Ex. IPv6), DARPA, NSF, and the research elements supported by the Comprehensive National Cyber Initiative all should be working to research and develop new capabilities. These could be researched, designed, implemented, piloted, and ultimately become operational on DoD and Intelligence networks, where attack attribution is far more important. After all, it was the original ARPANET where current internet protocols were developed and incubated before they ultimately flourished on today's internet.

New protocols based on the above research should be introduced through the IETF, as this process is the most likely to encourage commercial acceptance and deployment into worldwide networks. For security standards or algorithms, NIST is the appropriate agency.

Research in attack attribution would leverage many of the capabilities already developed. We have seen frameworks which securely embed the user ID, computer ID, process ID, institutional affiliation, and geo-location directly into the IP address. One way to do this is with cryptography and allows us to bind the above attributes to the IP address in a non-forgeable way. Continuous improvements in this area could also raise the bar significantly.

We envision transitioning to a multi-protocol internet infrastructure where services offered over DoD network segments would require transmission using these protocols, while other government services such as "Radio Free America" might be offered over network segments which allow or indeed welcome interaction with anonymous entities. Some incremental improvements in this arena are already being made, for example with Trusted Network Connect, which can be used to require machine-level attribution before network access is granted. Similarly, financial institutions might have far more stringent attribution requirements than a news media or marketing agency. Social networking sites would be adaptable to the needs of their constituencies which, I might add, will likely reflect generational differences over the need for privacy.

1.3 Distinguishing Factors between Anonymity and Privacy

Question 3: What are the distinguishing factors between anonymity and privacy? How should we account for both in the development and use of attribution technologies?

Privacy protections are usually given to people who are acting under their true identity while anonymity assumes that people are acting under an anonymous persona. Under privacy, public and private institutions have Personally Identifiable Information (PII) which is bound to other information they retain about their customers. This might be something as simple as the address of a customer who buys firearms. They have policies about protecting such information. Control objectives focused on privacy attempt to mitigate loss from:

a. Unauthorized Individual—Information systems are inadequately protected resulting in a release of data to unauthorized parties inside (or outside) the institution.

b. Authorized Individual—An authorized individual within the institution makes a unilateral decision to overstep their authority and release or sell privacy information.

c. Questionable Institutional Practices—Questionable (and generally accepted) institutional practices push the legal envelope too far by broadly interpreting the privacy laws pertaining to their business.

d. Systemic Institutional Corruption—Systemic institutional corruption results in the willful and unlawful release of privacy information.

In all the above cases, the institution has privacy information which it did not provide adequate protections for. This is not the case with anonymity which would have prevented the institution from knowing the identity of or having PII on the individual in the first place. This is quite different from well intentioned anonymizers which attempt to remove all PII information from data records so they can be used for other purposes, such as research, public health, crime statistics, etc. There have been some failures of anonymized data bases which revealed PII information through "data leakage" or "correlation handles".

There is very relevant research on the problem of working with Internet router flow records which were anonymized by having random substitutions applied to their IP address fields. Researchers were able to recover the actual IP addresses from a collection of anonymized records and known IP address segments. Since the purpose of attack attribution is to identify the attacker, the attacking computer, or the geo-location of the computer, this cannot be done successfully without unmasking someone or some computer who was attempting to be anonymous. Of course, this is not the case if the person was acting under a "anonymous persona" in the first place, in which case there is no persona to attribute the attack to.

Where true anonymity is allowed, attribution is neither desirable nor possible. Therefore a risk management decision has to be made as to how much anonymity is allowed and in which contexts. A news organization may consider it more important to allow anonymity to protect journalistic sources, while a DoD organization may see no need for others having anonymity but every need for security. Today's networks give us a mix between anonymity and security, but no fine-grained tools for managing the trade-off between them.

Many of the transactions on the internet are reasonably private but not anonymous. The financial institutions develop protocols which protect the integrity of the financial transactions, and the merchants may make some attempt to protect customer privacy information, but existing protocols don't allow anonymity where it may be called for. For example, I may wish to research AIDS treatments without letting my search agent know that it is me doing this research. I may even want to buy such treatment without revealing my identity to the merchant who is selling it to me, but I may want the supply chain and the public health officials to know what treatments are of interests to this anonymous purchaser. All of this is possible with the right protocols. In the standards section below we will demonstrate the type of research that is needed to develop such protocols.

In order for online commerce to flourish, there is a strong need for trusted entities to issue trustable and non-transferrable identity certificates. In this way people can be assured that when they communicate with the same online identity twice they are actually talking to the same person both times. Governments around the world already issue physical identity certificates, but in the online world governments came late to the game and private organizations such as Verisign have arisen to fill this gap. Any attempt by government to take back control of online identification, or even just to provide services in this space, will be met with resistance.

Leaving aside the issue of who is issuing identity certificates, and how they are secured so as to be non-transferrable, some of these should uniquely identify the holder while others should be able to provide less or even no identity information. It should be possible to acquire as many such identity certificates as are needed, and unless they contain personal information in common between them there should be no way to link one anonymous identity to another. Some organizations already provide physical analogs, in the form of pre-paid credit cards, or pay-as-you-go cell phones, that require little or no personal information to activate.

1.4 Need for Privacy and Attack Attribution Standards

Question 4: Is there a need for standards in the development and implementation of attack attribution technologies? Is there a specific need for privacy standards and if so, what should be the government's role in the development of these standards?

Technologies that are built into the network architecture need to be made in accordance with open standards, as this promotes interoperability and encourages broad adoption. Technologies for attack sensing and mitigation are more difficult to standardize, and standards may actually harm you because they give the attacker something to test their strength against before they come after you.

So, the military will always have to have secret capabilities for attack attribution in addition to the infrastructure standards discussed in the previous answer. These secret capabilities become problematic when the military is asked to apply them to other government agencies, critical infrastructure, ISPs, academia, and international corporations where transparency is vitally important. This is at the heart of the current Einstein debate which is considering the deployment of military intrusion detection capabilities to protect civil agencies. The only solution I see to this problem is a public-private partnership (or standing commission) where technical expert members have government security clearances while not required for other commissioners who, over time, learn to trust in the unclassified explanations given to them by the technical experts.

In the previous answer, we explained the need for standards involving authentication, integrity, confidentiality, non-repudiation, geo-location, institutional affiliation, and more at the infrastructure level which bind all these attributes to the IP address of the end user. We would add an anonymous persona standard as well as new standards to protect privacy. The government should invest in the development of these standards, but let the open standards groups such as IETF, NIST, ISO, WWC, and more run those standards though their respective processes. The government should have representation at the table.

There is a specific need for new and improved privacy standards. We can best illustrate this by introducing a suggested framework for two important areas where privacy is critical: medical records and on-line transactions. This framework should make it clear that existing protocols for on-line transactions focus on the integrity of the financial transaction rather than the privacy of the parties involved. The framework appears in the last section.

2. Full Discussion

2.1 Introduction

If we are to protect the Internet and its users from criminals, hostile nation states, and terrorists we will have to both design the Internet better and then be vigilant about monitoring it. The former will encourage technologies such as strong authentication, while the latter will likely force us to balance Security (attribution) & Privacy (anonymity) when designing new Internet protocols and host technologies. This may appear strange because, at some level, Security and Privacy (S&P) have a similar definition: ***The right to live out one's life without interference from others***. Indeed we can demonstrate many instances of best practices in computer & Internet security which result in enhancing both security and privacy simultaneously. The very existence of these synergistic outcomes, however, permits arguments that can be used to deflect the discussion away from other areas (like attack attribution) where we frequently have to make tradeoffs.

We say frequently above because it depends on the nature of the attack. Is it a National Security threat, or a criminal action and thus in the law enforcement domain? Attribution techniques sufficient to identify a Nation State initiator of an attack for appropriate political/military response need not impact personal privacy. If it is a criminal attack against banks or persons, "following the money" may be more effective in gaining forensic-quality evidence for court action, as opposed to machine identities used merely as clues as to where to start the hunt for physical evidence of crime.

Privacy and anonymity currently play a critical role to many of us here in the U.S. and to freedom fighters, whistle blowers, bloggers, and amateur reporters in both democratic and repressive regimes all over the globe. It's one of the few mediums where you can be relatively anonymous. Unfortunately, the trend line looks ominous for those capabilities and I think these traits will largely disappear in the Internet in 20 years independent of the best intentions of some governments. This prediction is a function of where the Net came from and the fact it's grown so fast and that it had to maintain the original assumptions which drove Internet plumbing (protocol and router development) in the first place and were friendly to anonymity interests. That said, the net is maturing, and as new protocols come online and a new generation of users grow up, the inevitable degradation of privacy is already well underway. In spite of the best efforts of civil libertarians, the current privacy issues are largely business driven. That is, you could still be anonymous if you wanted, but once you jump into the social networking or online commerce pool, it goes away quickly. It is highly likely that the next generation of internet protocols will have the capability to provide much stronger levels of attribution which will, as a byproduct, serve the interests of those seeking attack attribution. So our lack of privacy and anonymity in portions of the future internet may be inherent in the infrastructure, as well as a byproduct of the applications that ride on top of it, as is the case today.

Geo-location is perhaps one of the greatest threats to both privacy and anonymity. The trend towards wireless mobility is embedding location tags deep in the infrastructure which will be imposed by the new protocols that are difficult to circumvent. These protocols may also embed attributes such as personal identity, hardware identity, physical location, and institutional affiliation right in the internet protocol address. This trend will be business driven as national and international commerce will benefit from the stronger integrity and non-repudiation assurances for the transactions. Strong authentication of the person at the other end will be available from the infrastructure rather than from some application operating over it.

These capabilities will serve us well in emergencies caused by natural disasters, man-made accidents, or hostile foreign threats; tweeters, bloggers, and social media players will get their news and pictures from someone at ground zero, rather than having to first sort through the political rhetoric emanating from a distant corner of the globe. These capabilities will have many other benefits, such as providing parents with the real time location of their children. They will also be used for nefariously purposes by criminals, rogue nations, industrial competitors, and terrorists.

Wouldn't the terrorists like to turn the tables and know when key U.S. public officials or military commanders are dining in a restaurant?

When balancing the need for anonymity with attack attribution, there is no silver bullet, be it technology, policy, economic incentives, or cultural change, which will solve the problem. Even in cases where attack attribution is deemed more important, we don't currently have reliable ways of actually doing it. Furthermore, when we can identify the offending computer with high probability we may not know who the actual human offender is. This is true because the computer owned by the innocent user may have been previously commandeered by a malicious and anonymous adversary operating from a remote location anywhere in the world. For this reason corrective action such as quarantining the offender may actually be depriving the real computer owner of vital and even life supporting services delivered over the internet.

For the reasons stated earlier, it seems reasonable that individuals should have the right to have an "anonymous persona"—or as many of them as they need—which they can use for online interactions. One ought to be able to anonymously check out the prices in Amazon and Borders before making a purchase; one ought to be able to visit the VA STD site before registering for treatment information; one ought to be able to anonymously read about LAPD civil rights violations; one ought to be able to communicate privately and anonymously with others, while still having some assurance that when we talk to the same anonymous ID we are talking to the same person. Many information providers may chose to only release information to properly authenticated and authorized individuals, but what about sites giving guidance to political dissidents, whistle blowers, oppressed groups, freedom fighters, etc.? These sites, of course, want to share this information privately and without any strings.

In a world of insecure computers and botnets (commandeered armies of innocent computers) we will need attack attribution to point us to the offending computer, its owner or institutional affiliation, and its geographic location. But as computers become virtualized we will lose the ability to attribute action to specific computers and as we move to cloud computing we will even lose the ability to geo-locate the computer. This doesn't mean that we can't encode the user identity, computer ID, process ID, and institutional affiliation into the computer's (IP) address, because with the proper R&D we can move to a next generation of internet protocols which do precisely that.

2.2 Anonymity

As children, many of us watched a program called "The Invisible Man". Let's suppose that technology makes that a reality where one could take a pill and become invisible for the next hour. This technology might profitably be used to observe nature without disturbing it, visit public places without the fear of recognition and unwanted attention, associate with people we don't want to be linked to, etc. This technology is needed just as much by government entities as it is by citizens. Of course, it is also easy to envision how this technology might be used to commit crime, so we could surely expect a response which would, for example, make it illegal to enter a government building in the invisible state. Banks would respond by refusing ATM withdrawals to invisible people. While all of this sounds like an absurd policy debate, it is precisely what is being played out in cyber space today. Invisible actors from all of the threat groups are ever present in our computers, behind our locked doors, not in the jurisdiction of our courts, not in range of our guns, and overhearing both out thoughts and our private conversations.

2.3 Losing Transparency

As Americans we fiercely defend our right to privacy and security, and subsequently create a vision where we achieve both simultaneously. This vision embodies our protection from individuals, corporations, governments, cultural and religious institutions, subversive organizations, and common criminals. Through our human experience with these actors we recognize that we have reason to fear all of them. Our lives are played out in part through acts conducted by "perpetrators" and which have impact on "victims". While these words are pejorative, it is this concept of becoming a victim that drives our passion for achieving privacy and security. The problem with this logic is that the laws and tools which give potential victims privacy and security can also be used by the threat agents to achieve anonymity. The result is a world with very little transparency into what everybody, from criminals to nation states, are actually doing. Even when we can see the consequence of these actions we may never know who the perpetrators are. One might argue that the history of human social development (and even evolution) was driven by transparency of action. While human nature has remained largely unchanged, we have witnessed three transformations brought about by technology that are having a profound impact on human behavior:

- Attributable to anonymous
- Discoverable to forever hidden,
- Understandable to magical

Wherever we lost transparency, whether into governments, corporations, or individuals, bad actors eventually emerged and violated our trust and laws.

2.4 Who Should We Fear

In America we have a somewhat unique tendency to fear violation of our privacy from government above all. This stems from our beliefs and experiences that if we are wronged by an individual or a corporation we have recourse from damages in a court, while government has historically avoided such accountability. But, let us first explore the expanded threat to privacy and be specific about some of the (largely) foreign threats. Are we not concerned about the Chinese stealing our technology to produce less expensive versions, the Russians engaging in financial crimes, the Israelis' stealing our political intentions, the French stealing our competition-sensitive materials, the Nigerians conning our elderly, and so on? These actors are all foreign threats, and they represent official governments, large corporations, terrorists, and common criminals. And yet, to most of us, these actors are all beyond the reach of our American courts. Our security and privacy is threatened by all of them, yet many folks continue to focus primarily on government. I would suggest that more balance is needed in first identifying the real threat and then establishing the appropriate balance between privacy and security.

Finally, I would be remiss to exclude the fact that while many of these threats are foreign, many are domestic, and, in the past, violations of domestic civil liberties were justified by reference to foreign threat. These are very dangerous constitutional grounds we tread and the gravity of the legal and constitutional dimensions cannot be trivialized.

2.5 Conclusions

In conclusion my comments are not focused on promoting what the ideal balance between privacy and security should be, but rather a challenge to those embracing the utopian view that both may be simultaneously within our grasp. We need to put together representatives from both sides of the debate, allow them to frame the issue, and present the differences in a way our policy and law can respond appropriately. While we will continue to insist that private information remain just that, and that anonymous persona will be supported, the existence of a trusted third party such may be the only way to ensure that. So, the debate might eventually come to: can we trust government with the information it needs to protect our security or do we lose our privacy from a myriad of bad actors (the least of which may be government)? In my opinion government has not yet earned this trust and we will require a lot more transparency and oversight before giving that trust.

In summary, the privacy & security debate (and hence the anonymity and attribution debate) focuses us on only one aspect (albeit very important) of the problem and we need several initiatives to correct that. In parallel, we should also be using our status as a superpower to drive behavior by the Chinese on the internet, the French on business-competition practices, the Russians on stamping out financial crime, the Israelis on influencing our political system, and international crime-fighting organizations on establishing deterrents. This will require a U.S. policy with an enlightened international agenda which focuses on using what remaining superpower status we have to drive behavior. This is essential to balancing security and privacy at home while simultaneously promoting a robust ecommerce and human rights agenda globally. Once such behavior is agreed upon our policy must be "trust but verify" and will require some authorized (and transparent) monitoring of our information and telecommunications systems, while at the same time, embracing really strong mechanisms to protect privacy and anonymity. This monitoring will allow authorized governments to perform attack attribution with cooperation from the private sector. It will also require oversight by a trusted third party and considerable transparency on Main Street.

3. Appendix: New Privacy Standards Framework

We suggest a new framework to evaluate the security of an on-line transaction. We do this only to elaborate on the inadequacies of the current protocols which focus much more on security than privacy. Our transaction involves a buyer (Bob), a search agent (Goliath), a seller (Sam), a trusted identity provider (Ida), a bank (Betsy), manufacturers (Matt and Martha), the blind anonymity provider (Andy), and finally, Bob's roaming service (Robin). Bob wants to purchase specific goods and begins with asking Goliath to provide a list of sellers. Bob then selects a seller Sam and purchases a product using a credit card he was issued by Betsy. Ida provides some real time assurance that Bob and Sam are who they claim to be. Andy facilitates the sharing of some transaction details with manufacturers Matt and Martha who need to restock the shelves. Note that these latter details are not made available to Andy who is "blind" to the information needed by the wholesalers. Robin provides a roaming and/or backup service for Bob's secret credentials (Robin herself is blind to these credentials).

The security complexity of multi-party protocols grows rapidly as the number of parties in the transaction increases. Our problem potentially has eight distinct roles with some of the roles having multiple players within a specific transaction (such as merchants, manufacturers, or identity providers). Different parties talk both directly and indirectly to each other, security assertions are checked and passed along to other parties, and authentication, integrity, authorization, privacy, and non-repudiation are potentially important to each of the relationships.

We are now in a position to form a privacy framework based on the outcome of several assumptions:

1. Bob knows everything about his transactions.
2. Where Bob has shared his personal information with the other parties, he should still (legally) own that information and be able to update or revoke it at a later date.
3. Ida(s) has provided identity assurance to potentially all parties in the transaction.
4. Goliath knows the set of sellers that have the products Sam is interested in, and, may or may not know Bob's identity.
5. Sam has sold a product to Bob, and Sam may know Bob's identity and his bank account number (today's situation), or Sam knows Bob's identity and mailing address only, or Sam doesn't know anything about Bob.
6. Sam may keep a record of the purchase, but the customer data, and the account information may be kept by Bob only, or by both Bob and Sam.
7. Betsy knows that Bob has made a purchase from Sam, has completed the financial transaction, and may or may not know detailed information about the product that was purchased
8. Matt and Martha know somebody's "purchasing interest" or "purchasing profile", and may or may not know their identity.
9. Andy has facilitated the transfer of some encrypted data from Bob to Matt and Martha, but doesn't know what it is.
10. Robin has encrypted information about Bob, including his secret keys, so she can support his roaming, but knows little more than Bob's identity, and certainly can't decrypt his secret keys.

The choices in the above framework do not have one-size-fits-all answers, so the ultimate protocol selected must be tunable to the answers that fit the situation.

For brevity, we will not demonstrate a similar privacy framework for medical purposes, but we will point out that there are even more stakeholders in the communications and data retention aspects of any medical situation, and enumerate those stakeholders. They include patient, attending physician, treatment facility, pharmaceutical provider, nurses and other medical care professionals, consulting physician, insurance provider, public health officials, pharmaceutical and infectious disease research community, accounting and billing support staff, and several others. While there are currently many places where anonymizers are used today to share medical information, we believe those protections are woefully inadequate.

Acknowledgements

I would like to acknowledge the contributions by several people who made critical comments and constructive ideas during the drafting of this testimony. All the views expressed in the preceding text certainly do not represent the positions of the names listed below. Indeed, in some areas, their views represent alternate positions. Never-the-less, their contributions were invaluable.

William Crowell, Consultant, former CEO Cylink, former Deputy Director NSA

Jerry Dickson, former Director of the National Cyber Security Division (NCSD) at DHS

Kevin R. Fall, Ph.D.
Daniel E. Geer, Jr., Sc.D., CISO, In-Q-Tel
Susan Landau, 2010–2011 Radcliffe Fellow, Harvard
Ronald D. Lee, Attorney
James Lewis, Center for Strategic and International Studies
Mike McConnell, Booz Allen Hamilton, former DNI, former Director NSA
Vin McLellan, Consultant and Publicist in Security & Cryptography
Alan Paller, Director of Research, SANS institute
Bruce Potter, CTO of Ponte Technologies, SHMOO founder
Marcus Ranum, CSO of Tenable Network Security
Brian Snow, Cryptographer and former NSA Senior

Finally, this testimony would not have been possible without the content and editing contributions from Patrick Henry of Ponte Technologies.

BIOGRAPHY FOR EDWARD J. GIORGIO

Ed Giorgio is the co-founder and president of Ponte Technologies, a security and technology company. He is on numerous advisory boards, including the NSA Advisory Board and the Commission to advise the 44th president. He was formerly a principal at Booz Allen Hamilton, where he spent ten years working on information security and enterprise resilience issues for a variety of commercial clients and Federal agencies. Mr. Giorgio also has nearly 30 years of security experience with the National Security Agency (NSA). While at NSA, he pioneered developments in communications security, national intelligence policy and technology, and public key cryptography. Mr. Giorgio is the only person to have served as both Chief U.S. codemaker and, subsequently, as Chief U.S. codebreaker at NSA where he directly managed 1600 mathematicians and computer scientists. As a mathematician, he designed and delivered the first public key based e-mail privacy and authentication system on the worldwide intelligence network. Today he provides services which help clients bridge business innovation, technology, and security and delivers these services to government and commercial clients. He also advises investment bankers and VC's on the viability of early-stage security companies. Mr. Giorgio is considered a leading authority on cryptology and has extensive experience in cryptography, Internet security technology, wireless security, security policy, information warfare, privacy, and intelligence sources and methods.

Chairman WU. Thank you very much, Mr. Giorgio.

Mr. Rotenberg, please proceed.

STATEMENT OF MARC ROTENBERG, PRESIDENT, ELECTRONIC PRIVACY INFORMATION CENTER

Mr. ROTENBERG. Thank you very much, Mr. Chairman, Members of the Subcommittee. I appreciate the opportunity to be here today. I am President of the Electronic Privacy Information Center and I teach privacy law at Georgetown and I have been involved in most of the debates about cybersecurity and privacy going back 25 years.

My organization publishes an important report about privacy and human rights around the world, and I draw attention to this because in our testimony, we talk about the use of attribution by governments, not necessarily for the purpose of promoting cybersecurity but actually to monitor and track people with unpopular political opinions. China has the most advanced means of attribution today for Internet users. They require Internet users to individually register themselves, to provide their true names, their e-mail addresses and the list of news services from which they receive information on the Internet. They require Internet service providers to keep detailed logs on the activities of people who get access to the Internet through Chinese licensed ISPs, and they require the cyber cafes, which is the main point of access for people in China who want to get information on the Internet to track all the activity and keep these records for 60 days to make them available to the Chinese government, and most interestingly, because I also have a background in managing one of the Internet domains, the .org domain, when the .cn domain became available for website registration, the Chinese government also required that businesspeople who wanted to create an Internet website using the .cn domain provide their actual name and a photograph to the government so that they could also be identified.

Now, China, of course, is not alone, and I cite in my testimony similar examples involving Burma, Syria, Iran and Egypt. The point that I am trying to make here is that there is a real risk, which I think was suggested by one of the other witnesses, that attribution techniques through this means of keeping track of what people do online will be used for purposes unrelated to cybersecurity that has a real impact on human rights and freedom of expression because of course what attribution also does is make people think twice about saying things that might be unpopular or controversial.

Now, fortunately, in the United States, as I also describe in my testimony, we have a very strong constitutional right to speak anonymously, which is perhaps not surprising because the Federalist Papers that provided the basis for our country were written by people who made frequent use of pseudonyms. They understood that publishing their views in a way that could be easily attributable to them might quell their efforts to change the form of government that existed in the colonies at the time, and our courts have said repeatedly that anonymity is an important right that is protected within the First Amendment. More recently, we have also been involved in cases involving Internet freedom and the famous ACLU [American Civil Liberties Union] versus Reno case from 1996 that struck down the Communications Decency Act where the Supreme Court affirmed the very important role that the First Amendment plays in protecting Internet freedom.

Now, what I did in preparation for this hearing with the help of our excellent law clerks who are at EPIC this summer was to research the cases involving identification requirements for the Internet. We were trying to answer your very specific question, would it be possible in the United States to have an identification requirement, a mandatory requirement for anyone who goes online, which is certainly being talked about, and our conclusion is that we don't think it would be possible. In the one case where an identification requirement has been upheld, and this was in the State of Utah after an earlier effort had been struck down, it was permitted only for convicted sex offenders where there was narrow collection of personal data and used for very narrow purposes. That is the only case that we could find.

Finally, as I also set out in our testimony, looking at this problem of attribution turns out to be very difficult, as other witnesses have pointed out, primarily because it is so easy for people online to evade detection. Bruce Schneider, who is a noted security expert, said bluntly, "It is futile." What it will do is actually create new opportunities for people to hide because they will create new false credentials, and the recent report from the National Research

Council that also looks at the issue of attribution reaches a similar conclusion. This is not to say that we aren't aware that there are serious network threats which obviously implicate privacy and security interests but we think it is very important in this area to also consider the harmful impact that a broad attribution requirement might have for the freedom of Internet users.

Thank you again for the opportunity to be here.

[The prepared statement of Mr. Rotenberg follows:]

PREPARED STATEMENT OF MARC ROTENBERG

Mr. Chairman, Members of the Committee, thank you for the opportunity to appear today to discuss the topic of Cyber Security and Attribution. We appreciate your interest in this topic.[1]

My name is Marc Rotenberg. I am President of the Electronic Privacy Information Center (EPIC), a non-partisan public interest research organization established in 1994 to focus public attention on emerging privacy and civil liberties issues. Since our founding, we have had an ongoing interest in computer security, privacy, and identification. In fact, EPIC began in response to a proposal from the National Security Agency to establish a mandatory key escrow encryption standard that could have easily prevented the emergence of the Internet as a powerful force for economic growth and political change.

EPIC was founded in 1994 in part to address concerns about the role of the National Security Agency in computer security policy.[2] Since then EPIC has participated in numerous public debates regarding the protection of privacy rights on the Internet and elsewhere. EPIC is currently engaged in active litigation under the Freedom of Information Act with the NSA and National Security Council regarding National Security Presidential Directive 54, a secret document that governs the NSA's current authority over cyber security policy.[3] EPIC has also been involved recently in seeking information regarding the secret cyber security program known as EINSTEIN 3.0, as well as a new secret program within the NSA called "Perfect Citizen."[4] And I have participated in scientific workshops on such topics as "eDNA," a proposal to tie every user activity to their unique DNA, developed by Admiral John Poindexter the architect of Total Information Awareness, that was thankfully rejected.[5]

In my statement today, I will point to the risks and limitations of attempting to establish a mandatory Internet ID that may be favored by some as a way to address the risk of cyber attack. Such a proposal has significant implication for human rights and freedom online. It is not even clear that it would be constitutional to mandate such a requirement in the United States.

To be clear, there are real concerns about network security. Network vulnerabilities also have implications for privacy protection. But solutions to one problem invariably create new problems. As we learned in the early days of the Internet, a proposal to make it easier for the government to monitor network traffic will also make communications more vulnerable to criminals and other attackers. Similarly, proposals to mandate online identification will create new risks to privacy and security.

I. Internet attribution requirements have resulted in censorship and international human rights violations.

It may be that governments establish attribution requirements to address cyber security concerns. But it also clear that governments impose these requirements to

[1] EPIC Counsel Jared Kaprove and EPIC IPIOP clerks Matthew Lijoi, Laura Moy, Reuben Rodriguez assisted in the preparation of this statement. The views expressed are my own.

[2] *See* EPIC, *The Clipper Chip*, *http://epic.org/crypto/clipper* (last visited July 13, 2010).

[3] *EPIC* v. *NSA*, No. 10–196 (D.D.C. filed Feb. 4, 2010).

[4] *See generally* EPIC, *Cybersecurity and Privacy*, *http://epic.org/privacy/cybersecurity/* (last visited July 13, 2010).

[5] John Markoff, *Surveillance Agency Weighed, but Discarded, Plan Reconfiguring the Internet*, N.Y. TIMES, Nov. 22, 2002, available at *http://www.nytimes.com/2002/11/22/politics/22TRAC.html*. The project description of eDNA stated:

> We envisage that all network and client resources will maintain traces of user eDNA so that the user can be uniquely identified as having visited a Web site, having started a process or having sent a packet. This way, the resources and those who use them form a virtual 'crime scene' that contains evidence about the identity of the users, much the same way as a real crime scene contains DNA traces of people.

track the activities of citizens and to crack down on controversial political views. We know this from our research of identity requirements for Internet use outside of the United States.[6] The risk of mandatory attribution can be seen most clearly today in China. If fact, in just the last day, the Associated Press reported on efforts in China to crack down on anonymity and mandate identification requirements.[7]

Currently, China leads the world in Internet use. Over 360 million people access the internet in China, an increase of 1,500% since the year 2000, accounting for over twenty percent of the world's online population.[8] Despite these numbers, Chinese Internet users must abide some of the strictest identification requirements to get online. By making user Internet activity appear attributable to the individual, China's regulations generate user self-censorship.

The Chinese government identifies users who access to the Internet in three ways: (1) mandatory registration requirements, (2) requirements on Internet Service Providers, and (3) regulation of Internet cafes.[9]

China first began control over individual access to the Internet in 1996, and has since revised its policies several times;[10] many of these revisions entailed requirements that users provide identification when accessing the Internet or using certain Internet services. Chinese citizens wishing to access the Internet are required to obtain a license for Internet access. They must register with the local police by providing their names, the names of their Internet service providers (ISPs), their email addresses, and any newsgroups to which they subscribe.[11] In February of 2010, the Chinese government lifted a ban on registrations of domain names ending in the ".cn" suffix, but also imposed strict new requirements for their use.[12] Now, individuals individual wishing to set up personal websites using the suffix must verify their identities with regulators and have their photograph taken.[13]

Additionally, some local and provincial Chinese authorities currently require that individuals use their real names when accessing bulletin boards, chat rooms, or IM services.[14] The requirement also extends to university settings,[15] and in July 2005, all administrators and group founders of China's largest instant messaging service, QQ were told that they must use their real names to access the service.[16] A notice from the Shenzhen Public Security Bureau declared: "This year, at various internet chat rooms in our city, there were chat groups, forums, BBS, internet SMS and various internet public information services in which there were illegal assemblies, illegal alliances and obscene behaviors being observed. In order to protect national security and preserve social stability. . .we will be conducting clean-ups on network public information services."[17]

[6] *See generally* EPIC, PRIVACY AND HUMAN RIGHTS: AN INTERNATIONAL SURVEY OF PRIVACY LAWS AND DEVELOPMENTS (2006) [hereinafter "PRIVACY AND HUMAN RIGHTS."]

[7] Anita Chang, *China seeks to reduce Internet users' anonymity*, Associated Press, July 13, 2010, at *http://www.google.com/hostednews/ap/article/ALegM5goT1Hz28jUIOSMcwiJD9mX6GVZyQD9GUI6VO0* ("A leading Chinese Internet regulator has vowed to reduce anonymity in China's portion of cyberspace, calling for requirements that people use their real names when buying a mobile phone or going online, according to a human rights group.") *See also*, Rebecca MacKinnon, *RConversation: China's Internet White Paper: networked authoritarianism in action*, June 15, 2010, *http://rconversation.blogs.com/rconversation/2010/06/chinas-internet-white-paper-networked-authoritarianism.html*.

[8] Internet World Stats, *Internet Users—Top 20 Countries—Internet Use*, *http://www.internetworldstats.com/top20.htm* (last visited July 13, 2010).

[9] See Trina K. Kissel, *License to Blog: Internet Regulation in the People's Republic of China*, 17 IND. INT'L & COMP. L. REV. 229 (2007).

[10] Kristin M. Reed, Comment, *From the Great Firewall of China to the Berlin Firewall: The Cost of Content Regulation on Internet Commerce*, 13 TRANSNAT'L LAW. 451, 462 (2000). *See also*, PRIVACY AND HUMAN RIGHTS 349–51 (2006) ("China—Monitoring of Cybercafes").

[11] *Id*.

[12] Reporters Without Borders, *Internet Enemies: China*, at 3, Dec. 3, 2010, available at *http://en.rsf.org/IMG/article_PDF/china-china-12-03-2010,36677.pdf*.

[13] David Pierson, *China Steps Up Policing of New Websites*, L.A. TIMES, Feb. 25, 2010.

[14] Radio Free Asia, *China Tightens Grip on Cyberspace*, Aug. 17, 2005, *http://www.rfa.org/english/news/in_depthJ2005/08/ 17/internet_china/*.

[15] *Id*.

[16] Nanfang Weekend, *Fourteen Departments United to "Purify" the Internet*, Aug. 18, 2005, translated in EastSouthWestNorth, Purifying the Chinese Internet, *http://www.zonaeuropa.com/20050821_1.htm* (last visited July 9, 2010). QQ has 100 million active users, including 8 million users who are founders or administrators.

[17] *Id*.

Chinese state-licensed ISPs are required to track and store user activity.[18] ISPs must retain records on user identification, what sites the user visited, the duration of the user's visits, and the user's activity on those sites.[19] Though Chinese laws prohibit disclosure of this information generally, they make exceptions for a number of government purposes, including national security or criminal investigations.[20] Moreover, there are few formal procedures for requesting such data, and most of the time ISPs will disclose to the government an individuals internet usage and identification with just an informal request.[21]

Finally, Internet cafes in China abide by strict regulations that require them to identify their patrons.[22] Many Internet users in China rely on Internet cafes as a primary means of access.[23] All Internet cafes must install filtering software, ban minors from entering, monitor the activity of their patrons, and record patrons' identity and complete session logs for up to sixty days.[24] In many cities, Internet cafes are also connected by live video feeds to the local police department.[25]

The identification requirements China placed on Internet access cause users to police their own Internet usage. China's Internet users (justifiably) believe that all of Internet activity is attributable to the individual. Transgressing Chinese Internet policy is often met with harsh penalties.[26] Therefore, without anonymity, many Internet users in China steer well clear of any potentially controversial activity that might violate China's vague Internet prohibitions.

China is well known for directly filtering internet content within its borders;[27] however, the practice of attributing Internet activity to the specific user through identification requirements is even more effective in regulating Internet content than direct filtering.[28] China's identification laws are designed to make the user believe "that every bit of [her] activity is tracked."[29] Furthermore, China's enforcement of its Internet laws gives users reason to be concerned that if they violate the laws, they will be caught and the punishment will be severe.[30] Almost every internet-related imprisonment resulted from an accusation of subversion, a guilty verdict, and a two to twelve year prison sentence.[31] In this way, "[t]he manhunts for individual internet users, which often mobilize dozens of agents from the public security and state security ministries, serve as warnings for the recalcitrants and dissidents who continue to surf the internet."[32]

Given that individual users, content providers, and ISPs can all be held liable for illegal content,[33] each of these entities acts as a self-censor, avoiding, monitoring, or deleting content that might be illegal. Removing Internet anonymity and requiring identification to access the Internet means that China's "best censorship is self-censorship."[34]

[18] *See* Open Net Initiative, Internet Filtering in China (2009), *http://opennet.net/sites/opennet.net/files/ONI_China_2009.pdf* at 15.

[19] *Id*.

[20] *Id*. at 14.

[21] *Id*. at 14–15.

[22] *See id*. at 15. *See also*, Jill R. Newbold, Note, *Aiding the Enemy: Imposing Liability on U.S. Corporations for Selling China Internet Tools to Restrict Human Rights*, 2003 U. ILL. J.L. TECH. & POL'Y 503, 504 (2003).

[23] *See generally*, Audra Ang, *China Wants Web News 'Civilized'*, DESERET MORNING NEWS, Sept. 26, 2005, at A4, *available at* 2005 WLNR 15133888.

[24] Open Net Initiative, *supra* note 18 at 15.

[25] *Id*.

[26] *E.g.*, Kristen Farrell, *The Big Mamas are Watching: China's Censorship of the Internet and the Strain on Freedom of Expression*, 15 MICH. ST. J. INT'L L. 577, 578–85 (2007) (describing three examples of arrests and imprisonment for internet speech).

[27] *See, e.g.*, Open Net Initiative, *supra* note 18.

[28] *See generally*, Congressional-Executive Commission on China, 2005 Annual Report, at III(e), *http://www.cecc.gov/pages/annualRpt/annualRptO5/2005_3e_expression.php* (last visited July 9, 2010).

[29] Tim Johnson, *In China, Sophisticated Filters Keep the Internet Near Sterile*, MCCLATCHY, July 13, 2005, *http://www.mcclatchydc.com/2005/07/13/12100/in-china-sophisticated-filters.html*.

[30] Congressional-Executive Commission on China, 2005 Annual Report, at III(e), *supra* note 28. *See also* Farrell, *supra* note 26; Kissel, *supra* note 9 at 243–46.

[31] *See* Bobson Wong, *The Tug-of-War for Control of China's Internet*, *http://www.hrichina.org/fs/downloadables/pdf/downloadable-resources/a3_Tugofwar.2004.pdf?revision_id=8986* (last visited July 9, 2010) (describing Chinese citizens who were imprisoned for posting information on the internet).

[32] Reporters Without Borders, *Living Dangerously on the Net: Censorship and Surveillance of internet Forums*, May 12, 2003, *http://www.rsf.org/article.php3?id_article=6793*.

[33] *See* Open Net Initiative, *supra* note 18 at 15.

[34] Matthew Forney, *China's Web Watchers*, TIME, Oct. 3, 2005, available at *http://www.time.com/time/magazine/article/0,9171,501051010-1112920,00.html*.

In addition to China, several other countries have used Internet identification requirements to limit or control their citizens' speech. In Burma, internet cafes are required to take screenshots of their patrons' screens every five minutes, and must be able to provide every users ID number, telephone number, and address if the police request them.[35] In Egypt, Internet cafes must be licensed by the government, although what the requirements and stipulations of obtaining a license are unclear.[36] Additionally, although no formal policy demands it, Internet cafe owners are often coerced through licensing raids into recording customer IDs and maintaining them on file. The records are not sent to a central database.[37] In Iran, ISPs are liable for their users' activity, and are also responsible for recording all user information and IP addresses.[38] All Internet traffic is also routed through the Telecommunications Company of Iran, so it can easily be monitored.[39] In Syria, although other ISPs are available, users wishing to use the government-owned Syria Telecommunication Establishment (STE) must apply with their government issued identity card and supply their username and password.[40] Internet cafes are also heavily monitored, with cafe managers required to take customers' personal information (up to and including mother's and father's names) and to keep a record of what sites their customers visit. Additionally, cafe managers must report any overtly illegal activity.[41] Just like in China, all these identification and tracking requirements must lead to self-censorship of politically sensitive speech.

II. In the United States, a government-mandated Internet identification requirement would likely violate the First Amendment.

Anonymity is an important protection to shield the speakers of unpopular or controversial opinions. It is settled law that the First Amendment incorporates a right to speak anonymously.[42] A government mandated identity requirement would pose a significant threat to the ability of users to engage in political speech online. In order to place such a burden on the ability of individuals to express political speech, the government must show that the proposed burden is the least restrictive means of advancing an overriding state interest. Under this standard, a program to deter and investigate cyber attacks in which all users are required to identify themselves before accessing the Internet is unlikely to be constitutional in practice.

A. The First Amendment protects the right to speak anonymously online.

Anonymous and pseudonymous speech has a long history in the United States. Before the American Revolution, much political writing was distributed in the form of anonymous pamphlets and later, during the debate surrounding adoption of the Constitution, the Founders published essays under names such as "Publius," "Cato," and "Brutus."[43] In light of this history, the Supreme Court has recognized a First Amendment right to anonymous political speech.[44] As the Supreme Court said in the McIntyre case, while this right to remain anonymous "may be abused when it shields fraudulent conduct. . .our society accords greater weight to the value of free speech than to the dangers of its misuse."[45] Courts have also recognized that in the area of speech, the interest in anonymity outweighs other competing interests, such as the interests in preventing fraud, false advertising, and libel.[46]

In the current age, the Supreme Courts has recognized the important role the Internet plays as a means of communication.[47] People use the Internet for a wide

[35] Reporters Without Borders, *Internet Enemies—Burma*, at 3, *http://en.rsf.org/internet-enemie-burma,36676.html*.

[36] *See* Eric Goldstein, et al., *False Freedom: Online Censorship in the Middle East and North Africa*, Human Rights Watch Vol. 17, No. 10(E) at 33 (2005) (hereinafter *False Freedom*).

[37] *Id.*

[38] *See False Freedom*, *supra* note 36 at 47.

[39] Open Net Initiative, *Internet Filtering in Iran, 2009*, *http://opennet.net/sites/opennet.net/files/ONI_Iran_2009.pdf* at 3.

[40] *False Freedom*, *supra* note 36 at 75.

[41] Reporters Without Borders, *Internet Enemies—Syria*, at 3, *http://en.rsf.org/IMG/article_PDF/syria-syria-12-03-2010,36689.pdf.*

[42] *McIntyre v. Ohio Elections Comm'n*, 514 U.S. 334 (1994).

[43] *See McIntyre v. Ohio Elections Comm'n*, 514 U.S. 334, 368 (1994)(Thomas, J. concurring).

[44] *Id.* at 342.

[45] *See id.* at 357 (citing *Abrams v. United States*, 250 U.S. 616, 630–31 (Holmes, J., dissenting)).

[46] See, e.g., *Talley v. California*, 362 U.S. 60, 65 (1960).

[47] *See Reno* v. *Am. Civil Liberties Union*, 521 U.S. 844, 870 (1997) (finding that Supreme Court precedent "provide[s] no basis for qualifying the level of First Amendment scrutiny that should be applied to [the Internet]").

range of political and social purposes.[48] Through the use of the Internet, "any person with a phone line can become a town crier with a voice that resonates further than it could from any soapbox."[49] Anonymity is an important part of Internet communication. "The 'ability to speak one's mind' on the Internet 'without the burden of the other party knowing all the facts about one's identity can foster open communication and robust debate."[50] Knowing they might face retaliation, ostracism, or embarrassment, users were forced to identify themselves before engaging in speech on the Internet might be deterred from expressing unpopular ideas or seeking sensitive information.[51] As a result of the Internet's importance as a communication tool, courts have extended the protections of the First Amendment, and specifically the right to anonymity, to online speech.[52]

B. Courts have found broad identification requirements on Internet use to violate the Constitution.

A broad requirement for all users to identify themselves before being able to access the internet would almost certainly be considered overbroad, insufficiently narrowly tailored to achieve its purpose, and unconstitutional. In *ACLU* v. *Miller*, the Northern District of Georgia considered a state law that criminalized knowingly transmitting data while falsely identifying oneself.[53] The state asserted that the statute's purpose was fraud prevention. The court agreed that this was a compelling interest, but held that the statute was not sufficiently narrowly tailored to achieve its purpose because the statute would apply whenever anyone falsely identified themselves, even when there was no intent to defraud or deceive. Furthermore, the court noted that "the act prohibits such protected speech as the use of false identification to avoid social ostracism, to prevent discrimination and harassment, and to protected privacy. . ."[54] As a result, the court held that the statute was overbroad and unconstitutional.

Whereas *Miller* merely prevented people from falsely identifying themselves, in *Doe* v. *Shurtleff* the state of Utah sought to require a convicted sex offender affirmatively submit his "internet identifiers" to the state for inclusion in its sex offender registry. This would include all of the offender's email addresses, chat user names, instant messaging names, social networking pages, and passwords. Once the information was submitted, there were no restrictions on how the Department of Corrections could use or disseminate it. There were no statutory limits which prevented the Department of Corrections from "using the information to reveal the identity of a registrant who had spoken online in a non-criminal manner, or to release the information to others who wish to do so." Although he was a convicted sex offender, Doe retained his First Amendment right to speak anonymously online and the statute implicated criminal and protected speech alike.[55] Thus, the court held that the statute was not sufficiently narrowly tailored to achieve its purpose of protecting children from Internet predators and investigating online crime.[56]

These two cases show that where the government attempts to install a mandatory identification requirement without limits as to how the information can be used, the courts are likely to strike the requirement down as overbroad and unconstitutional.

[48] See DAVID KIRKPATRICK, THE FACEBOOK EFFECT: THE INSIDE STORY OF THE COMPANY THAT IS CONNECTING THE WORLD 1–8 (describing the use of Facebook to promote an anti-FARC group in Columbia).

[49] *Id.*

[50] *Doe* v. *2theMart.com*, 140 F. Supp. 2d 1088, 1092 (W.D. Wash. 2001) (citing *Columbia Ins. Co.* v. *Seescandy.com*, 185 F.R.D. 573, 578 (N.D. Cal. 1999)).

[51] *See McIntyre*, 514 U.S. at 334; *Am. Civil Liberties Union* v. *Miller*, 977 F. Supp. at 1230.

[52] *See e.g., Sinclair* v. *TubeSockTedD*, 596 F. Supp. 2d 128, 132 (D.D.C. 2009) ("Generally speaking, the First Amendment protects the right to speak anonymously. Such rights to speak anonymously apply, moreover, to speech on the Internet." (citations omitted)); *Doe* v. *2TheMart.com*, 140 F. Supp. 2d at 1093 (holding "the right to speak anonymously extends to speech via the Internet"); *Am. Civil Liberties Union* v. *Johnson*, 4 F. Supp. 2d 1029, (D.N.M. 1998) (holding that a state statute requiring website operators restrict access to indecent materials through use of a credit card, debit account, or adult access code violates the First Amendment "because it prevents people from communicating and accessing information anonymously").

[53] 977 F. Supp. 1228, 1230 (N.D. Ga. 1997)

[54] *Id.* at 1233.

[55] *Id.* at 21.

[56] *Doe* v. *Shurtleff*, No. 1:08-CV-64 TC, 2008 U.S. Dist. LEXIS 73787, at *23 (D. Utah Sept. 25, 2008).

C. Courts have only found Internet identification requirements to be constitutional in extremely limited circumstances involving convicted sex offenders.

The only courts that have found Internet identification requirements not to violate the Constitution have been considering extremely limited situations involving the tracking of convicted sex offenders on specific websites. The best example of this is the sequel to the *Shurtleff* decision. After the original decision, the Utah legislature went back and amended the statute requiring the sex offender to submit his Internet identifiers to include new limits on how the information could be used and disseminated. The Department of Corrections would only be able to use the information "to assist investigating sex-related crimes."[57] In accordance with Utah's Governmental Records and Management Act, they would also be able to disclose the information to the subject of the record, to anyone authorized by the subject, or when the information is subject to a court order or legislative subpoena. With these new restrictions in place, the court held that the identification requirements "no longer intruded into Doe's ability to engage in anonymous core political speech."[58] Because the information could no longer be used to monitor Doe's speech, the chilling effect on his speech was diminished and the registry was in compliance with the First Amendment.[59]

In a similar case, *White* v. *Baker*,[60] the court struck down a requirement for sex offenders to submit all of their Internet identifiers as overbroad, however, it provided suggestions for how such a statute would pass constitutional muster. The court held that the Georgia statute at issue went wrong by requiring *all* of the offender's Internet identifiers. First, the court noted that "a regulatory scheme designed to further the state's legitimate interest in protecting children from communication enticing them into illegal sexual activity should consider how and where on the internet such communication occurs."[61] A requirement to turn over all Internet identifiers would include an offender's identification on blogs or on shopping websites where communication with children would be unlikely or impossible.[62] Furthermore, there were few limits as to how the information, once submitted, could be used or disseminated.[63] The statute allowed the information to be used for undefined "law enforcement purposes" and even to be disclosed to the public. This opened up the possibility that the offender's speech could be monitored by government or private citizens, disclosing protected speech that the offender chose to engage in anonymously.[64] Concluding the opinion, the court noted that, because the state had a compelling interest, it had the ability to enact regulation, provided it was sufficiently narrowly targeted at the kind of interactive communications that entice children into illegal sexual conduct and the disclosure provisions of the statute were narrowed.[65]

Investigating cyber attacks is a broad use compared to investigating sex crimes and one could easily imagine it turning into monitoring of political speech on anonymous message boards or similar communications platforms. This would be an especially prevalent concern if the government required individuals to submit all of their Internet identifiers, as in *White*. Finally, there would be the ever-present specter of a data breach in the government's database, thereby risking the exposure of the identities and activities of all Americans on the Internet. Given the difficulties in narrowly tailoring the law to meet some ill-defined interest in cyber attacks, a mandatory identification scheme for Internet use may be possible, but it would probably be unconstitutional in practice.

III. Most research makes clear that attribution techniques have significant limitations.

So far, I have described how countries will deploy Internet attribution techniques for purposes unrelated to cyber security. I have also suggested that it would be unconstitutional for the United States government to impose an identity requirement for Internet users in the United States. Still, there is a clear need in the instance of a cyber attack or other types of malicious Internet use to determine the source of an attack. As one commentator has said, "[w]ithout the fear of being caught, con-

[57] *Doe* v. *Shurtleff* No. 1:08-CV-64 TC, 2009 U.S. Dist. LEXIS 73955, at *5 (D. Utah Aug. 20, 2009) [hereinafter "Shurtleff II"].

[58] *See id*. at *9–10.

[59] *Id*.

[60] No. 1:09-cv-151–WSD, 2010 U.S. Dist. LEXIS 25679 (N.D. Ga. Mar. 3, 2010).

[61] *Id*. at 48–49.

[62] *Id*. at 49–50.

[63] *Id*. at 50–54.

[64] *Id*. at 52.

[65] *Id*. at 55.

victed and punished, individuals and organizations will continue to use the Internet to conduct malicious activities."[66] But the problem is not easily solved. As Internet security expert Bruce Schneier has bluntly stated:

> Any design of the Internet must allow for anonymity. Universal identification is impossible. Even attribution—knowing who is responsible for particular Internet packets—is impossible. Attempting to build such a system is futile, and will only give criminals and hackers new ways to hide. . . .
>
> Attempts to banish anonymity from the Internet won't affect those savvy enough to bypass it, would cost billions, and would have only a negligible effect on security. What such attempts would do is affect the average user's access to free speech, including those who use the Internet's anonymity to survive: dissidents in Iran, China, and elsewhere.[67]

As I said earlier, improved attribution techniques may chill speech, including dissenting speech in repressive political and organizational regimes. This has been acknowledged by many of the current participants in the cyber security debate. One group stated that the absence of attribution, or "non-attribution," can be "vital to protecting radical ideas and minority views in oppressive regimes,"[68] and cautioned that the "[m]echanisms developed to facilitate attribution must enforce non-attribution for the purposes of sharing opinions and ideas."[69] Another group pointed out that attribution exposes political dissidents and whistleblowers to potential reprisals.[70] The Department of Homeland Security has itself made clear the need to balance attribution against the need for anonymity and free speech.[71]

Second, no matter how good attribution technologies are, attribution will probably still fail to identify the most sophisticated attackers. In the words of one expert group, "[w]hile anonymizers can be defeated in theory, there are numerous practical difficulties to achieving attribution when a sophisticated user desires anonymity."[72] Another commentator notes that "[s]mart hackers . . . route attacks through countries with which the target's government has poor diplomatic relations or no law enforcement cooperation, and exploit unwitting, third-party networks."[73] Because sophisticated attackers often obscure their trail by routing activities through multiple countries, complete attribution capability would require the implementation of coordinated policies on a near-impossible global scale.

Finally, improved attribution techniques will probably not be effective against non-state enemies, such as the al-Qaeda terrorist network. As an initial matter, non-state actors are unlikely to have access to the resources necessary to launch successful cyber attacks. As Mr. Knake has said "al-Qaeda lacks the capability and motivation to exploit. . .vulnerabilities" in our country's critical infrastructure.[74]

On the other hand, some scholars believe that terrorist groups may well have access to the sort of sophisticated computer technologies needed to conduct cybercrime.[75] Even if terrorists could get their hands on the tools needed to launch a successful cyber attack against the United States, improved attribution techniques probably wouldn't help us deter them because one of the biggest problems with non-state terrorists is that they aren't deterred by the threat of retaliation.

The National Research Council ("NRC") recently undertook an extensive review of cyber security and considered the problem of attribution in several instances.[76]

[66] Jeffrey Hunker, Robert Hutchinson & Jonathan Margulies, *Attribution of Cyber Attacks on Process Control Systems*, in CRITICAL INFRASTRUCTURE PROTECTION II 87, 88 (Mauricio Papa & Sujeet Shenoi eds., 2008). [Hereinafter "CRITICAL INFRASTRUCTURE PROTECTION II."]

[67] Bruce Schneir, *Schneir on Security: Anonymity and the Internet*, Feb. 3, 2010, available at *http://www.schneier.com/blog/archives/2010/02/anonymity_and_t_3.html*

[68] CRITICAL INFRASTRUCTURE PROTECTION II.

[69] *Id.*

[70] MATT BISHOP, CARRIE GATES & JEFFREY HUNKER, THE SISTERHOOD OF THE TRAVELING PACKETS 4 (2009), available at *http://www.nspw.org/papers/2009/nspw2009-gates.pdf.*

[71] U.S. DEP'T OF HOMELAND SEC., A ROADMAP FOR CYBERSECURITY RESEARCH 69 (2009), available at *http://www.cyber.st.dhs.gov/docs/DHS-Cybersecurity-Roadmap.pdf.*

[72] Hunker, Hutchinson & Margulies, *supra* note 66, at 91.

[73] Kenneth Geers, *The Challenge of Cyber Attack Deterrence*, 26 COMP. L. SEC. REV. 298, 301 (2010).

[74] Robert K. Knake, Expert Brief: Cyberterrorism Hype v. Fact, *http://www.cfr.org/publication/21434/cyberterrorism_hype_v_fact.html* (last accessed July 13, 2010).

[75] *See, e.g.*, CLAY WILSON, CONG. RESEARCH SERV., BOTNETS, CYBERCRIME, AND CYBERTERRORISM: VULNERABILITIES AND POLICY ISSUES FOR CONGRESS 16 (2008), available at *http://www.fas.org/sgp/crs/terror/RL32114.pdf*; Geers, *supra* note 73, at 302.

[76] NAT'L RESEARCH COUNCIL COMM. ON OFFENSIVE INFO. WARFARE, TECHNOLOGY, POLICY, LAW AND ETHICS REGARDING U.S. ACQUISITION AND USE OF

The NRC identified three reasons that deterrence by retaliation may be particularly ineffective against non-state actors:

> First, a non-state group may be particularly difficult to identify. . . . Second, a non-state group is likely to have few if any information technology assets that can be targeted. Third, some groups. . .regard counterattacks as a challenge to be welcomed rather than something to be feared.[77]

The NRC concluded:

> The bottom line is that it is too strong a statement to say that plausible attribution of an adversary's cyberattack is impossible, but it is also too strong to say that definitive and certain attribution of an adversary's cyberattack will always be possible.[78]

Based on our review of the costs and benefits of attribution techniques, there are a few key points to consider:

- The attribution of cyberattacks would greatly assist in facilitating counterattacks.
- The law of war requires an attacked body to attribute the initial attack before a counterattack will be permitted.
- Improved attribution methods would probably increase the ability to deter attacks; however, deterrence would only be effective against individuals or groups who fear retaliation.
- Attribution of activities carried out over the Internet is extremely difficult, and in many cases impossible, to achieve.
- Improvements to attribution methods will most likely fail to prevent technically sophisticated attackers from hiding their identity.
- Because Internet activity may be routed through multiple countries, including those with limited network security resources, complete attribution capability will require the implementation of coordinated policies on a near-impossible global scale.
- Improved techniques for achieving attribution of Internet activities will chill dissenting speech in repressive political and organizational regimes.
- Critical infrastructure administrators ought to be more concerned about vulnerability to internal attacks than about vulnerability to attacks from the outside.

Conclusion

Steve Bellovin, another security expert, noted recently that one of risks of the new White House plan for cyber security is that it places too much emphasis on attribution.[79] As Dr. Bellovin explains:

> The fundamental premise of the proposed strategy is that our serious Internet security problems are due to lack of sufficient authentication. That is demonstrably false. The biggest problem was and is buggy code. All the authentication in the world won't stop a bad guy who goes around the authentication system, either by finding bugs exploitable before authentication is performed, finding bugs in the authentication system itself, or by hijacking your system and abusing the authenticated connection set up by the legitimate user.[80]

While I believe the White House, the Cyber Security Advisor, and the various participants in the drafting process have made an important effort to address privacy and security interests, I share Professor Bellovin's concern that too much emphasis has been placed on promoting identification.

I also believe that online identification, promoted by government, will be used for purposes unrelated to cyber security and could ultimately chill political speech and limit the growth of the Internet. Greater public participation in the development of

CYBERATTACK CAPABILITIES (William A. Owens, Kenneth W. Dam & Herbert S. Lin eds., 2009).

[77] *Id.* at 313.

[78] *Id.* at 41.

[79] The White House, *National Strategies for Trusted Identities in Cyberspace: Creating Options for Enhanced Online Security and Privacy (Draft)*, June 25, 2010, *http://www.dhs.gov/xlibrary/assets/ns_tic.pdf*

[80] Steve Bellovin, SMBlog: Comments on the National Strategy for Trusted Identities in Cyberspace, July 11, 2010, *http://www.cs.columbia.edu/~smb/blog/2010-07/2010-07-11.html*

this policy as well as a formal rulemaking on the White House proposal could help address these concerns.

Thank you for the opportunity to testify today. I will be pleased to answer your questions.

BIOGRAPHY FOR MARC ROTENBERG

Marc Rotenberg is Executive Director of the Electronic Privacy Information Center (EPIC) in Washington, DC. He teaches information privacy law at Georgetown University Law Center and has testified before Congress on many issues, including access to information, encryption policy, consumer protection, computer security, and communications privacy. He testified before the 9–11 Commission on "Security and Liberty: Protecting Privacy, Preventing Terrorism." He has served on several national and international advisory panels, including the expert panels on Cryptography Policy and Computer Security for the OECD, the Legal Experts on Cyberspace Law for UNESCO, and the Countering Spam program of the ITU. He chairs the ABA Committee on Privacy and Information Protection. He is a founding board member and former Chair of the Public Interest Registry, which manages the .ORG domain. Rotenberg is editor of "The Privacy Law Sourcebook" and co-editor (with Daniel J. Solove and Paul Schwartz) of "Information Privacy Law" (Aspen Publishing 2006). He is a graduate of Harvard College and Stanford Law School. He served as Counsel to Senator Patrick J. Leahy on the Senate Judiciary Committee after graduation from law school. He is the recipient of several awards, including the World Technology Award in Law.

Chairman WU. Thank you very much, Mr. Rotenberg.

Now it is in order for questions, and first I want to note that we in Congress sit on multiple Committees, and as is frequently the case where there are two flies flying in the Grand Canyon, they collide, and I have votes occurring right now in my other Committee and I will have to excuse myself after asking this first set of questions, and I aspire to come back because this is a very, very important topic that I care about very much.

Secondly, I would like to welcome our friends from Russia TV Today. I understand that Russia TV Today has also broadcast one of our NASA hearings. It is not unusual for foreign media to take a stronger interest in topics of importance to the United States more so than American media does at times, and we welcome our Russian friends. But we also want to note that the usual process is to accredit into the Committee prior to attendance, but you are welcome to stay today.

Now, I think that each of the witnesses referred to both in your spoken and oral testimony that there may be some limited role for deterrence and that there may be some greater role for attribution in protecting legitimate interests on the Internet, but that both deterrence and attribution to different extents are overplayed in the current discussion. I would like each of the witnesses to the extent you can or want to address first that opening query about deterrence and attribution.

Mr. ROTENBERG. Well, I will jump right in and I am sure the other witnesses will make comments. I cited in my testimony the conclusion of the National Research Council report because I thought this was a very thoughtful point they were making, particularly with non-state actors. They said attribution would be difficult. We are talking about entities that are typically outside of the United States so you would need an attribution technology that is global, not easy to identify outside the United States, not much of a technical infrastructure, which means that there is not much opportunity to respond, and with some of the non-state actors, it is not even clear they wouldn't mind being identified. It is almost the exact inverse of the model that we had during the Cold War in our relationship with the Soviet Union, and I think the National Research Council report makes this point very well.

Mr. GIORGIO. Yes, I would like to add, even in the hearing background that was put together by the staff, we talk about attribution not only from a point of view of identifying the person who is on the other side but perhaps just identifying at least the location they are coming from. So if you have a purist view of attribution, I certainly agree that it is extremely difficult technologically to guarantee you know who the human person is on the other end, but that doesn't mean that some attack attribution technology wouldn't give us lots of information which could be used for other purposes such as shutting down the computer at the other end independent of who is on it. Thank you.

Dr. WHEELER. If I may speak as well, as I noted earlier, there is no possibility of having absolutely perfect defenses, so I believe there is value for attribution. On the other hand, we have to admit that attribution itself is difficult and there are some serious limitations to that as well. You know, attackers can cause attacks to be delayed and perform their attacks through lots of intermediaries and often can make it very difficult to attribute when they don't want to be attributed. And so basically I think computer network defense shouldn't depend on attribution, it should be part of a larger strategy having basically multiple tools in the toolbox.

Mr. KNAKE. The only comment I would add is that for the last decade our strategy for preventing another major terrorist attack on U.S. soil has both been effective and does not in any way materially rely on deterrence so I think that may be a better model for how we deal with the cyber threat, to focus on prevention, to focus on protection, to focus on resiliency rather than to focus on trying to deter cyber actors. The only other point I would make is that in a lot of cases we don't lack attribution, we lack response options. We don't know what we should do when we discover that the Chinese have hacked into Google in 30 other countries. We seem to have fairly good evidence that they did that. We have traced the attack back. We have then asked for an explanation and we have not received it. I am not sure how better attribution one further layer down would help resolve that problem. Similarly, with French intelligence or Russian criminals, Nigerian scammers, we know their national origins. We simply lack response options and a mechanism for cooperating and requiring cooperation internationally.

Chairman WU. Thank you very much. Because there are votes going on and not only votes for me in my other Committee but I am told close votes, I am going to ask one further question and then I am going to step out and aspire to return promptly after those votes.

Thank you for your answer to the deterrence and attribution question and its utility. Following up on that, I think several of you, perhaps all of you have noted that to the extent that there is a deterrent utility and that there is a capability for attribution, that there is also potentially or there is a drastic effect on speech and free flow of information, and I think, Mr. Giorgio, you stated in your written testimony that there is a necessary tradeoff, and I don't know if others put it quite that crisply, but can you address that issue to the extent that we put attributability capability into the backbone of the Internet that we would be decreasing anonymity, freedom of speech and freedom of inquiry? Whoever wants to start with that.

Mr. GIORGIO. Chairman, since you referenced me, let me also say that I do believe that we need protocols with a lot more privacy in them, and I am very troubled by the situation today because frankly a lot of people learn information about us that they shouldn't need to know in, for example, a financial transaction. So it is very important that we build new protocols to protect anonymity or privacy, I should say, when it is called for.

Mr. ROTENBERG. I should say also, Mr. Chairman, that many businesses that operate on the Internet have identification requirements. In fact, there is a big controversy right now involving the company Blizzard, which offers World of Warcraft, and they are now requiring the use of true names for people who come in the forums and it has, you know, provoked a big discussion about, you know, identity requirements as a way to make people a little more hospitable online, but the key point here is that whatever decisions private companies might make about identification is really very different from a government-mandated identification requirement, because what a government-mandated identification requirement does is basically hold out the specter that if you say something that is unpopular and the government can trace it back to you, the government can hold you accountable, and I think that is really anathema to our view in the United States of freedom of expression, and so it concerns us, of course, that a government-mandated identification requirement wherever it may be imposed in the world could have a similar impact on political speech.

Mr. KNAKE. I think I would echo those comments, but I would also add that I see the equation in need of being reversed. I actually think government needs to do a better job of protecting the privacy of users in the commercial arena. That is where the biggest threat to privacy is today. The reliance on anonymity, which is still very, very useful for protecting freedom of speech and is useful for protecting freedom to access information, is not useful in the context of communicating, banking and interacting the way we do online and increasingly commercial web operators are tracking their users without telling them by downloading cookies onto their computers, some very insidious forms, and using other geolocation technologies that your browser, your computer, your Internet service provider and the services that you are using online are all by default not going to tell you that that is going on so essentially you surrendered your anonymity without knowing it, and in my view, government needs to step in to create some form of disclosure that is upfront and obvious to the average Internet user that for the free content they will be tracked and that will be used to target advertising at them.

Dr. WHEELER. If I may jump in also, first of all, getting back a little bit to the original question, clearly attribution technologies have potential to greatly harm anonymity, pseudonymity, privacy and so on but it is not the same for all the different technologies.

Some technologies are much riskier than others. I cite probably the more egregious example, recording every bit that goes back and forth between a user and everything else has radically different effects than storing much smaller pieces of information, you know, fingerprints and so on. So depending on what is stored and how it is stored makes a big difference on the effect on anonymity and privacy and pseudonymity.

Mr. GIORGIO. May I make an additional——

Chairman WU. Mr. Giorgio, yes.

Mr. GIORGIO. Thank you. You know, I think credibility is very important when we decide who to listen to, so whether it is the distinguished Members of this Committee or my distinguished colleagues, when they speak, I want to listen because I know what they have gone to get to the position they are in today. So all of that is lost when people speak with anonymity, and so I would—and even during emergencies, it would be very important to me, for example, if somebody who is reporting from ground zero if I have some confidence that they are actually at ground zero. So the credibility of listening to what people have to say is tied up to some extent in being able to attribute who they are, what their past is, how they came to be in that position and why we should listen to them, and where they are. Thank you.

Chairman WU. Thank you all very much. I am going to hand over the gavel to the gentlelady from Maryland, Ms. Edwards, and before I do that, I will recognize Mr. Smith for his questions.

Mr. SMITH. Thank you, Mr. Chairman, and I appreciate the opportunity, and I would also like to briefly note that it is my understanding a follow-up hearing in which we hear from NIST, the National Science Foundation and other relevant Federal agencies is under consideration, and I would certainly like to offer my support for holding such a hearing.

Regarding the questions that I have, I was wondering if you could just share what you think are the best methods for tracing the attacks, anyone? Maybe start with Dr. Wheeler.

Dr. WHEELER. That actually turns out to be more difficult than you'd like. I would like to give you a very simple, "there it is, there is the one solution," and of course, life is often more complicated than we wish it could be. Actually, what is intriguing, when I started writing this particular paper that I mentioned earlier and I submitted as testimony, I didn't expect there to be many different possibilities to do this, and it turned out in fact there are a very large number, and although I haven't worked on this particular area more recently, the number can only go up. So there turns out to be a remarkably large number of ways, and unfortunately what it really turns out to be is, I suspect people aren't surprised when you go to technologies, there are various tradeoffs. Some of the techniques are particularly helpful for tracking down what is called denial of service attacks. You are being attacked, sent a lot of messages, maybe from many different places, and there is basically constant streaming of data. In that case, the very fact that someone is constantly sending messages to you and trying to overwhelm your systems means that you can try to track back, "well, I just wait for the next one and start looking backwards that way," for example. But of course, those techniques that depend on that don't work for many kinds of attacks where in fact that isn't what happens, it is a few messages and all of a sudden your systems are down or something terrible has happened. So I don't believe there is a single answer. There is a set. And one other good thing about that from the point of this particular hearing is that some of them are much more egregious or concerning in terms of privacy and attribution. Probably one of the more extreme examples I guess would be what is informally called hack backs where you actually say, "I am being attacked, I am breaking into the computers backwards to find out where that comes from." Unsurprisingly, that is severely restricted by U.S. laws, as well it should be. But sometimes, particularly if those systems are under control of outside powers and it is really critically important and nothing has been pre-positioned that may be one of the few techniques available.

I will quickly note, though, that a number of these techniques fundamentally require pre-positioning. You can't wake up in the morning and say, "I would like to know where this attack came from." Many of these techniques require systems to be already in place before you can do the attribution, and I think that is one of the reasons why discussions and hearings like this are necessary, because if we the United States wish this kind of capability, we are going to need to put things in place and thus that requires this kind of discussion that we are having today.

Mr. SMITH. Thank you.

And since I have limited time, I also want to note, Mr. Rotenberg, in your testimony you said that no matter how good attribution technologies are that it will probably still fail to identify the most sophisticated attackers. So I guess I have to ask the question, are our efforts futile, and if other attribution technologies will not be able to get the job done, what are the other options for protecting us from cyber attacks?

Mr. ROTENBERG. Congressman, thank you for the question. I don't think they are futile, and I think it is important particularly for us to improve our security through education and open standards. I think it is important to develop better forensic techniques so it is possible to trace back attacks, as Dr. Wheeler described. I will also mention that, you know, one of the key problems here which was uncovered in a workshop shortly after 9/11 that I participated in where people were talking about attribution, Admiral Poindexter brought us together and said well, how do we solve this problem, and someone said well, you could, you know, hash a person's unique DNA against every keystroke so that everything that went from your keyboard, every single stroke was uniquely defined to, you know, tied to a biometric identifier, and people said "wow, we have solved the attribution problem, isn't that great," and someone said "well, what if you have a guy standing next to the user with a gun telling someone who is authorized to type into the keyboard, now what do you do?" In other words, you can have perfect attribution in a hostage situation, and by the way, probably a good plot for a movie, and still not be able to prevent a smart attacker, which I think reveals really how difficult this challenge is. I am not saying we shouldn't improve security or pursue good forensic techniques. I just think it would be a mistake for practical reasons in addition to human rights reasons to place too much emphasis on attribution.

Mr. SMITH. Okay. Thank you.

Ms. EDWARDS. [Presiding] Thank you, and thank you to all the witnesses today. I just have basic questions kind of as a consumer. All these questions revolve around balancing the need for security against the protection of privacy and so where do you strike that balance.

Mr. Rotenberg, I wonder if you could tell me, almost every website on the Internet uses cookies to collect data over activity. As a consumer I know I get to make a decision, do I really want to type in all of that personal information that they ask me or go through the list of things until I find out that I actually don't have to give them that information at all unless, if I check the box way down at the bottom after scrolling and scrolling and scrolling, and then you get free services in exchange for turning over all of your information and so there are instances, for example, where the user wants to do that and so they make a decision. There are other instances for some reason to get something sent to your home, the commercial enterprise has to have it, otherwise they can't mail what it is that you want. And so how is that the need to protect the user privacy being as important as it is can the Federal Government help me, the average Internet user, understand what my options are and what the consequences are for sharing that information, for sharing it at that moment, but also the longer term consequences once that information is housed someplace or other or shared with some other source?

Mr. ROTENBERG. Congresswoman, thank you for the excellent question. While on the national security side I imagine there is a sense that there is not enough attribution, I can tell you on the consumer side, there is a sense that there is way too much attribution, which is to say that when someone does a Google search, you simply type in, you know, apartments, Virginia, because you are interested in trying to find an apartment in Virginia. I bet no one has any understanding or very few people do that at that moment in time Google will record the time and the day when the search was made, the search query, the cookie tied to the user ID. If they have a unique identity, the IP [Internet Protocol] address for the device, that will also be recorded. All of this information will be collected and stored by the company for every single search and kept for months and maybe years building this enormous profile, and from the privacy perspective, we think that is very invasive. It even creates some security risks if the information is misused. In fact, part of the great concern about network vulnerability, Google's experience in China was that they essentially lost control over a lot of sensitive information because of internal vulnerabilities that were exploited. That information that they lost control of included a lot of personal data on Google users. So we think on this side, the government actually has a role in protecting consumer privacy by limiting the amount of data that is being collected and giving people more control over that data.

Ms. EDWARDS. Thank you.

And then Mr. Giorgio, you mentioned in your testimony that the bulk of the privacy concern is actually directed at our own government. I was reading, I think just in the last day or so, about the National Security Agency program, Perfect Citizen, and while there is this need obviously to safeguard our infrastructure, whether it is our nuclear plants, the power grid, etc., there is a concern that using a tool like that could then really impede on all of our individual privacy giving up that anonymity that you have described as a constitutional protection but we have to rely on the government to really protect us from all the bad actors. So I wonder if you could discuss the difficulties in achieving both security and privacy, especially when the bad guy of one concept is the protector of the other and in an environment where if the bad guys are operating in concert, that is kind of one thing, but we have a whole bunch of just bad actors, whether they are from Nigeria trying to get my mother's money or from someplace else, and those set of actors may be uncoordinated, they may be individuals, and to draw a national security concern around trying to protect against those kind of actors is, I think, a little complicated.

Mr. GIORGIO. Yes. Thank you, Congresswoman. I couldn't agree more. When Mr. Rotenberg just made his point, I agree with him that we may fear government least of all. It is these companies who have all these databases that are a true threat to us. And if we look at what is happening in many of these databases that are being collected, for example, all the databases that bind our physical location to our use of wireless devices such as cell phones, these are all in the hands of the private sector, and it is quite easy, and in this country they are in the hands of the private sector. I wouldn't go overseas and wander about with a cell phone turned out, you know, if I wanted to protect my anonymity or privacy, and so I see it over and over again that there is a myriad of bad actors out there, the least of which may be government, and as you point out, government does have a role to protect our critical infrastructure but I am not sure they are the greatest threat to our privacy.

Ms. EDWARDS. Mr. Rohrabacher, I think you are up.

Mr. ROHRABACHER. Thank you very much.

You know, the last point that was made was very interesting. If you are in a relatively free society, that may be true. In a relatively dictatorial society, the opposite is true. And the idea of how you—what you demand of people who involve themselves in this arena of affairs in a society, it is a very complicated issue and it is, for example, where I happen to believe in the maximum degree of individual freedom. I can also understand that in France, for example, they don't want to say women shouldn't wear a burka, all right, but there are some national security implications to that rather than just cultural implications as well. We don't permit people to go around hiding their identity as they are walking around the street, or do we? Do we in this society?

Mr. ROTENBERG. Well, it is a very interesting point, Congressman. Actually the United States unlike most other countries does not allow its police to ask people on the street to present identity documents.

Mr. ROHRABACHER. Right.

Mr. ROTENBERG. There actually has to be some suspicious activity that provides a reason for the police to be able to say to someone, may I see, you know, some identification. It is not true in most countries. In many countries, you can be asked without suspicion to identify yourself.

Mr. ROHRABACHER. I am wondering if a person wearing a mask, if that would be suspicious activity.

Mr. ROTENBERG. Yes, it is, and we actually do have anti-mask laws in many states in the United States, so that is generally not permitted. But as for your identification, that is something that we tend to allow people to keep to themselves.

Mr. ROHRABACHER. This is of course what we are talking about, cyber attacks. It is very similar to the idea, the challenge faced by the entertainment industry of people who are unlawfully making copies and downloads of material. I guess that is sort of a cyber attack. Is there technology that any of you know about that you believe that—is this a technological solution rather than a government regulatory solution?

Mr. GIORGIO. So there are problems that require authentication and authorization, knowing who people are and what they have access to do, and there is a tremendous amount of very good security research and in fact solutions today that provide these strong access controls. Digital rights management, which protects music, you know, is one form of those controls. The goal of those controls is not dissimilar to the DoD goals of trying to protect information. So as technology gets developed in various places, it is frequently leveraged for other purposes.

Mr. ROHRABACHER. Is the technology solution a wall or is it a retaliatory strike, you might say, against someone who has come into your system?

Mr. ROTENBERG. Well, in the copyright arena, it is actually a tracking technique. As Mr. Giorgio mentioned, digital rights management is much like a watermark and it basically allows an entity both to assign its ownership of a product, of a digital product and also identify who the appropriate user is. So if it is in the possession of someone who didn't properly acquire the song or the movie, they will essentially be tracked down through that digital watermark.

Mr. ROHRABACHER. Is it possible in dealing with the hackers and dealing with these types of cyber attacks to have a situation if someone doesn't have an authorization to be where they are electronically that there is an instant retaliation against their own equipment, meaning a disintegration of the system that is the vehicle for this aggression?

Mr. GIORGIO. So that capability is possible. You know, whether or not it is actually done anywhere, I don't know.

Mr. ROHRABACHER. Is that something that we should strive for?

Dr. WHEELER. This is David Wheeler. Is it possible? I agree with him, yes. Should we do it? I would be extremely hesitant. As I noted in my paper, attribution is something that although it can be done, there is also the risk of misattribution, and indeed, for some attackers, that may be actually their primary goal is to try to accomplish misattribution, perform their attack and cause misattribution of the attack.

Mr. ROHRABACHER. Oh, I see.

Dr. WHEELER. And so therefore that doesn't mean under no possible circumstance could we never imagine this but I would be very hesitant about installing such an automatic counterattack system generally for most kinds of—you know, certainly for military systems you want a human in the loop double-checking first.

Mr. ROHRABACHER. Well, just one note, and I know my time is up after this, and I don't know how to pronounce your—is it——

Mr. KNAKE. Knake.

Mr. ROHRABACHER. Say it again.

Mr. KNAKE. Knake.

Mr. ROHRABACHER. Okay. I have surfer's ear in this ear and I have trouble——

Mr. KNAKE. I am sorry. It is Knake.

Mr. ROHRABACHER. Knake. You mentioned that efforts made after 9/11 actually identifying methodologies actually had a major impact in preventing another 9/11. I would suggest it is not just identification, however. It is identification and retaliation. If we just had identified potential al Qaeda terrorists since then and let them be, we would have had another 9/11. We aggressively sought them out and in some cases killed them, which was good, or sent them to Guantanamo, which is debatable, but there was actually an action taken so the identification isn't the only step that needs to happen if we are to protect ourselves from the electronic type of aggression. You can answer that if you would like.

Mr. KNAKE. Thank you, sir. I think that is absolutely right, and I think I would go a step further. Prior to 9/11, the United States roving ambassador for counterterrorism, Michael Sheehan, delivered a very stern message to the Taliban which was essentially, if we are attacked by al Qaeda who plan their attack on your soil, we will hold you responsible for that. The Taliban did not get that message until after 9/11 but we followed through on that. So essentially we assigned responsibility to the Taliban for the activities carried out by a terrorist organization on their soil. Their failure after 9/11 to cooperate with apprehending bin Laden resulted in the invasion of their country. So I think it is actually very analogous to the situation we want to move to in cyberspace where if a country refuses to cooperate in an investigation that attributes the attack to a system or an individual in their country, we in turn hold them responsible for it.

Mr. ROHRABACHER. Thank you very much. That was very astute, and I appreciate you permitting me, Madam Chairman, the right of questioning because I am not a member of this subcommittee. But thank you for allowing me to do that.

Ms. EDWARDS. Thank you, Mr. Rohrabacher.

I just have one question. We are going to take one question. We have been called for votes. The Chairman will come back and so we are actually going to recess. He is on his way back and so I am just going to stall and ask my question.

Mr. Giorgio, it is actually an important question. You discussed the need for standards in a lot of areas and you say that government should actually invest in this development but allow standards development organizations like the Internet engineering task force to develop them through normal processes, but Mr. Knake has testified to the difficulties involved in using these processes to produce standards, specifically new protocols and advocates for more government involvement. How can the Federal Government better protect the development of consensus-based standards?

Mr. GIORGIO. So Mr. Knake is quite accurate on that point. It is extremely difficult to get these standards pushed through the standards bodies, even when various governments are behind them. So I think—but first and foremost we have to develop the technology that will allow us to propose those standards in the first place. In parallel, we have to work with the standards committees, however difficult that is, and try and influence the course of those standards.

Ms. EDWARDS. Mr. Knake, there are just so many different agencies, though, whether you are talking about the DoD, the FBI, I mean, just all of these various agencies that all use so many different tools. I mean, it does feel very daunting to then create a standard for the multiple tools that are used within these agencies. Do you have any comment about that?

Mr. KNAKE. I certainly would recognize the problem that you are highlighting. I think in a couple of areas, however, it is a narrower issue, particularly for the main suite of Internet protocols which are universal, and I think we have a fairly good set of what are the security problems with those protocols and how they should be addressed, essentially how do we secure them to a standard to which they cannot be abused but not to a standard in which attribution becomes ironclad across the Internet, and so that is the area where I think we need to return to a situation of more government intervention. These protocols were initially developed for the Defense Department with U.S. government funding. I think a similar initiative now would be in order in an effort to address the vulnerabilities that were introduced in that original protocol suite.

Ms. EDWARDS. Thank you very much, and I see the Chairman has returned and so I will let him take it from here, and thank you very much.

Chairman WU. We have about seven minutes before Floor votes, and I frequently talk about having three rings going in this particular circus at any given time, at least when we are here in Washington, and that is why it takes more time when we are home in our districts because we can only do one thing at a time there. I have several more questions. If the minority does not, I will try to get my questions in before we go vote on the Floor, but let us see how we do.

Based on both your spoken but particularly your written testimony, I get the impression that you all are of the opinion that there is limited utility of any particular security technique, and that some combination of techniques would afford us potentially the best combination of security and privacy. Is that roughly accurate?

Mr. ROTENBERG. Yes.

Dr. WHEELER. Yes.

Chairman WU. Okay. If that is the case, is it further sort of what you overtly state or what you imply that perhaps we have a system of networks in our country or in the world which are best served by different degrees of security and privacy/anonymity, that is, we might set a different standard for those networks dealing with publicly available information or journalism or blogs and opinions, we might set a higher standard for networks dealing with utilities, the power grid or banking or financial transactions and we might set again an even higher standard for, let us say, DoD or NSA types of networks. Can you address that?

Mr. ROTENBERG. Well, Mr. Chairman, I think there are a couple different ways to think about it. Certainly we have within the United States and in the military community, for example, secure networks that are essentially not connected to the public open Internet, but with respect to the public open Internet, I think as much as possible we want to keep systems connected because of all the benefits that the Internet provides and place the added security obligations at the end points. In other words, if there are applications or organizations or entities that have needs for enhanced security, for example, a password and user ID is a simple one, you know, place the responsibility there, and as much as possible maintain the common protocols of the public Internet for general use. Now, that is not to say, as I said at the outset, that clearly there will be segregated networks for specialized purposes but I am concerned as, you know, Vint Cerf and others have expressed concern about the possible balkanization of the Internet if we start carving things up too much. Literally separating parts of the network out from other parts, we will lose a lot of the benefit.

Mr. GIORGIO. Sir, I am on the DARPA [Defense Advanced Research Projects Agency] oversight board with Vint Cerf on an issue related to this, and I completely agree with Mr. Rotenberg that, you know, we have to preserve as much as possible for common use, okay? However, when somebody is providing a service at one end of the network and somebody somewhere else in the world is trying to use that service, it is the responsibility of that endpoint to enforce the protocol that they will demand that person to use. So they might be on the same backbone but we might have very different protocols running through that and effectively have different networks, but we don't want to physically separate them, and I think Marc said the same thing.

Dr. WHEELER. If I can jump in here also, I very much by the way agree that there are different levels of anonymity, privacy desires comparing, say, the public Internet versus, say, you know, a network inside the DoD that involves classified information or weapons systems or something. You would expect a whole lot less anonymity in the latter situation. I think the interesting thing is that there is somewhat odd good news that attribution often tends to be a lot easier against insiders. We were talking about this before while you were out, Congressman Wu, but many of these attribution technologies fundamentally require pre-positioning. You have got to put the technology in place ahead of time. That tends to be easier to do inside a smaller closed network. The DoD is of course large but nevertheless it is certainly not as large as, say, the United States as a whole or some such and therefore when you have a smaller network, you can treat it as inside an organization. It is much easier pre-positioning things. And so in that sense, at least, you can put attribution technologies available that perhaps at least will tell you well, he is inside and there he is, or he is outside and now at least maybe I should start closing off the gates for them to come inside.

Chairman WU. Some of you have addressed the need for standards for the operation of anonymity services like Hotspot Shield, and I think the argument is that because these services make it easier for folks to do all sorts of things anonymously that there is an interest in different forms of access or identifiers in order to gain this level of anonymity, and there may be a difference of opinion on this issue and I would like to have that specifically addressed.

Mr. ROTENBERG. Well, let me say that, you know, pure anonymity means that you really can't trace back to the user. Now, there are a lot of escrow-style configurations where you can allow people to conceal their public identity but still put a responsibility on a service provider to say, for example, with a warrant we now need to know who this person is and this isn't true anonymity but it gives, you know, many of the elements of anonymity. Here is the hard problem. You know, true anonymity, which we think is important, will protect the political dissident in a country that is hostile to the person's views and may in fact imprison the person if his identity is known. Pure anonymity will also protect the pedophile who is trying to distribute images on the Internet and should be prosecuted and imprisoned. And do you see in this one tool, you know, there is one application that we would value very much and another application that we would try to prevent, and if we go the half step in and we say, well, maybe we should allow this through a pseudonym escrow service, it will be easier to catch the person engaging in the transfer of child pornography but it will also be easier to catch the human rights advocate. It is not a simple problem.

Chairman WU. Well, that is what I was thinking about in reading the testimony. One of the trapdoors is, if you get a legitimate judicial decree asking for identification in connection with a crime, well, we in our society would view pedophilia as very legitimate for such a judicial decree, and it is my impression that there are other countries where for what we view as vague crimes like breach of state security which can cover a whole host of activities that in this country we view as legitimate that that may result in the issuance of a valid judicial decree, and the question is, how does the third party respond to such a judicial decree which on its face these two decrees are indistinguishable?

Mr. ROTENBERG. That is the dilemma.

Mr. GIORGIO. I think we need to rely on other types of third parties in these circumstances. It might be perfectly okay for me to positively identify myself to my identity provider but then perhaps that identity provider could enable me to talk to a search agent, for example, and maintain my privacy. The identity provider might be blind to everything I do and the search—the service doing the searching for me doesn't know who I am but yet because that privacy is provided to me by a third party.

Mr. KNAKE. I would only add that if what you are looking for is anonymity, there is a limited number of reasons that you really need that. It is freedom of speech, it is access to information. So restricting the ability to use these services for transactions can cut down on a lot of criminal behavior and a lot of network infiltration.

Chairman WU. If there is no further answer on this question, the rules of this Committee preclude us from recessing and reconvening without a minority Member present, and since that apparently is not possible, I am going to adjourn this meeting momentarily. I do want to point out—well, there are many additional questions, many additional topics to be covered. You all have prepared very thorough presentations, and it is normally the practice of this Subcommittee in addition to asking many questions to give you all an opportunity to say anything in addition that has not been asked. We apparently will not have that opportunity today. There will be written inquiry of each of you. In particular I am curious as to the confidence that the legal analyses that some of you all have presented, your level of confidence since these are district court opinions, and I also want to commend the law clerks for having done a fine job. I just want to add that I think there is enough material here for an interesting law review note or maybe several law review notes, and also in particular I would like to have addressed the role of international agreements, international standards and agreements about what constitutes a breach, what constitutes an attack, and what kind of standards there should be for the various technologies for attribution or otherwise, and finally, I think that addressing the issue of standards in general needs to be further fleshed out.

I want to thank you all for your presence, for your tolerance for the wrinkles in Congressional operation, and as I said to some of you before the hearing began, you prepared very, very thoughtful, thought-provoking and dense materials. It is as if I were trying to reduce to five or ten pages how Congress really works, the version that is not in your high school civics textbooks. It would require a lot of parsing of what is between the lines.

I want to thank you all very much for being here today. The Subcommittee hearing is adjourned.

[Whereupon, at 11:19 a.m., the Subcommittee was adjourned.]

Appendix:

Answers to Post-Hearing Questions

ANSWERS TO POST-HEARING QUESTIONS

Responses by Dr. David A. Wheeler, Research Staff Member, Information Technology and Systems Division, Institute for Defense Analyses

Questions submitted by Chairman David Wu

Q1. Information sharing is critical for success in cybersecurity, whether it supports attribution of attacks or awareness of vulnerabilities. How important is it to have common nomenclature, common metrics, and standard sharing methods for success in information sharing? How should these different elements be developed, which government agencies should be involved, and what roles should they play throughout the process?

A1. In any technical endeavor it is important to have some common nomenclature, common metrics, and standard sharing methods in the areas most important to the task. In many cases these should be developed through a partnership between government, industry, and academia. The government organizations that should be involved should include those in charge of defending the country and/or involved in information technology (IT) standards. These government organizations include the Department of Defense (DoD), the Intelligence Community (IC), the Department of Homeland Security (DHS), and the National Institute of Science and Technology (NIST).

Q2. Many of you have discussed the need for new internet protocols to be built on the concepts of security, authentication, and attribution. What parties would help develop and implement these protocols and what would their roles be? Who would use these new protocols and would multiple protocols diminish the utility of the internet?

A2. I do not believe there is a need to *replace* the existing suite of Internet ("TCP/IP") protocols with radically different protocols. Even if this were desired, the cost and effort to make this switch would exceed any likely benefits. For example, organizations are currently adding support for version 6 of the Internet Protocol (IP), in addition to version 4, yet this *minor* change is taking more than a decade to complete. Thus, instead of wholesale replacement, there is primarily a need to develop new protocols (for new functionality) that build on *top* of the *existing* protocols. In a few cases there may need to be extensions of existing protocols (to add new capabilities) but this is still different from replacement.

There are already standards-setting bodies whose purpose is to develop and promulgate Internet protocols, such as the Internet Engineering Task Force (IETF) and the World Wide Web Consortium (W3C). The government, industry, and academia should gather within these standards-setting bodies help develop the specifications of these protocols. Where attribution-related standards are involved, "attribution techniques that pose less danger to privacy should be the ones most encouraged."[1]

The internet already has many protocols; as long as each protocol performs a specific task not performed by others, this is not a problem. However, having multiple incompatible protocols with the same functionality *does* bear the risk of diminish the utility of the internet, due to incompatibilities between parties.

The key mechanism to countering such incompatibilities is for users to insist that their systems, including all network protocols, must be built using *open standards*. "Standards should be publicly defined and held. This way, no single vendor controls others, permitting competition."[2] Any patents possibly present on parts of the standard must be made irrevocably available on a royalty-free basis. This is because a "standard that cannot be implemented without a patent license gives a special advantage to the patent holder(s). Such patents constrain or prevent competition, and thus undermine the advantages of standards listed above"[3]). There must be no constraints on the use and re-use of the standard (since such constraints would threaten to balkanize the Internet). The standard's specification document should be available without fee over the Internet (the IETF and W3C already do this), enabling all to copy, distribute, and use the standard freely.[4]

[1] Wheeler, David A. and Gregory N. Larsen, "Techniques for Cyber Attack Attribution," Institute for Defense Analyses Paper P–3792, October 2003 (hereinafter referred to as "IDA 2003 "). Section 3.13.

[2] [IDA 2003], section 3.7.

[3] [IDA 2003], section 3.7.

[4] This definition from Digistan is available at *http://www.digistan.org/open-standard:definition*, and is a clarification of the definition by the European Union (EU) European Interoperability Framework (EIF).

Many attribution "techniques are immature and will require funding before they are ready for deployment. If the [government] wishes to have a robust attribution capability, it must be willing to fund its development and deployment."[5]

Q3. Please discuss how the level of confidence can have an impact on the utility of attack attribution. Please relate the level of confidence to the spectrum of available responses including diplomatic, economic, cyber, and kinetic.

A3. Responses that are especially damaging or non-reversible, such as kinetic responses, should be avoided unless the attribution confidence is extremely high, typically through confirmation by multiple methods.

One issue that must be kept in mind is that attackers may "wish to cause misattribution as their primary purpose, rather than actually be successful at the attack. For example, if there is already tension and conflict between two adversaries (e.g., two countries A and B), a third party (C) could try to attack one (A) and cause the attack to be misattributed to the other party (B). Thus, the third party could escalate a conflict between others simply by forging attacks."[6]

Ideally, "an attribution process would also report the confidence level in the attribution, but this information is often not available."[7] In some cases, using multiple techniques and using techniques that resist misattribution can increase confidence. Fundamentally, however, "computer network defense should not depend on attribution. Instead, attribution should be part of a larger defense-indepth strategy."[8]

Q4. Are there any other thoughts or issues you would like the share with the Committee on attack attribution and cybersecurity?

A4. As noted in my paper, a good first step would be to "*change the terrain*" of our computer networks so that attacks are less likely to be successful or are more difficult to hide. We need to harden our information technology (IT) systems (including clients, servers, and network components) to resist attack far better than they currently do. This is partly because this reduces the need for attribution, and partly because this makes them more difficult to exploit as intermediaries. We should harden our routers and hosts so that attribution is easier (e.g., limit the use of spoofable protocols and disable broadcast amplification/reflection). Finally, we should consider implementing network ingress filtering on government networks at all levels, so that data packets cannot cross between networks unless they truly could be from the claimed network.[9]

We should decrease the number and impact of vulnerabilities in commercial software (both proprietary and open source software) we use, via:

1. *Education*. We should try to ensure that all software developers know how to develop *secure* software. This knowledge includes knowing the common mistakes and methods to prevent these mistakes. Since the U.S. economy depends on software and nearly all software connects to a network or uses data from a network, practically all software developers now need this knowledge. Unfortunately, secure software development education is often available only as an optional graduate-level course.
2. *Improved tools and standards*. We should enhance software development tools (such as programming languages and key libraries) and their standards so that writing secure software is much easier, mistakes leading to vulnerabilities are much less likely, and mistakes are easier to detect before the software is released to users.

The government should consider becoming even more involved in the development and deployment of open standards. It is currently government policy to encourage the use of commercial items where applicable, for reasons that are well-understood. However, commercial items are less likely to support government needs and concerns if the standards they are based on were not developed with those considerations. The government has unique needs and concerns, both as a user and as a representative for the people of the United States, including issues around cybersecurity, privacy, and anonymity. It should be noted that in some cases the government is already involved in standards development, and in some cases the government asks if the commercial products it buys meet the relevant standards. However, to ensure that commercial products will be suitable for its own use and use in the country, the government should ensure that it has "a seat at the table"

[5] [IDA 2003], section 4.

[6] [IDA 2003], section 3.15.3.

[7] [IDA 2003], section 3.15.3.

[8] [IDA 2003] section 4, conclusion 2.

[9] See [IDA 2003], especially section 4, conclusion 6.

when key information technology standards are set, ensure that those standards are open standards, and require that the commercial items it purchases correctly implement the relevant standards.

Questions submitted by Vice Chair Ben R. Luján

Q1. The Fourth-generation of cellular wireless network standards being developed uses the internet protocol suite and would extend the internet to cellular devices. What are the implications of this 4G standard for this discussion on privacy and attribution?

The Internet protocols have long been demonstrated and used for wireless communication. Indeed, DARPA experiments in the 1970s demonstrated that packet radio networks could interact with other networks using protocols that eventually became the Internet protocols. However, I have not evaluated the 4G standards in depth for their implications on privacy and attribution, so I cannot give a specific answer about the 4G standards. If the government is concerned about the privacy or attribution affects that 4G standards could have on itself or its citizenry, it should be involved in the development of those standards.

ANSWERS TO POST-HEARING QUESTIONS

Responses by Mr. Robert Knake, International Affairs Fellow, Council on Foreign Relations

Questions submitted by Chairman David Wu

Q1. Information sharing is critical for success in cybersecurity, whether it supports attribution of attacks or awareness of vulnerabilities. How important is it to have common nomenclature, common metrics, and standard sharing methods for success in information sharing? How should these different elements be developed, which government agencies should be involved, and what roles should they play throughout the process?

A1. In my view, we need to move beyond information sharing as the answer to addressing cybersecurity. Along with "public-private partnerships", information sharing has been called out as the solution to cyber security for the last two decades. The idea is that once companies and individuals are informed about threats and vulnerabilities, they will be armed with the information they need to improve security. That was a good theory but it is one that has turned out to be proven wrong by the facts. Information sharing is in fact quite good in cybersecurity. At last count, there were more than thirty partnerships between the Federal Government and the private sector to share information on cyber security. The National Institute of Standards has done a excellent job of providing standard nomenclatures for policy makers and practitioners. Efforts such as the National Vulnerability Database and the Common Vulnerabilities and Exposures naming standard provide the technical means for exchanging information. Information sharing is good. It is getting better. We now need to take a hard look at why better information sharing hasn't led to better cybersecurity and then develop remedies.

Q2. Many of you have discussed the need for new internet protocols to be built on the concepts of security, authentication, and attribution. What parties would help develop and implement these protocols and what would their roles be? Who would use these new protocols and would multiple protocols diminish the utility of the internet?

A2. I believe that the current iterative, consensus-based process through the Internet Engineering Task Force for the development of protocols is broken. By way of example, look at DNSSEC. The security flaws in the Domain Name System (DNS) that DNSSEC is designed to address were first discovered in 1990. It took another decade to develop the first specification for DNSSEC. In 2010, we are just taking the first meaningful steps to implement the solution and it will likely take another decade for widespread adoption. In my view, government needs to set the goals, fund the research, and then require implementation. The argument that the pace of innovation is too fast for government regulators to keep up with is patently untrue given the thirty-year timeframe to develop and implement DNSSEC. I believe that the U.S. government should layout a technical challenge to the IETF on a strict timeframe to develop a secure suite of protocols, fund the development, and require implementation.

Q3. Please discuss how the level of confidence can have an impact on the utility of attack attribution. Please relate the level of confidence to the spectrum of available responses including diplomatic, economic, cyber, and kinetic.

A3. With existing technologies, we can have a high degree of confidence in our ability to trace an attack back to a system. The difficulty is in determining both the originating system and the human at the keyboard. In almost every conceivable cyber attack, we will be able to trace the attack back to at least the first system and then ask the host country for assistance with further investigation. If they refuse, we can say with confidence that they are uncooperative and assign them responsibility. Ultimately, attribution back to the originator of the attack may take time, particularly for the President and Congress to authorize diplomatic, economic or kinetic responses outside the cyber domain; however, as in our response to the terrorist attacks of 9/11, we may respond "at a time of our choosing", once we have enough confidence to act.

Q4. Are there any other thoughts or issues you would like the share with the Committee on attack attribution and cybersecurity?

A4. Not at this time.

Questions submitted by Vice Chair Ben R. Luján

Q1. The Fourth-generation of cellular wireless network standards being developed uses the internet protocol suite and would extend the internet to cellular devices. What are the implications of this 4G standard for this discussion on privacy and attribution?

A1. I am not familiar enough with this issue to provide a meaningful response.

ANSWERS TO POST-HEARING QUESTIONS

Responses by Mr. Ed Giorgio, President and Co-Founder, Ponte Technologies

Questions submitted by Chairman David Wu

Q1. Information sharing is critical for success in cybersecurity, whether it supports attribution of attacks or awareness of vulnerabilities. How important is it to have common nomenclature, common metrics, and standard sharing methods for success in information sharing? How should these different elements be developed, which government agencies should be involved, and what roles should they play throughout the process?

A1. Common nomenclature and metrics are extremely important to move the current state forward. Standards have been very difficult to achieve in this area due to the vested interests of the private security service companies who want to develop these standards as their individual intellectual property and only make them open source after they have achieved sufficient market penetration. In some cases these private companies have no interest in standards at all because they don't want their systems to easily interoperate with competitor systems as that might cause them to eventually be marginalized. This resistance can be overcome by government activities such as the Security Content Automation Protocol (SCAP) currently underway by NIST, NSA, and others.

SCAP details can be found on the NIST web site. In short, SCAP is a synthesis of interoperable specifications derived from community ideas and is initially focused on vulnerability management. Subsequent activity will expand to include compliance, remediation, and network monitoring. Existing SCAP standards include Common Configuration Enumeration (CCE) , Common Vulnerabilities and Exposures (CVE), Open Vulnerability and Assessment Language (OVAL), Common Vulnerability Scoring System (CVSS) and others.

Q2. Many of you have discussed the need for new internet protocols to be built on the concepts of security, authentication, and attribution. What parties would help develop and implement these protocols and what would their roles be? Who would use these new protocols and would multiple protocols diminish the utility of the internet?

A2. As mentioned in my testimony, government cannot by itself mandate changes in underlying infrastructure technologies (Ex. IPv6). DARPA, NSA, NSF, and the research elements supported by the Comprehensive National Cyber Initiative all should be working to research and develop new capabilities. These could be researched, designed, implemented, piloted, and ultimately become operational on DoD and Intelligence networks, where attack attribution is far more important.

New protocols based on the above research should be introduced through the IETF, as this process is the most likely to encourage commercial acceptance and deployment into worldwide networks. For security standards or algorithms, NIST is the appropriate agency.

As for using multiple protocols, we've done this for decades with considerable success. The challenge is to make sure that different protocols complement each other rather than cause uncertainly, confusion, and even counter productivity. The way to reduce this risk is to make sure the standards development processes are not done in isolation as has frequently happened in the past.

Q3. Please discuss how the level of confidence can have an impact on the utility of attack attribution. Please relate the level of confidence to the spectrum of available responses including diplomatic, economic, cyber, and kinetic.

A3. If we have a legally meaningful level of confidence in attack attribution then the utility of this goes beyond mere attribution, as some would-be attackers will be deterred by the ramifications of that attribution. We should have fine-grained control over what level of identification and authentication we require before access is granted. This in turn will give us control over the level of confidence we have in attribution. Perhaps for a low value target we would just accept that it's going to be attacked and not bother so much with attribution.

The level of confidence one can have using attack attribution technologies varies dependent on the:

1. Type of hardware the attack is emanating from,
2. Specific operating system and application software in use,
3. Level of user authentication used on that system,

4. Internet protocols, including security protocols such as IPSEC, and
5. Cooperation from the Internet Service Providers (ISPs)

If the identity of the individual is required, that is harder than just knowing the machine from which the attack is emanating, and that, in turn, is much harder than knowing the geo-location of the that machine. As mentioned in my testimony, trying to pinpoint the exact individual who is willfully committing the attack cannot be done with a high level of confidence due to problems with the security on the system the attack is emanating from.

Consideration of all the above attributes will be required to obtain a level of confidence suitable for the appropriate diplomatic, economic, cyber, and kinetic response. A *diplomatic* response such as a formal state department demarche does not appear to be much of a deterrent at all, as countries like China and Russia will simply deny it. *Economic* responses could be very valuable, but will require an international approach which does not impinge on the individual nation state sovereignty. *Cyber* responses are certainly unclear as to their effectiveness, especially since the U.S. is the most dependent on cyber and has the most to lose in a cyber conflict. Finally, a *kinetic* response of course escalates any cyber attack to a much higher level conflict and cannot be done without absolute certainty of where the attack is coming from. Even then, I doubt there would be much national or international support for such an action and this response should be avoided.

Lastly, in answering this question, it is important that research & development be done in all the five areas listed above as advances in these areas will both stop some attacks and deter others. DARPA, NSF, NIST, and NSA all have a role in accomplishing this.

Questions submitted by Vice Chair Ben R. Luján

Q1. The Fourth-generation of cellular wireless network standards being developed uses the internet protocol suite and would extend the internet to cellular devices. What are the implications of this 4G standard for this discussion on privacy and attribution?

A1. There has been an explosive growth in the availability of location databases that associate building and emitter identifiers (IDs) with geographic coordinates. While these capabilities are assisting in solving the attribution problem, they are also enhancing criminal activity and adversely impacting our personal privacy and national security. This is especially troublesome since the data is (primarily) in the hands of private and frequently multinational corporations.

Examples of these data bases include information about 4G cell phones & PDAs, IP addresses, WiFi and WiMax emitters, cell towers, routers, gateways/points of presence, physical addresses, among others. Additional clues to location can be derived from the above plus timing calculations and measurements within data and voice traffic.

These data bases exist in many different forms today and are perpetually updated, some in real-time. Furthermore, these data bases are held in the hands of multiple distinct parties, including:

1. Classified government data bases
2. Private commercial data bases (e.g., cell phone, PSTN, ISP, and utilities),
3. Open-source data bases (e.g., Internet registrars, Google Maps),
4. Unclassified (but sensitive) government data bases, and
5. Foreign government or foreign corporate data bases.

For example, the above data bases can be correlated and combined to discern coordinates for various scenarios, such as tracking individuals in real-time by overlaying their current position on a satellite image or street view to follow their every movement and make notes of where they went, at what time, who they met with, who they emailed or phoned, what they purchased, and so on. As mentioned in my testimony, these capabilities pose both an opportunity to do attribution when we need it, but a potentially catastrophic vulnerability when it is used for foreign cyber attacks, corporate espionage, criminal activity, and, potentially, terrorism.

ANSWERS TO POST-HEARING QUESTIONS

Responses by Mr. Marc Rotenberg, President, Electronic Privacy Information Center

Questions submitted by Chairman David Wu

Q1. Information sharing is critical for success in cyber security, whether it supports attribution of attacks of awareness of vulnerabilities. How important is it to have common nomenclature, common metrics, and standard sharing methods for success in information sharing? How should these different elements be developed, which government agencies should be involved, and what roles should they play throughout the process?

A1. There are technical standards that enable data exchanges but it is critically important to keep in mind that there are also legal standards that help ensure trust and confidence in the collection and use of personal information by the Federal Government. This problem is already clear in the use of "cookies," i.e. persistent identifiers, by government agencies in the management of Federal web sites.

The Federal Privacy Act sets out a framework for all Federal Government agencies collecting and using the personal information of American citizens. That framework embodies a set of principles that any new Federal attribution system is bound to adopt. The Privacy Act limits most agencies to maintain records of individuals only which are "relevant and necessary" to accomplish specific purposes derived from statute or executive order.

More generally, the framework prioritizes the individual citizen's right to request and view all government records about him or her that do fall under a set of specific statutory exemptions, and for that citizen to sue the government for violations of the statute.

Clearly, there is a need to strengthen the application of Privacy Act across the Federal Government. The original draft bill considered by Congress contemplated an independent Federal privacy agency to oversee enforcement of the Act. We would still favor this approach. Short of new legislation, the OMB should play a more active role ensuring compliance with Privacy Act provisions.

Q2. Many of you have discussed the need for new internet protocols to be built on the concepts of security, authentication, and attribution. What parties would help develop and implement these protocols and what would their roles be? Who would use these new protocols and would multiple protocols diminish the utility of the internet?

A2. The ideal security model for new Internet protocols should focus on end-to-end encryption and dynamic addressing instead of attribution and surveillance. End-to-end encryption translates data into a secret code, thereby protecting it from the moment it leaves the sender computer until the moment it is received by the intended recipient computer (and decoded). This kind of comprehensive encryption is essential for protecting personal data that travels over vulnerable channels, such as the public Internet.

Dynamic addressing serves a similar purpose in a different way. The term refers to Internet Protocol (IP) addresses, which computers use to direct bits of data across the web. There are two ways to assign IP addresses. A dynamic addressing system assigns each computer a random selection from a preselected pool of addresses. A static addressing system assigns each computer a single, permanent address. The latter is based on the same philosophy as attribution systems, and shares its inherent flaws.

The most recent version of widely used Internet Protocols is IP version 6 ("IP v.6"). IP v. 6 enables, but does not require, network administrators, IT professionals who run individual networks for companies and other large organizations, to use static addressing. This could create new risks to users. Permanently tracing personally identifiable online conduct to individual users serves to provide hackers additional targets. Alternative protocols can take advantage of IPv6 functionality while minimizing the privacy risk.

There are numerous organizations that can assist in developing and implementing protocols that reflect a more resilient, open approach to internet security that rely on end-to-end encryption and dynamic addressing. I would recommend the Internet Engineering Task Force, the Internet Architecture Board, and the Internet Corporation for Assigned Names and Numbers (ICAAN).

Q3. Please discuss how the level of confidence can have an impact on the utility of attack attribution. Please relate the level of confidence to the spectrum of available responses including diplomatic, economic, cyber, and kinetic.

A3. Attribution programs do not prevent highly skilled attackers from remaining anonymous. They do create vulnerable repositories of personally identifiable information, but only for those Internet users who are not trained in frustrating attribution systems. In fact, these repositories would soon become tempting new targets for the hackers who are outside the attribution system.

Furthermore, the National Academy report that I cited in my testimony said, "It is not known how much the smooth operation of society depends on such things, or on the assumption that they are possible. There is a risk, however, that they would be lost, or at least significantly impaired, if a broadly used nationwide identity system came into existence."

Again, current schemes of attribution are inherently limited, which significantly diminishes the levels of confidence we can invest in them. Still, one useful mechanism of attribution is called Domain Name System Security Extensions, or DNSSEC. DNSSEC reduces the risk of phishing by focusing attribution efforts on authenticating websites. That is a distinctly different approach than tracking individual users, and in 2008, the Electronic Privacy Information Center endorsed this approach in administrative comments relating to ICANN's adoption of DNSSEC for websites ending in ".org" (the .ORG Domain).

"Phishing" is a hacker term for malicious websites that pose as legitimate ones to fraudulently acquire sensitive information about Internet users. The primary mechanism DNSSEC uses to prevent phishing is a new form of authentication built into the Domain Name System. The Domain Name System translates the computer language identifiers for Internet addresses into language human users understand. DNSSEC adds a level of security to this process by requiring sites to use digital signatures. Digital signatures are mathematical messages which allow the users' computer to discern whether or not the site is the one it claims to be or instead a fraudulent intruder.

Beyond bounded approaches like DNSSEC, the Federal Government probably not design diplomatic, economic, cyber, and kinetic approaches to foreign policy around the attribution systems currently available. They are not very reliable, and suffers from the limitations I've described in my testimony and in response to questions.

Q4. Are there any other thoughts or issues you would like to share with the Committee on attack attribution and cybersecurity?

A4. Cyber security is a transnational problem that requires resilient solutions. The primary function of a national attribution system, in the abstract, would aim to solve more problems than it creates by extending the range of our country's foreign policy tools and domestic policing techniques. In practice, however, available systems can yield ambiguous results at best, which will frustrate security efforts instead of bolstering them.

Moreover, there are fundamental privacy rights at stake. Building the capacity to track American citizens has always been two-edged. Large scale, preventative surveillance invites abuse. In this case, it invites the malicious users we are fighting to participate in the abuse. Cyber attackers can operate outside of any available attribution system, and use our system against us.

Invariably, solving one problem in the cyber security field will create a new problem. A smart strategy must anticipate this dynamic.

Questions submitted by Vice Chair Ben R. Luján

Q1. The Fourth-generation of cellular wireless network standards being developed uses the internet protocol suite and would extend the internet to cellular devices. What are the implications of this 4G standard for this discussion on privacy and attribution?

A1. As mobile phone companies such as Verizon and AT&T Mobility transition to the 4G wireless standard, there is the possibility that the "Internet of things"—familiar communications devices, such as cell phones, as well as many objects, such a refrigerators, identity cards, and clothing—will become uniquely identifiable and locatable.

Some may favor this capability because it will make possible new forms of real-time attribution. But for the determined attackers, it will also create new opportunities to conceal identity and to turn the techniques of attribution against us. Robust security systems should not rely on the perfectibility of attribution.

○